EMYS: a social robot that plays “Sueca”

Filipa Correia

Instituto Superior Técnico
Universidade de Lisboa

Abstract. Regarding the needs of elder population, technology can be the answer to some of them. Existing robots already guide the elderly walking through the house, or are companions for cuddling and petting. In the same way, this work aims to create a companion robot that plays the card game *Sueca*. The idea is to develop an entertaining and pleasuring environment that can, additionally, stimulate their reasoning. After reviewing the related work, this report proposes an artificial player based on Monte-Carlo methods and the architecture that will use it. Considering the game state, the chosen move, and some of the environment perceptions, this robot must produce an appropriate behaviour to be considered socially present during the game.

Keywords: Artificial Intelligence, Trick-taking Card Game, Hidden Information, Interactive Companions, Socially Intelligent Behaviour

1 Introduction

The world population is ageing dramatically and as such, the society needs to change and embrace this problem in many different ways. The elderly have specific needs, physical and cognitive, that are often not considered in the way we, as a society, organise our lives. Some of these concerns are recently being solved with the help of technology and may range from computer programs to intelligent robots. For instance, elderly with mobility disabilities may be guided by a robot while walking through the house [28]. Another example is software for stimulating and training memory problems of Aphasia, a language disorder caused by brain damage [29].

However, existing technology with elderly purposes is commonly focused on health care. When dealing with aged people with no serious health problems and that are still capable of doing their regular daily tasks, there is still a need to occupy their free time with pleasuring activities. Finding appropriate tasks supported by technology may include the training of their cognitive functions or just accompany them. In order to join these two purposes of accompany and train their reasoning, an artificial game player could be a suitable solution. Several embodied agents for game playing exist, such as the iCat chess tutor [19] and the EMotive headY System (EMYS) Risk player [27].These examples have inspired the idea of exploring a card game scenario, that can exemplify an activity that aged people enjoy doing.

Overall, a card game is an entertainment activity that aged people are used to do and, at the same time, might help them training their cognitive functions. As a result, considering some of existing card games are still unsolved challenges for Artificial Intelligence (AI), the goal of this project is to integrate a social robot with aged humans in a card game scenario. This project is a great opportunity to relate all the concerns mentioned above. A game-playing companion for elder people must, at the same time, (i) be able to play competently the card game; and (ii) interact socially. The first skill requires such agent to include an AI module that is able to reason strategically about the game. The second skill requires an emotional/social module that enables the agent to behave in a manner that is socially believable.

Computer programs that play games have been an interesting challenge for AI. From board to card games, or even role-playing games, the goal is to create rational agents capable of evaluating the game and achieving the best outcome. Deep Blue, Chinook and Watson are good examples that have raised the bar for developing this kind of agents. Deep Blue is a remarkable chess player and has defeated the human world champion in 1997 [8]. Schaeffer et al. have solved Checkers with Chinook program and proved the game leads to a draw with two optimal players [34]. Lastly, Watson is the Question Answering (QA) system that has beat the two highest ranked Jeopardy players in 2011 [11]. All these agents are good baselines to improve AI in games.

Besides building programs that try to think rationally or humanly, AI has also another branch that aims to act humanly [33]. This concern arises from the inclusion of robots in humans' life and influences the way they interact and communicate with people. Consequently, robots have to behave properly in those environments, considering they are surrounded by humans. The Human-Robot Interaction (HRI) field explores these concerns of integrating robots with humans in a social environment. Since this field descends from Human-Computer Interaction (HCI), it also inherits the user centred development that establishes the prior need of doing user studies.

Lastly, the chosen card game is *Sueca*, a well known Portuguese game among the ageing population. It is a four-player game with two teams and involves an opponent and a companion role for the agent. Regarding HRI concerns, these two roles together have not been studied yet in an artificial embodied game player.

1.1 Goals

The main goals of this project are:

- To develop a robotic agent capable of playing competently the *Sueca* card game;
- To include social behaviours on an embodied agent in order to act according to the game state;
- To evaluate the correctness and advantages of the proposed system.

*

The next section presents some background research that helps the reader understand the problems further mentioned (Section 2). The report proceeds with the state-of-art of playing card-games and human-robot-interaction with the elderly (Section 3). Additionally, it reveals a pilot user study (Section 4). Finally, it presents the architecture, its evaluation methodology and the final conclusions (respectively, Sections 5, 6 and 7).

2 Background

The current section introduces the discussion of relevant research in order to understand some concepts and terminology further mentioned.

2.1 Game theory concepts

Game theory studies decision making problems involving multiple decision makers. A problem of this nature is usually called a game and defines a set of constraints to the players' actions. It also studies the strategies these players might take and the properties of each game.

Each decision maker tries to maximise the payoff/reward of his possible actions and one possible approach to do that is to consider the opponents' actions. The Nash-equilibrium [24] of a game is a stable strategy for every player and occurs when each player chooses the best strategy for himself, considering their opponents have the same behaviour. Moreover, each player cannot have a better benefit by changing his strategy unilaterally.

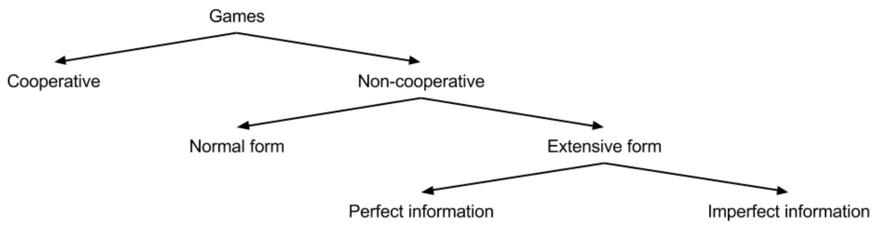

Fig. 1: Hierarchy of games

Figure 1 shows how games can be hierarchically categorised [25]. In a cooperative game, players cooperate with one another in order to achieve a common goal. Alternatively, in a non-cooperative game, each player works independently for its own purposes. Non-cooperative games can also be branched into two forms: normal and extensive form [35]. A normal form game can be defined as the tuple $(N, (A_k)_{k=1}^N, (u_k)_{k=1}^N)$, where:

- N is the number of players;
- A_k is the finite set of available actions for the k-th player;
- u_k is the payoff for the k-th player.

Additionally, considering the players' payoffs, another relevant concept is the zero-sum game, where the sum of all players' payoffs is zero. For instance, in a zero-sum 2 players game, $u_1 = -u_2$. Although the normal form games assume that players' actions are made simultaneously, in the extensive form games, the players' actions are sequential. This evidence leads to another branching in the hierarchy of games and, consequently, an extensive game can be considered as a perfect information and an imperfect information game. In perfect information games, each player knows exactly the real state of his opponents, (e.g. Chess). In imperfect information games, the game state is not fully observable to the players. For instance, in a Poker game, a player only knows its own cards, the cards in the table, and the bets of all players.

In imperfect information games, an *information state* or *information set* for a player k corresponds to a set of all games states that yield the same observation to player k. For example, in a Poker game, an information set consists in all game-states that lead to the same observed cards in the table and in the players hand.

2.2 The game of *Sueca*

Sueca is a card game categorised as trick-taking, which means the game has a finite number of rounds, called tricks. In this case, there are ten tricks, since the deck has forty cards equally distributed among the four players. This game uses the standard French card deck, excluding the rank 8 through 10. Although most trick-taking card games count the number of winning tricks to determine the winner, *Sueca* assigns points to the cards, according to Table 1. The most significant difference, compared to other games, is the card with rank 7 being higher than the King (K) and lower Ace (A).

All valued cards sum 120 points, which means a team with more than 60 points wins the game. Moreover, each player is paired with the player in front of him, and the two adjacent players form the opposing team. Hence, the game involves both cooperation and competition.

Table 1: Rank of cards per suit and respective reward values

Cards	2-6	Q	J	K	7	A
Points	0	2	3	4	10	11

After the deck has been shuffled and divided, the dealer chooses the top or bottom card to be the trump suit, leaves it on the table, and distributes the remaining cards among all players. The remaining rules are quite similar to any other trick-taking games:

- Follow the suit of the first card played in the turn (lead suit), if possible;
- A player wins the trick if his card has the highest value belonging to the lead suit or the trump suit.

Sueca is a nondeterministic game, since it includes what is called the element of chance by the cards being dealt randomly at the beginning. Additionally, since the cards of each player are hidden from the other players, this is considered as an imperfect information game. There are almost 1.9×10^{22} possible card distributions[1].

2.3 Monte-Carlo Tree Search

Monte-Carlo Tree Search (MCTS) is a family of algorithms whose goal is finding optimal decisions, by combining standard tree search with Monte-Carlo sampling [6]. This method incrementally builds a search tree according to the results of previous iterations. The search tree is expanded by randomly sampling the nodes. Usually, it is divided into four steps, as described below.

- Selection: To select a child node through a selection policy. This policy must balance between unexplored areas of the tree and promising nodes that may lead to higher rewards.
- Expansion: To expand the selected node to add one or more nodes to the tree, according to the available actions.
- Simulation: To select an expanded node through a simulation or default policy to produce an outcome.
- Backpropagation: To propagate the reward value of all the selected nodes in order to update their statistics.

The main strength of MCTS is its little knowledge requirement, and so, it can be applied to many different domains. This algorithm can also be easily paralleled, since each simulation process can be done independently.

According to Browne et al., finding a suitable variation of MCTS is the greatest challenge of applying the algorithm to a specific environment. The most popular algorithm of MCTS family is Upper Confidence Bounds for Trees (UCT), and this variation differs from the original in the selection phase. It uses a maximisation function to evaluate the available nodes, according to the following equation:

$$UCT = \overline{X}_j + c\sqrt{\frac{\ln n}{n_j}}, \tag{1}$$

where $\overline{X}_j$ represents the average value associated with option j in the current state, n_j is the number of times that option j was selected in the current state, n is the number of visits to the current state and c is an exploration parameter.

It models the reward of each child node as an independent multiarmed bandit problem. This means that, besides considering the already obtained average

[1] $4 \times {}^{40}C_{10} \times {}^{30}C_{10} \times {}^{20}C_{10}$

reward of a child node ($\overline{X}_j$), it also considers the maximum expected gain of that same child node ($c\sqrt{\frac{\ln n}{n_j}}$). This function establishes an equilibrium between exploitation and exploration. Exploitation is evaluated in the first term of the equation by considering the average reward of a child node. The exploration is manipulated in the second term by balancing the amount of times the parent node and the child node j have been visited (respectively, n and n_j). A considerable amount of iterations will approximate the UCT to a minimax tree. Consequently, the produced results are nearly optimal with high probability.

3 Related Work

This section presents the state of the art related to this work. Since no relevant studies on *Sueca* have been found, the research is focused on algorithms used in similar card games. It also presents existing companion robots that will allow the analysis on human-robot interaction relevant for this work.

3.1 AI in games

AI has been solving many games over the years, however, the definition of "games" usually refers to zero-sum and perfect information games. These kind of games are commonly solved by creating a tree representing all possible states and searching for the optimal, or a nearly optimal, solution. The greatest achievements related to perfect information games are generally based on finding good heuristics to refine the search and also good prunings to reduce the search space. Deep Blue can exemplify this idea [8], it uses an iterative-deepening alpha-beta search and the key of its success is mostly the null move heuristic and the futility pruning. Another example is Chinook, also a perfect information game that was solved using alpha-beta search [34].

Nevertheless, *Sueca* is considered an imperfect information game, as described in Section 2, and this class of games is usually solved by one of three different approaches [10]. The first one, and the most popular, is based on Monte-Carlo Methods. Then, another possible approach is trying to compute a Nash equilibrium strategy or an approximation thereof. Lastly, belief distributions involving game state inferences and opponent models can also be used. The first two mentioned approaches are mutually exclusive, while the last one can be used as a supplement. The following two subsections detail how Monte-Carlo Methods and Belief Distributions can be applied in hidden information games. The second pointed approach will not be addressed due to the imposed limitations of our domain, considering, for instance, that the maximum known number of states for computing a Nash equilibrium is 10^{12} [42], which is much lower than the number of possible states in a *Sueca* game.

Monte-Carlo Methods

The popularity and acceptance of Monte-Carlo based methods have increased since its success on Bridge. Ginsberg's Intelligent Bridgeplayer (GIB)[2] was the first computer bridge champion using Monte-Carlo Methods, and subsequently, another two successful domains were Skat[3] and Computer Go [15]. Since some of these domains remain a challenge for traditional AI techniques, this method seems to be very promising.

In order to solve a hidden information game, the first challenge is to deal with information sets. The most used approach to solve it is determinization, which samples choice nodes instead of considering all of them in an unique set. Applying this approach to MCTS is known as Perfect Information Monte-Carlo (PIMC). For instance, in a card game scenario, each iteration of PIMC samples the cards distributions for all players and the simulation process of the game behaves as a perfect information game. In other words, during the simulation each player makes decisions as if his opponents' cards are visible. The first successful implementation of this technique was GIB [16].

In 1998, Frank & Basin produced an analysis on PIMC's limitations [12]. They identified two distinct problems: *strategy fusion* and *non-locality*. Due to the repeated minimaxing architecture that PIMC has and its evaluation of possible distributions with the best strategy, applying this knowledge, when information is missing, might produce suboptimal decisions. This is called the *strategy fusion*. For instance, when having a move with a guaranteed reward and another move with a possible reward of the same value although depending on the current world, PIMC equally considers both moves.

The second problem, *non-locality*, results from the propagation of values. The value of a game tree node only considers its children' values, however, in an imperfect information game, some guesses might be done using values of the non-local subtree. For instance, considering 2 different worlds, the player 1 can guarantee a winning trick in the world 1 by making a certain move, and if in that state, he makes another move instead, player 2 might assume they are in world 2. PIMC cannot make such an inference.

Despite the satisfying outcomes of PIMC, there were still difficulties in understanding the strong results of this algorithm. As such, Long et al. have analysed the previously mentioned problems of PIMC search, and they have shown how three different properties of a game can influence the success of PIMC [23]. The first property is *leaf correlation*, which refers to how likely it is to affect a player's payoff in the neighbourhood of a leaf. When the probability of all siblings having the same payoff values is higher, the correlation value increases. Secondly, *bias* indicates the chance of a player being preferred over another. Finally, the last game characteristic that has been pointed is *disambiguation factor*, that denotes how rapidly the hidden information is revealed.

These properties have been tested in a set of experiments in both PIMC and a random player against an optimal Nash equilibrium player. Results shown

[2] http://www.gibware.com/
[3] https://skatgame.net/

the performance of PIMC increases as the correlation value is higher, bias does not considerably affect its success, and, finally, disambiguation has the greatest impact on the results of the algorithm. When this last value is higher, it means the game turns more quickly into a perfect information game. Additionally, the authors demonstrate these properties on real game examples, such as Skat and Kuhn poker. Skat indicates a considerably good performance of PIMC, due to its values of *leaf correlation*, *bias*, and *disambiguation factor*. Since Skat presents strong similarities to *Sueca*, it is expected that PIMC also has a good performance when applied to *Sueca*.

Cowling et al. have also investigated the application of MCTS to hidden information games [10]. Their research supports a new descendant family of algorithms, Information Set Monte-Carlo Tree Search (ISMCTS). ISMCTS works with information sets, instead of game states and uses determinization to sample the game, however producing a single tree. The main advantages are the computational budget efficiency and the fact of suffering less from *strategy fusion* than PIMC. The authors also presented some experiments in three different games, including a card game. Their results on the card game Dou Di Zhu were very similar to UCT and did not introduce any improvement to the playing strength. The authors explained these results with the high branching factor this domain produces, which has discouraged the usage of this technique on the domain of *Sueca*, since in the information set tree, the initial branching factor would also be high (10^8, ${}^{40}C_{10}$).

Recently, Furtak & Buro [14] presented a new search algorithm called Imperfect Information Monte-Carlo (IIMC) that can be suitably applied to hidden information games and reduces the *strategy fusion* problem. During the simulation phase, each player's move is chosen inside a player's module and the game behaves as an imperfect information due to this encapsulation. Additionally, the players' modules allow the differentiation of players using different strategies. The authors revealed the great potential of this approach when applied to trick-based card games, considering it has been successfully tested in the Skat scenario.

Table 2: Advantages and disadvantages of the mentioned Monte-Carlo algorithms

Algorithm	**Advantages**	**Disadvantages**
PIMC	Offline Computation Easy to parallel	Strategy fusion Non-locality
ISMCTS	Offline Computation Easy to parallel Computational budget	Strategy fusion (less than PIMC) Non-locality Complexity
IIMC	Offline Computation Easy to parallel Allow a different player model per player	Strategy fusion (less than PIMC) Non-locality Complexity

The advantages and disadvantages of the mentioned MCTS variations are clearly summarised in Table 2. Both of three techniques are easy to parallel and allow an offline computation. However, ISMCTS uses the computational budget more efficiently than the other two techniques, and IIMC allows different player models per player. Disadvantages show all the three techniques have the *non-locality* and *strategy fusion* problem, although *strategy fusion* is lower in ISMCTS and IIMC. Another relevant disadvantage is the computational burden that ISMCTS and IIMC add, when compared to PIMC.

Game State Inference & Opponent Modelling

While discussing imperfect information games, belief distributions, game state inference and opponent modelling are another relevant subjects to consider. Predicting some of the opponents' cards or other clues would be beneficial to select better actions at each state of the game. Additionally, inferring hidden information, while using a Monte-Carlo based method, can also decrease the *non-locality* problem [10].

Buro in 2009 [7] presented his work on state evaluation and inference that has been included in his Skat player. His approach combines two techniques, one for evaluating the bidding and another for selecting hypothetical worlds during the game play. The former technique uses a logistic regression to evaluate the winning probability of each hand and it has 22 million Skat games as data base. This winning probability determines the strength of a hand and can, therefore, be used on the bidding.

The second technique is mainly based on two heuristics. Fastest-cut-first search heuristic evaluates each move according to its beta-cutoff value and minimises the expected number of visited nodes. Additionally, in order to reduce the tree exploration, another heuristic groups cards by their strength value and considers, for example, 7♣ and 8♣ the same move, when holding both cards in a player's hand. The author compares his work to other similar ones and concludes the strength of his techniques lies in two central points. First, determining the $P(\text{world}|\text{move})$ on offline data, instead of doing it in runtime. Sencond, his formulation is generalised in a way that it is possible to perform it on high-level features. Since the main difference between *Sueca* and Skat is that the first one does not have the bidding phase, Buro's first technique would not be appropriate for the *Sueca* game. However, the search enhancements could be suitably applied, considering the game trees are identical.

Usually, opponent modelling uses optimal strategies to predict the other players' actions and these models tend to be overly defensive. Consequently, Long & Buro in 2011 [22] suggested a post-processing analysis that is able to infer opponent's qualities based on their decisions in a certain environment. The main idea is to classify each opponent with a mistake rate and use that value to be more or less defensive. This approach, called Perfect Information Post-Mortem Analysis (PIPMA), computes a procedure after each game episode (in a trick-taking card game, it would be after each trick) to incrementally update the mistake rate of each opponent. The authors made some experiments in a Skat player with very good results, where they used the mistake rate to adjust the bidding behaviour during the game. Despite the fact that *Sueca* does not have the bidding phase, classifying opponents with a mistake rate can useful to other purposes. As a result, it would be interesting to model the opponents in a similar way in the domain of *Sueca*, in order to make better decisions or even for the embodied agent to produce adequate behaviours.

Another highly suitable card game to make opponent models is Poker, since predicting the players' moves can naturally affect the outcome of this game. In order to predict players' cards and their future actions, Posen et al. in 2010 [30] have investigated this subject. They proposed an opponent model that starts with a prior distribution and changes over time with a differentiating function. The prior distribution allows it to make reasonable inferences while having insufficient information. In addition, the relational probability tree algorithm TILDE builds a decision tree with the stored samples of a player. This decision tree represents the differentiating function that will adapt the initial prior distribution. Besides this opponent model, the authors explain how to integrate this function with MCTS. Instead of sampling the cards randomly, MCTS uses card predictions and, therefore, the algorithm does not need a numerous amount of iterations to reach a uniform card distribution. Furthermore, the probabilities of action predictions are used in the selection phase of the MCTS, according to the state of the game and the sampled cards. Since MCTS can be used in the *Sueca* domain, a similar opponent model can also improve the capabilities of this algorithm, as shown in Poker.

Table 3: Techniques signed with ×, ✓ and ∼ symbols are, respectively, not suitable, suitable and conditionally suitable to the *Sueca* domain.

Tested domain	**Technique**	**Goal**	**Suitable to *Sueca***
Skat	Determine the winning probability of a hand	Improve the bidding	×
	Fastes-cut-first heuristic	Order moves	✓
	Considering similar states equally	Reduce tree exploration	✓
	Calculate the mistake rate of each player	Improve the bidding	∼
Poker	Opponent model	Improve MCTS policies	✓

Table 3 summarises what techniques have been reviewed, their purposes, and, finally, if they can be applied to the *Sueca* domain. The technique of determining the winning probability cannot be used for the exact same purpose, since our domain does not include a bidding phase. The next two search enhancements can naturally be used due to the similarities between Skat and our domain game trees. The mistake rate was signed as conditionally suitable because it can also be used, although with a different purpose. Thinking in the embodied agent of our work, it can assign a mistake rate variable to each player and produce appropriate behaviours according to their values. A similar approach might be thought to use the winning probability, however, opponents' hands are not visible and the agent should not reveal its own information. The last technique also can be an addition to the MCTS base policies.

3.2 Human-robot interaction

Regarding the goals of this project, it is crucial to investigate and evaluate the state of the art of HRI, in particular in the context of robot companions or players. Since it aims to interact with aged people in a card game scenario, there are two clear branches that must be studied. Firstly, the existing robots with an elderly care purpose. Secondly, how social agents have been integrated into games. The next subsections will address these points.

Robots in elderly care

The greying of population is an undeniable demographic fact and, consequently, assisting the elderly in their daily living is a worrying subject. In order to address this concern, robots can be a valuable aid, however, considering the limitation of current robotic technology, their purposes are present in more specific tasks.

In 2009, Broekens et al. analysed and reviewed the most relevant literature about social robots in elderly care [5]. The authors categorised assistive robots for elderly as shown in Figure 2. The first division distinguishes social robots from nonsocial robots. The nonsocial ones are used for rehabilitation purposes and physical assistance, such as a smart wheelchair or an artificial limb, however, regarding the main purposes of this work, nonsocial robots will not be discussed. Social robots should be perceived as social entities due to their interaction with humans and can also be divided into two different sets, service type and companion type. The intersection of these two sets represents some of the robots that are used for both purposes and cannot be strictly categorised.

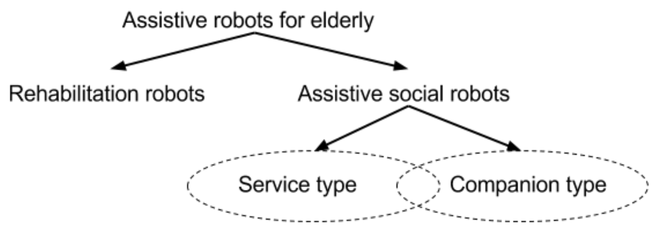

Fig. 2: Categorization of assistive robots for elderly

A well known social service robot is *Pearl* (Figure 3a), developed in the Carnegie Mellon University within the Nursebot Project [28]. This autonomous robot's duties are to guide the elderly through their environment, and to remind them about their daily activities, such as eating or taking their medicine. In other words, this functional assistant is capable of giving advice and providing cognitive support. When analysing *Pearl* through a more general AI point of view, this robot is equipped with many different technologies. Firstly, it has a speech recognition module and also has speech synthesis. Secondly, it has stereo camera systems and performs a fast image processing including face recognition. Lastly, *Pearl* also provides a navigation system and its body is touch sensitive.

Another two similar service robots are *RoboCare* [2] and *Care-O-bot II* [17]. They both are autonomous and provide indoor guidance to the elderly and, due to their advanced domotic components, strong planning, and scheduling frameworks, they can improve the independence of their owners. Since the aid these service type robots may grant to the elderly covers most of their daily basic activities, the involved concerns are amplified when compared to the proposed robot that plays a card game. These worries are reflected, for instance, in the extensive amount of sensors these robots should include.

(a) Pearl (b) Paro (c) Aibo

Fig. 3: Service and companion robots for the elderly.

Paro is a seal shaped companion robot used as medical therapy for the elderly (Figure 3b). Since 2003, the work by Wada et al. provides a very good psychological and physiological evaluation of *Paro*'s effects on the residents of a care house [39–41]. This robot contains a behaviour generation system that provides proactive, reactive and physiological reactions, such as, poses or motions, looking at the direction of a sound, and sleeping. Their studies of both three weeks and one year have shown improvements in residents' moods, depression, stress levels, and social interactions with other residents. The goal of such a robot is fully inspired in animal-assisted treatments, which have studied benefits in humans' health. However, hospitals and health centres do not allow animals due to hygienic and safety reasons. Hence, researchers found a great opportunity to build similar robotic animals.

Another example of a purely companion robot is the *Huggable* [37], a teddy bear shaped covered of extremely sensitive touch sensors. The *Huggable* not only detects hard and soft touches, but also distinguishes between an object and a human touch. Considering experiments in an hospital, this robot was connected to a computer in the nurses' station and allowed the staff to access the sensory input data. Nurses could detect fear or insecurity by the way people hold the robot and provide appropriate assistance.

Purely companion robots in elderly care have only been applied to people with some kind of psychological or physiological disorder. As a result, these studies have distinct target audiences and also different concerns when compared to the purposes of our proposed embodied agent.

Aibo illustrates a robot that can be assigned to both the service type and the companion type (Figure 3c). It is considered by its creators as an entertainment type due to its puppy shaped body [13], and its appearance tries to maintain a lifelike experience to its owners. Tamura et al. started to study the acceptance and effects of this robot on elders with severe dementia [38]. Their study revealed a relevant increase of social actions, emotions and feelings of comfort about past memories.

Table 4: Robots for the aged population, their type and purposes

	Pearl	RoboCare	Care-O-Bot-II	Paro	Huegable	Aibo
Service type	✓	✓	✓			✓
Companion type				✓	✓	✓
Guidance	✓	✓	✓			
Advice	✓	✓	✓			
Therapy				✓	✓	✓

Table 4 groups all the previously mentioned robots and their purposes. This information strengthens the pertinence of our work, since existing robots for the elderly are focused on their physical and mental disabilities. Providing pleasuring activities for the aged population, that are still capable of reasoning, should also be a concern.

Social robots in games

The idea of entertainment robots, which was previously mentioned, is expanding and becoming more frequent. Its general goal is to create a social robot to interact with humans through a specific entertainment activity. These activities should be lifelike experiences providing pleasure and enjoyment feelings. Depending on the target audience, they can also be included in more challenging or even pedagogic activities.

Leite et al. uses the *iCat* robot in a chess game scenario with children [9, 20, 21]. This chess companion also has the role of a tutor due to the help it provides during the game, for instance, it expresses opinions about children’ moves so that they can improve their chess skills. After their first pilot studies, the authors revealed the need of including social and cognitive abilities, commonly referred as empathy. Their further studies introduced into the *iCat* affect recognition in order to improve the robot’s social cues. The way they address this point includes recognising users’ expressions and considering others’ affective states. For instance, when a child is losing, the *iCat* comments about his moves should not cause embarrassment. In addition, and considering their goals were also focused on long-term interactions, this chess player recognises faces and greet people mentioning past events.

This agent has some similarities and differences with the proposed agent of this work. On one hand, including empathic behaviour to robots usually leads to more engaging, natural and likable experiences to users. On the other hand, the *iCat* in this scenario needs access to more details of users’ emotional state because of its tutoring advices. Our *Sueca* player will not advise other players about their actions, instead it will comment the game state. Additionally, the target audience is clearly different and may lead to different concerns, and their work was also focused on long-term interaction.

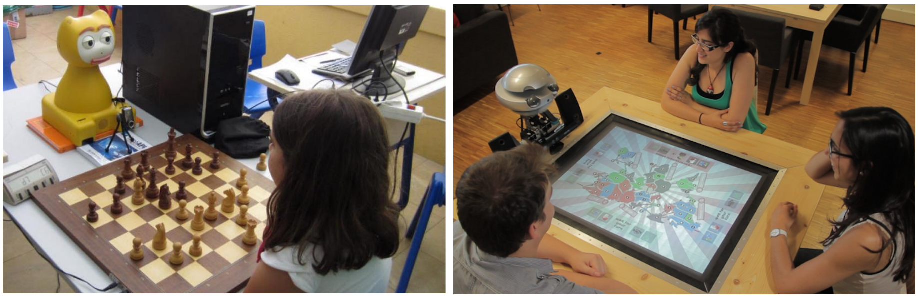

(a) iCat - Chess tutor (b) EMYS - Risk player

Fig. 4: Companion robots in game playing scenarios.

Another example of a robot integrated into a game scenario is the Risk player by Pereira et al. [26]. The goal of their work was to create a robot that interacts with humans and is perceived as socially present in long-term interactions. Firstly, the authors presented how physical embodiments can provide interactivity and, therefore, cause the belief of social presence and improve face-to-face interactions. They also presented some guidelines in order to improve social presence and how they implemented them in the *EMYS* robot for the mentioned scenario [27]. In the Risk scenario, the agent produces non-verbal interactions through a gazing system and a speech direction detector, and it is capable of giving verbal feedback using a topology of speeches according to the game state. Moreover, the authors included an emotion or appraisal system that considers the values of some variables to improve the agent’s behaviours, for instance, every event is rated with a relevance value and the robot only comments important moves. Another example is measuring the power of each player and, since Risk is about conquering and controlling, this power measure is used to shape the robot's mood and defining its strategy to play. Equally important are the simulation of social roles and the luck perception when rolling the dice. All the described behaviours were fully inspired by user studies.

Pereira's work is by far the most similar to the purposes of our goals. It demonstrates how to enrich the Risk game experience with a robot capable of social behaviours at a human level. The main difference from the proposed *Sueca* player is the game. Since no relevant user studies have been done with *Sueca*, applying the Risk' constraints to the *Sueca*'s scenario would lead to inconsistencies. However, an analogous approach might be taken, considering the domain data collection and the following development of the game player architecture.

4 User Centred Studies

Developing a robot for aged people brings some delicate questions. The potential users sometimes have few, or nonexistent, experience with technology, which makes it is difficult for them to understand how robots work and what they can actually do. As a result, understanding their needs, expectations, and fears is another concern [1].

The current section explains the methodology, procedures and preliminary results of an already developed user study in a care home. It involved two different activities, a focus group and a pilot card game study, as a result of two distinct motivations: to understand the elderly' concerns about robots, and to analyse the set-up to further collect information in the game domain.

4.1 Focus Group

A focus group seems to be a good approach for a first meeting due to the informal and conversational way of interacting with participants. The goal of this activity was to introduce to the elderly the robots' theme, and to understand their opinions and expectations. To accomplish this purpose, used techniques were a Brainstorming and a Storytelling.

Methodology

The elderly participants were divided into groups of 5. There were 2 researchers per group commanding and guiding all the process. The list of materials used, per group:

- An illustrative video of existing robots;
- 6 photographs of different robots, including 3 of service type and 3 of companion type (Paro, EMYS, Pleo, Pearl, PR2, and Care-O-Bot);
- Two white boards and three pens (black, red and green);

- Three hypothetical stories of robots;
- An audio recorder;
- Four lavalier microphones;
- A video camera.

The last three items will only be used for a further analysis of this focus group. The video tries to answer the questions: what is a robot, what can robots do, how do they work, do they fail and how do science fiction movies present robots to us. In order not to bias their thoughts, we tried to gather positive and negative aspects of existing robots. The three hypothetical stories aim to bring ethical discussions to the focus group [18, 36]. For instance, an elderly that owns a robot in his home tells him a secret. If that robot is questioned about the secret, should it or should it not tell other people the truth?

Procedures

All the materials enumerated in the previous list were arranged as in Figure 5a. Firstly, each person in the room briefly introduces himself in order to make everyone feeling more comfortable. Secondly, the video is shown. Then, everyone starts discussing about robots' purposes and they are registered in one of the white boards with the black coloured pen. People also express a positive or negative impression of each robot's purpose and their opinions decide the colour of the surrounding line (Appendix A). For instance, the sentence "Call an ambulance" written on the board is surrounded by a green line if they think it is a good purpose for a robot. After finishing this task, one of the group leaders writes all the sentences previously collected in the second board but without the surrounding green or red lines. The other group leader starts reading the hypothetical stories and opens a new discussion about what the robots of each story should do. He also presents the photographs and tries to understand which robot is more suitable for each purpose in their opinion. When bringing the new board to the room, the idea is to understand if their positive and negative opinions about each purpose have changed.

Preliminary Results

This focus group is an ongoing activity that has not yet been fully analysed. Information has already been collected from 3 different focus groups with a sum of 15 participants. For instance, contrasting with what was expected, the elderly do not feel uncomfortable and disrespected with a robot calling the doctor and revealing improper behaviours about its owner (e.g. a diabetic elder eating a chocolate cake slice). Instead, they think it is a valuable aid in their lives and might save them while disrespecting strict instructions. In addition, their safety is their prior worry, and when walking through the house and sometimes due to physical disabilities, they fear about falling and not being noticed.

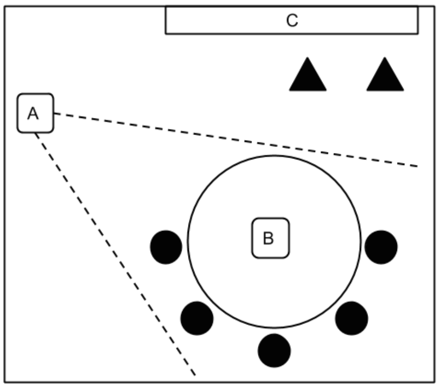

(a) Focus Group

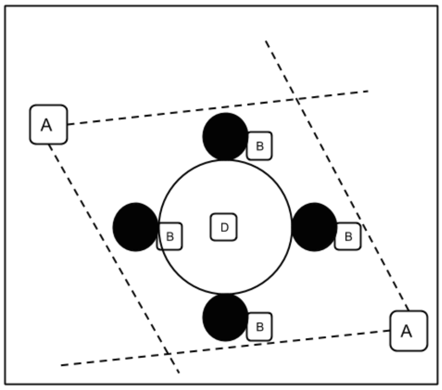

(b) Card game

Fig. 5: Setting of the user study activities. **A** - Video recorder **B** - Audio recorder / Microphone **C** - White board **D** - cards ● - Aged person ▲ - Group leader

4.2 Card Game

A pilot card game activity with the elderly aims two distinct purposes. On one hand, to collect all kind of behaviours and interactions between players during the game. On the other hand, to rehearse and check the technical set-up for further studies. It is important to understand difficulties, for instance, lightning conditions might affect video recording, or the room acoustic and noise might affect audio recording. Testing these conditions is essential to guarantee the analysis of additional studies. Moreover, this activity aims to perceive how the elderly play, what they say and how they behave in certain game states.

Methodology

Recording each *Sueca* card game requires four players, a card deck, a table and chairs for the four players, two video cameras and an audio recorder with four microphones.

Procedures

All the previously enumerated material was arranged as in the Figure 5b. Each video camera was positioned to capture the hands of two adjacent players. Players were recorded during a tournament of several games. They were told to play as long as they wanted with a maximum duration of one hour.

Preliminary Results

The session took only 40 minutes, since the 4 players were feeling weary. Ten games, with an average duration of 3,75' each, were collected. From the average duration, 1’ belongs to the initial setting of shuffling, distributing, and rearranging the cards in each hand. The points per team were being counted during the game.

As expected, players did not frequently talk during the game. *Sueca* is indeed traditionally called a silent game. After a game, paired players frequently discuss extremely good or bad moves from each other. However, the analysis was focused mostly on interactions during the game. Table 5 illustrates the collected expressions, the game stage and its intention. Considering players said specific domain words, expressions were not translated in order not to lose their meaning and regarding the future usage of these sentences in a Portuguese environment.

Table 5: Examples of expressions collected during the card game activity and its respective classification.

Expression	**Game Stage**	**Intention**
Joga [player-name]*!*	Before a play	Speed up a play.
Anda [player-name]*!*	Before a play	Speed up a play.
Podes jogar, [player-name]*!*	Before a play	Speed up a play.
Quase que livrámos.	After collecting the first points	Hopeful/ironic comment.
E eu puxo trunfo.	Initialising a turn with a trump card	State an action.
Outro trunfo!	Play a trump card after an already played trump card	State an action.
Outro(a) [suit]*!*	Play a [suit] card after an already played [suit] card	State an action.
O trunfo é [suit]*.*	Anytime	Give game information. Answer a question.

Regarding the video recording, a relevant aspect that has been noticed was the lighting reflection through the cards. If this technique will be adopted to collect game play information, lighting conditions should be well prepared.

5 Proposed Architecture

The current section describes how to address the development of an artificial *Sueca* player and its integration into an embodied agent that interacts with other players during the game. First of all, it presents the chosen approaches in order to build the artificial *Sueca* player (Section 5.1). Lastly, it introduces the conceptual model and the architecture of the physical embodied agent that reacts socially according to the game state (Section 5.2).

5.1 *Sueca*

Currently, there are no artificial players of the *Sueca* card game and the review of the related work was focus on other hidden information games. Research has shown that the state of the art of imperfect information games is based on Monte-Carlo Methods. To build the *Sueca* card game, the chosen approach is similar to what Buro et al. have done in the Skat card game, since these two games are identical, excluding the nonexistent bidding phase on *Sueca*. Moreover, Furtak et al. have explored how PIMC's results vary according to some of the game properties and have proved its benefits on Skat [23], and, due to the affinity between the two games, it is predictable that the results of applying PIMC to *Sueca* are also satisfactory.

The idea is to evaluate PIMC responses in our domain, since no previous studies have been done. However, if further improvements are needed, ISMCTS and IIMC are still available options. These two approaches were not our first choice due to the computational burden they introduce.

Furthermore, this *Sueca* player will play against aged people. Since they are not world champions or are not even at a professional level, the power of the artificial player must be balanced. On one hand, the idea is to create a challenging environment for the elderly. On the other hand, an existing concern is not to devastate their self-esteem. The motivation of this work is to create a pleasing and, at the same time, stimulating activity for the elderly. As a result, to reinforce the sampling phase of PIMC, an opponent model similar to what has been done in Poker will be used [30]. With this technique, we aim to approximate our artificial agent to a common *Sueca* player.

This model will include cards and actions predictions. Instead of using the random sampling method of the original PIMC, cards probabilities will influence the cards sampling in each iteration of the algorithm. After sampling each world, actions predictions will be used in the simulation of a game to influence the opponents' moves. This technique also aims to introduce some of the common mistakes that the usual player makes, instead of always considering optimal moves.

In order to model opponents, several instances of *Sueca* games will be collected. To easily register this game playing logs, an additional platform must be created.

5.2 The social robot in the game context

Along with the *Sueca* player, this work aims to develop a robot that is socially present in the environment of the game scenario. In order to achieve this goal, many concerns arise. The model presented in Figure 6 tries to solve and organise all the components involved.

This model distinguishes physical components from virtual ones. Some entities are not detailed on the scope of this project and are presented as both physical and virtual components. The human players, Users, play with physical cards on top of a Touch Table, and their game actions are managed by

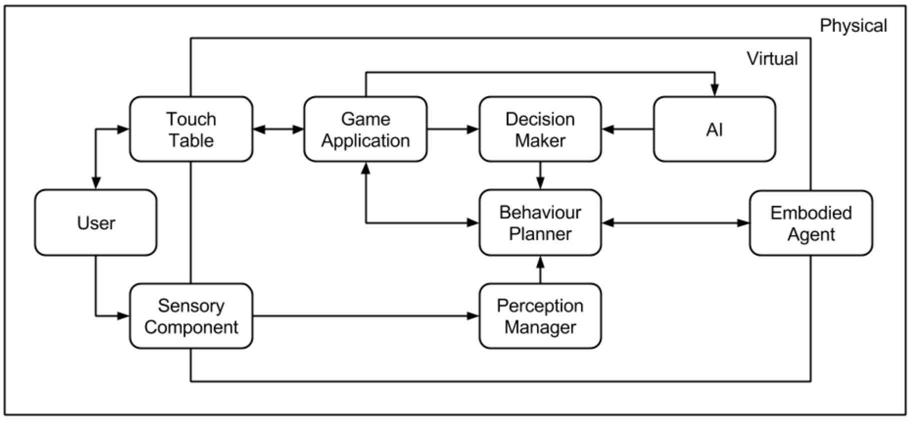

Fig. 6: Structure of the social robot that plays *Sueca*

Game Application and communicated to both the AI module and the Decision Maker. The Perception Manager receives information from all the Sensory Components. The AI module includes all the reasoning about the game and decides the robot's next move. However, the embodied agent's actions also involve social behaviours, and Decision Maker is the responsible module for this management. The Decision Maker balances the AI move and game information, depending on the situation, in order to produce an appropriate sequence of behaviours and inform them to Behaviour Planner. For instance, being the last player of a trick and taking the two highest cards of a suit, by playing a trump card, is an exciting move and should produce an equally exciting reaction on the robot. Lastly, the Behaviour Planner, after receiving high-level intention-directed instructions, builds a suitable plan to execute the chosen instructions, considering the state of the embodied agent, information from Perception Manager, and additional game information from the Game Application.

The architecture that will instantiate the model described above is partially decided. The virtual layer will be mainly covered by Thalamus framework, which enables the communication between the mentioned entities [31]. The chosen Behaviour Planner is Skene [32]. The AI module will be processed offline with the algorithms described in Section 5.1. The embodied agent will initially be EMYS due to its expressiveness, although other robots will be considered, depending on the users' preferences. Finally, the undecided components are the sensory inputs. Since Pereira et al. have shown the importance of collecting data from user studies, these components will be settled further, after some field research with *Sueca* players.

6 Evaluation Methodology

The current section aims to explain the methodology that will be used in order to evaluate the correctness and benefits of the proposed work.

In order to evaluate the proposed artificial player that will be developed using the PIMC algorithm, there are two main considerations: parametrization and performance. The parametrization issue will be addressed by measuring the time of both the offline pre-computation and runtime decision, varying the value of the sampling limit parameter. The average points per tournament will be used as a performance measure, and it will be compared to naive approaches (e.g. rule-based). In addition, our artificial player will play against humans in order to evaluate its performance. The last mentioned evaluation will not be developed with aged people, considering that finding a group of elderly to do it is not simple, and it will instead use the university community.

Concerning the integration of a social embodied agent into the game scenario, a proper user study of elderly playing *Sueca* with EMYS will be settled. Each group will play a tournament with two different conditions of the embodied agent:

- An agent that plays the game with few or nonexistent social behaviours;
- An agent that plays the game and reacts according to the game state with verbal and nonverbal cues.

After the tournament, each person will answer a questionnaire in order to evaluate the individual experience. This questionnaire aims to measure the participants' perception of the robot and also their presence perception of the embodied agent, using, respectively, the Godspeed questionnaire [3] and Networked Minds [4].

7 Conclusion

The pertinence of this work has been demonstrated considering the related work presented. The state of the art robots for elderly is focused on service robots to guide them in their daily tasks. Considering the companion type robots for aged people, existing technology is focused on therapy especially for the disabled. However, the proposed robot aims to interact with elderly in a specific scenario for an entertaining and stimulating activity.

The *Sueca* player will be based on Monte-Carlo Methods, that are currently being successfully used on similar card games. Additionally, in order to balance the power of this artificial player, an opponent modelling will also be included. This artificial player should also be a challenging opponent for the elderly players, while interacting socially with opponents according to the game state.

The development of the proposed work will follow the schedule presented in Appendix B.

References

1. Alves-Oliveira, P., Petisca, S., Janarthanam, S., Hastie, H., Paiva, A.: How do you imagine robots? Childrens expectations about robots (2014)
2. Bahadori, S., Cesta, A., Grisetti, G., Iocchi, L., Leone, R., Nardi, D.: RoboCare: an Integrated Robotic System for the Domestic Care of the Elderly (December 2002), 1–16
3. Bartneck, C., Kulić, D., Croft, E., Zoghbi, S.: Measurement Instruments for the Anthropomorphism, Animacy, Likeability, Perceived Intelligence, and Perceived Safety of Robots. International Journal of Social Robotics 1(1), 71–81 (Nov 2008), `http://link.springer.com/10.1007/s12369-008-0001-3`
4. Biocca, F., Harms, C., Gregg, J., Interface, M., Labs, N.D.M.I.N.D.: The Networked Minds Measure of Social Presence : Pilot Test of the Factor Structure and Concurrent Validity Co-Presence (2001)
5. Broekens, J., Heerink, M., Rosendal, H.: Assistive social robots in elderly care: a review Assistive social robots 8(2) (2009)
6. Browne, C.B., Powley, E., Whitehouse, D., Member, G.S., Lucas, S.M., Member, S., Cowling, P.I., Rohlfshagen, P., Tavener, S., Perez, D., Samothrakis, S., Member, G.S., Colton, S.: A Survey of Monte Carlo Tree Search Methods 4(1), 1–43 (2012)
7. Buro, M., Long, J.R., Furtak, T., Sturtevant, N.: Improving State Evaluation, Inference, and Search in Trick-Based Card Games pp. 1407–1413
8. Campbell, M., Hoane, A.J., Hsu, F.h.: Deep Blue 134, 57–83 (2002)
9. Castellano, G., Leite, I., Pereira, A., Martinho, C., Paiva, A., Mcowan, P.W.: Affect recognition for interactive companions: challenges and design in real world scenarios pp. 89–98 (2010)
10. Cowling, P.I., Powley, E.J., Whitehouse, D., Member, S.: Information Set Monte Carlo Tree Search 4(2), 120–143 (2012)
11. Ferrucci, D., Brown, E., Chu-carroll, J., Fan, J., Gondek, D., Kalyanpur, A.A., Lally, A., Murdock, J.W., Nyberg, E., Prager, J.: Building Watson: An Overview of the DeepQA Project pp. 59–79 (2010)
12. Frank, I., Basin, D.: Search in games with incomplete information: a case study using Bridge card play. Artificial Intelligence 3702(97) (1998)
13. Fujita, M.: AIBO: Toward the Era of Digital (1) (1983)
14. Furtak, T., Buro, M.: Recursive Monte Carlo Search for Imperfect Information Games
15. Gelly, S., Silver, D.: Monte-Carlo tree search and rapid action value estimation in computer Go. Artificial Intelligence 175(11), 1856–1875 (Jul 2011), `http://linkinghub.elsevier.com/retrieve/pii/S000437021100052X`
16. Ginsberg, M.L.: GIB: Imperfect Information in a Computationally Challenging Game 14, 303–358 (2001)
17. Graf, B., Hans, M., Schraft, R.D.: Care-O-bot II Development of a Next Generation Robotic Home Assistant pp. 193–205 (2004)
18. Kahn, P.H., Ishiguro, H., Friedman, B., Kanda, T.: What is a human? - Toward psychological benchmarks in the field of human-robot interaction. In: Proceedings - IEEE International Workshop on Robot and Human Interactive Communication. pp. 364–371 (2006)
19. Leite, I.: iCat, the Chess Tutor An Affective Game Buddy Based on Anticipatory Mechanisms. Ph.D. thesis, Instituto Superior Técnico (2007)
20. Leite, I., Mascarenhas, S., Martinho, C.: Why cant we be friends? An empathic game companion for long-term interaction

21. Leite, I., Pereira, A., Martinho, C., Castellano, G.: Towards an Empathic Chess Companion
22. Long, J., Buro, M.: Real-Time Opponent Modelling in Trick-Taking Card Games pp. 617–622 (2009)
23. Long, J., Sturtevant, N.R., Buro, M., Furtak, T.: Understanding the Success of Perfect Information Monte Carlo Sampling in Game Tree Search pp. 134–140 (2010)
24. Nash, J.: Non-Cooperative Games. Ph.D. thesis (1950)
25. Osborne, M.J.: A Course in Game Theory (2011)
26. Pereira, A.: Socially Present Agents for Tabletop Games. Ph.D. thesis (2014)
27. Pereira, A., Prada, R., Paiva, A.: Socially Present Board Game Opponents
28. Pollack, M.E., Engberg, S., Matthews, J.T., Dunbar-jacob, J., Mccarthy, C.E., Thrun, S.: Pearl: A Mobile Robotic Assistant for the Elderly (2002)
29. Pompili, A., Abad, A., Trancoso, I., Fonseca, J., Martins, I.P., Leal, G., Farrajota, L.: An on-line system for remote treatment of aphasia pp. 1–10 (2011)
30. Ponsen, M., Gerritsen, G., Chaslot, G.: Integrating Opponent Models with Monte-Carlo Tree Search in Poker Monte-Carlo Tree Search (Coulom 2006), 37–42 (2008)
31. Ribeiro, T., Tullio, E., Corrigan, L.J., Jones, A., Aylett, R., Castellano, G., Paiva, A.: Developing Interactive Embodied Characters using the Thalamus Framework: A Collaborative Approach pp. 1–10
32. Ribeiro, T., Tullio, E., Paiva, A.: From Thalamus to Skene: High-level behaviour planning and managing for mixed-reality characters
33. Russell, S., Norvig, P.: Artificial Intelligence: A Modern Approach. chap. 5 - Advers, pp. 161 – 190. 3rd edn. (2009)
34. Schaeffer, J., Lake, R., Lu, P., Bryant, M.: CHINOOK The World Man-Machine Checkers Champion 17(1), 21–29 (1996)
35. Shoham, Y., Leyton-brown, K.: Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations (2010)
36. Should, I.: Granny and the robots: Ethical issues in robot care for the elderly. pp. 1–43 (2010)
37. Stiehl, W.D., Breazeal, C., Han, K.h., Lieberman, J., Lalla, L., Maymin, A., Salinas, J., Toscano, R., Tong, C.H., Kishore, A., Berlin, M., Gray, J.: The Huggable: A Therapeutic Robotic Companion for Relational, Affective Touch p. 2005 (2005)
38. Tamura, T., Yonemitsu, S., Itoh, A., Oikawa, D., Kawakami, A., Higashi, Y., Fujimooto, T., Nakajima, K.: Is an Entertainment Robot Useful in the Care of Elderly People With Severe Dementia ? 59(1), 83–85 (2004)
39. Wada, K., Shibata, T.: Living With Seal Robots Its Sociopsychological and Physiological Influences on the Elderly at a Care House 23(5), 972–980 (2007)
40. Wada, K., Shibata, T., Saito, T., Sakamoto, K., Tanie, K.: Psychological and Social Effects of One Year Robot Assisted Activity on Elderly People at a Health Service Facility for the Aged (April) (2005)
41. Wada, K., Shibata, T., Saito, T., Tanie, K.: Effects of Robot Assisted Activity to Elderly People who Stay at a Health Service Facility for the Aged. Conference on Intelligent Robots and Systems (2003)
42. Zinkevich, M., Bowling, M., Johanson, M.: Regret Minimization in Games with Incomplete Information pp. 1–14

A Brainstorming board

B Planning

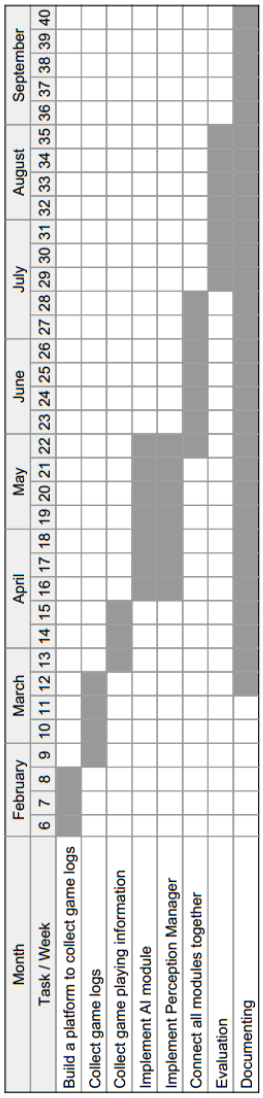
Month
Task / Week
February
March
April
May
June
July
August
September
6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31 32 33 34 35 36 37 38 39 40
Build a platform to collect game logs
Collect game logs
Collect game playing information
Implement AI module
Implement Perception Manager
Connect all modules together
Evaluation
Documenting