

Deceptive actions to improve the attribution of rationality to playing robotic agents

Ewerton de Oliveira · Laura Donadoni ·
Stefano Boriero · Andrea Bonarini

Received: date / Accepted: date

Abstract Play is a common activity, providing not only pleasure but also physical and cognitive development. In the quest for new playing experiences, there is an increasing tendency to develop robots playing with people. Making believable playing robots able to keep human players engaged and satisfied by the playing experience is the main challenge.

In this work, we investigate the possibilities of a playful interaction between a human player and a mobile robot. In particular, this paper focuses on the applicability of deception as a means to support engagement and the attribution of rationality to playing robotic agents. By analyzing the interaction situation between the human and robot players, by identifying the need for deception, and by deciding whether and how to deceive, we aim at increasing self-reported engagement and fun, which are also related to the perception of the robotic opponent as smart enough to compete at an appropriate level.

Experiments were conducted on a sample of 78 subjects facing two different deceptive behaviors and a basic behavior without any deception. All participants responded to a post-interaction questionnaire from which it was possible to observe a positive acceptance of the perception of the robot as a rational agent aimed at winning. In general, deception was perceived by most of the players as one of the robot's abilities, when actuated, and contributed to the reported fun.

Keywords Robogames · Human-robot interaction · Deception · Game theory · Rational agent · Play · Robogame

Ewerton de Oliveira · Laura Donadoni · Stefano Boriero · Andrea Bonarini
Politecnico di Milano, Dipartimento di Elettronica e Informazione, Piazza Leonardo da Vinci, 32, 20133, Milano, Italy
Tel.: +39-02-2399-3525
E-mail: andrea.bonarini@polimi.it

1 Introduction

Play is a basic need for people that supports the development of abilities and provides reward and amusement [10, 34, 42]. Recent technological developments make it possible to include robots as players also in games that require physical engagement. This opens interesting research opportunities in the field of human-robot interaction.

Robotic toys and animated plushes enter our homes at a rate of millions every year [1, 2]. In most cases, the interaction is limited to stimulus-response reactions, but the introduction of low-cost sensors and computer power have made it possible to introduce a richer interaction and games specifically designed for autonomous robots. Robogames are expected to be one of the next robotic products for the mass technological market [14], but, in order to respond to demands about engagement support [48–50], we still need to explore how to enable autonomy, intelligence, and adaptive behavior in such a highly dynamical activity.

The present work focuses on a challenging type of games, where human players are involved in a physical and quite demanding activity with robots. This type of games has been introduced as *Physically Interactive RoboGames* (PIRG [29]), and also as a possibility for *Phygital play* [27]. In PIRG, physical, autonomous, (often moving) robotic agents are actively engaged with people in a game based on physical interaction, while acting in the physical world. The task is quite challenging for the robot, since, by using data collected in real-time on the field, in addition to the usual tasks for a mobile robot (navigation, self-localization, data fusion, etc.), it must comply to the game rules, play a convincing role, and, at the same time, preserve the player’s physical integrity.

In our research, we are focusing on adversarial PIRG games, where the human player and the robot are opponents pursuing contrasting goals. An important aspect that supports engagement in this type of games is the perception of the robotic opponent as a rational agent, an entity that aims at achieving its goal, i.e., win the game. This goes over the simple concept of *animacy* [30], since the robot should not only be perceived as animated, but also as purposeful. A competitive game is an ideal situation to study the perception of rational behavior in a robot, since interaction is needed, and goals are clearly defined by shared rules.

A general requirement for enjoyable games is to maintain an even possibility to win for both players, thus making the game interesting and engaging at the right level [15]. This can be made possible by generating and properly managing a player model good enough to adapt the robot’s behavior to the specific player, while playing the game.

In general, a human player has expectations that should be matched by a clear exhibition of intention of the robot, in particular, the one related to its aim to win the game [29]. The equilibrium between the exhibition of a rational behavior and the covert aim of playing an even game are goals for a good PIRG design.

In the study presented in this paper, we focus on deception and we propose a framework for deciding and communicating deception in PIRG applications.

Possible issues about the opportunity for a robot to deceive a human might be raised. However, in this type of playful situation, a deceptive behavior is part of the game, as well as an obvious expectation from players, and public. Most adversarial games based on ability, from soccer to boxing, are defined by rules that limit the type of acceptable deception, but also leave space to the possibility to exploit abilities to trick the opponent's senses. This is also an intrinsic aspect of many games played by animals to learn to face real hazards. Therefore, a robot trying to deceive a human in a PIRG is actually playing its role as expected, and will not be blamed for this. The experiments reported in this paper show that the interaction with a deceptive robot in a PIRG can improve the perception of animacy in the robot, as this is the way interaction is expected to occur. This is on the line of designing interaction to enhance the perception of animacy [4, 18]. In other games, not involving physical interaction, deception has been recognized as a way to increase attributions of mental states to the robot [38].

This paper is organized as follows. Section 2 introduces related works, section 3 explains in details the game we are considering to evaluate our approach, in section 4 the method for detecting the need of deception and how to translate it into motion commands are described; experiments and validation of the proposed method are presented in section 5, and, finally, discussions and conclusions follow in sections 6 and 7.

2 Related Works

2.1 Physically Interactive Robogames

One of the critical aspects in a PIRG concerns adaptation to the players' abilities, so to engage them in a situation that is challenging enough, but not too demanding, as stated by the commonly accepted theory of *flow* [15]. Since we would not like to ask the player to select a priori a degree of difficulty for the game, which would reduce the credibility of the robot as a smart companion, it is relevant to be able to model the player's characteristics so to adapt the robot's behavior to obtain even possibilities to win the game. The research activities supporting the video game industry produced several studies, most of which use artificial intelligence and machine learning algorithms to achieve this result. A general model is presented in [51].

There have been attempts for presenting robots as toys over the years, where, in most cases, the robot acts more or less like a mobile pet (e.g., [9, 17]). In these cases, interaction is often seen as limited to almost static positions, not exploiting rich movement, nor high level of autonomy; the credibility of these toys to lively engage people, such as kids, are said to be constrained [29]. The lack of engagement actually lead, in some recent, significant cases, to commercial failures [19].

Some PIRGs have been reported over the last few years. Jedi Trainer [29] was a PIRG mimicking the light-saber training situation included in the first Star Wars saga movie “Episode IV - A New Hope”: it was able to show some apparent adaptation to the player’s playing style as well as some realism of the drone’s behavior, which could be perceived by the human players as a kind of rational behavior.

Many robots are used with people with disability, in particular autistic children, for which engagement is critical. Some of them have been explicitly designed to play physical games. Queball [37] and LEKA [21] are robotic, mobile balls proposed to engage autistic children in different games that require also movement. IROMECA [7, 24] was another mobile robot that included play activities involving movements. Teo [8] is a huggable, mobile robot designed to play games with autistic children, and to provide the possibility to implement free play [6], as well as structured play experiences.

Despite being examples of successful applications, none of the mentioned examples includes any online player modeling activity, and the abilities and behaviors of the robots are pre-defined.

RoboTower [31–33], is a game where human players physically interact with an omni-directional robot by trying to conquer towers while preventing the robot to push them down. A player model is maintained and used to adapt the game to the performance of the player.

An approach based on genetic algorithms for capturing and modeling individual entertainment is described in [50], where the main goal is to construct a model of a child playing a Playware game. “Playware is the use of intelligent technology to create the kind of leisure activity we normally label play” [26]. The system can predict the answers to a question asking which variants of the game are more or less “fun”. The model is built from physiological signals measured during play.

In summary, pretty much all the related works agree that being able to classify the player’s behavior in a game may enable the robotic agent to modify its interaction and playing style in order to adapt to the skills of its human counterpart. This is a general statement shared with most human-robot interaction applications. With the present work, we go further and investigate the introduction of deceptive behavior as an additional factor to support engagement, and to increase the opinion about the robot as a smart device against which it is nice to play.

2.2 Deception and robots

Models to describe how people adjust their interactive behavior as consequence of their perception of a social situation have been proposed in Sociology. We selected the *interdependence theory* [23] as a suitable framework for our purposes. The adjustment of interaction depends on rewards and costs related to each action choice. Every situation can be expressed as a matrix of rewards associated to each choice, called *outcome matrix*. This model can be seen as

an equivalent of a normal form game in Game Theory. Furthermore, the authors introduced a four dimensional space where mapping of social situations is supported by outcome matrices that can be defined from the analysis of the opportunities. *Interdependence*, *correspondence*, *control*, and *symmetry* are the four dimensions that define this space. We use the first two dimensions to decide when to apply deception, as described in Section 4. The interdependence dimension measures the extent to which each individual's outcome is influenced by the other's actions in a given situation. Correspondence describes the extent to which the outcomes of an individual in a situation are consistent with the outcomes of the other individual [43]; this happens when the selected actions bring similar outcomes to both individuals, for example when they act to collaborate to achieve the same goal. In our case, we apply this concept to make the robot selecting actions that bring it towards the common goal of obtaining a game with even possibilities to win.

In [44] an algorithm based on the interdependence theory for detecting when and whether a robot should deceive was proposed; the decision is based on the mapping of the outcome matrix to the interdependence-correspondence space. This algorithm has been used in [45], where the analysis of the social situation is performed to provide robots with the capacity of determining whether deception is needed. The works reported in [23, 43–45] have laid the foundation for many research activities in social aspects involving humans and robots, in particular on the aspect of *deception*.

The method we propose takes inspiration from [23] as well, but in opposition to [43], revisits the interdependence space and the way of mapping the outcome matrices, as it will be discussed in Section 4.

Deceiving a human is not trivial. Studies about a robot able to deceive a human have been reported in [40], where the authors discuss whether the feeling of being deceived by a robot would be an indicator that the human treats the robot as an intentional entity. Further experiments about increasing engagement in a game due to a robot able to deceive, have been conducted in [39]. In this work, a cheating robot plays “rock-paper-scissors”, and takes advantage of the deception for its own benefit. The authors found a noticeably increased level of engagement by the participants when the robot cheats. Many could be the motivations for this, including increased attention to catch cheating. However, a robot that does not follow the rules of a game cannot be accepted in real applications: for the same reasons a cheating human player would be excluded from the game. Differently from the above mentioned situation, the robot presented in [41] was programmed to follow a multi-player, reaction-time, conflict game, where its main role was to establish who was winning. The robot did not play, but it sometimes changed the wins in order to make the game more socially engaging when the players discover the cheating behavior of the robot.

The most important part in a deceiving algorithm is to properly transmit to the deceived the false communication. This is a delicate aspect, because the deceived should not perceive to be deceived.

In [16] the authors propose a way to study the communication of false information (deception) using a robotic arm. The authors reinforced the application of robot deception, thus making games against the robot more engaging. The paper focuses on robot deception in goal-directed motion, by learning deceptive trajectories in which the robot is concealing its actual goal. Our proposal is similar to this, since it explores goal-directed motion, but differs from it by using a mobile robot in a real game scenario, and by the decision mechanism.

3 The game environment

In this section we describe the experimental setting adopted to evaluate our hypotheses: Robotower2.0¹.

3.1 Playground and game rules

The playground is a rectangular area of 4 m × 4 m. On each corner, 1 meter high props, representing towers, are placed.

The robot and the human player have conflicting goals: the robot should hit a tower and make it fall, while the human player should secure all the existing towers without letting a single tower be torn down by the robot. Each tower is equipped with a button and four LEDs. When the player presses the button for at least 2.5 seconds continuously, a new LED is turned on. When all the four LEDs have been turned on, the tower is captured by the human player and cannot be considered by the robot any longer. Such LEDs are a representation of the progress of the human player in the attempt of capturing a specific tower. The activity of pushing the button can be distributed in different moments, so that the players will not lose their progress if they leave the button before a tower is captured, for instance to defend another tower attacked by the robot. Tower information such as tower status (whether active, fallen, or captured) and button pressing activity are transmitted to the robot via wireless communication, so to inform it about the game status.

If, at anytime, a tower falls (because of the robot or the player), the game ends, and the human player loses.

In figure 1 we can see a picture taken from a game.

3.2 Robotic agent

The robot is a holonomic platform, 115 cm high, with a triangular base which can be inscribed in a circle of 50 cm in diameter. The robot can run at a speed of up to 1 m/sec. It is capable of navigating autonomously in the environment,

¹ A video illustrating the basic game scenario and its rules is available at <https://www.youtube.com/watch?v=3azXf8V64iM>



Fig. 1: Human-Robot interaction scene from our experiments.

using nodes developed in the ROS² framework. The robot localization is performed by using Monte Carlo localization algorithms through the AMCL³ node, using laser range scans. For the management of sensor information we have implemented custom ROS nodes. By relying on its laser scanners, the robot can also perceive the human players during the game, so to track them, and avoid hitting them while moving. For the player, a possibility to defend a tower is to stay between the robot and the tower.

Since the robot is holonomic, it is able to move across the entire playground in a way similar to the human player, and its goal of tearing down towers is constrained by the fact that an already captured tower, or one whose button is currently being pushed by the player, cannot be attacked. Moreover, the players can also block the robot path by staying on it, since the robot would avoid them. Notice that, while the player is trying to capture a given tower, the robot can change its target and try to tear down another one, so that the human player has to run to defend the new target abandoning the tower she/he was trying to capture.

4 Method

In this section, we describe how to decide when to deceive and the different strategies to implement deception. The presented approach relies on the definition of outcome matrices for both players and the computation of *interdependence* and *correspondence* [44]. First, the way to decide whether to deceive is presented, then how deceptive behaviors have been implemented.

² <http://www.ros.org>

³ <http://wiki.ros.org/amcl>.

4.1 Deciding to deceive

During the game, the robot continuously (re)calculates the payoff estimates for each of its possible goals in a process that depends on the specific game. In our game, the payoff for attacking a given tower is considered. At each time, the goal with higher expected payoff is selected. Given the set of tower coordinates $\mathcal{T} = \{\tau_i\}_{i=1}^{N=4}$, and the vectors $\mathbf{l}_r \in \mathbb{R}^2$ and $\mathbf{l}_p \in \mathbb{R}^2$ representing, respectively, the robot's and human player's spatial position, and $\delta(\mathbf{a}, \mathbf{b})$, computing the Euclidean distance between the two vectors given as arguments, it is possible to define two vectors used to compute the estimated payoff for the robot ($\boldsymbol{\rho}_r$) and the human player ($\boldsymbol{\rho}_p$). In this specific game, both payoffs are defined by spatial relations. $\boldsymbol{\rho}_r$ estimates the robot's payoff as a difference between the distance to each tower, respectively from the human player and from the robot: the higher this difference is for a tower, the higher the potential payoff. $\boldsymbol{\rho}_p$ estimates the player's payoff as the sum of the inverse of the respective distances: the higher this sum the closer are the two players, situation that is desirable by the human since it makes her/him controlling the robot. The formal definition of the payoff vectors is reported below.

$$\boldsymbol{\rho}_r = \begin{bmatrix} \delta(\tau_1, \mathbf{l}_p) - \delta(\tau_1, \mathbf{l}_r) \\ \delta(\tau_2, \mathbf{l}_p) - \delta(\tau_2, \mathbf{l}_r) \\ \delta(\tau_3, \mathbf{l}_p) - \delta(\tau_3, \mathbf{l}_r) \\ \delta(\tau_4, \mathbf{l}_p) - \delta(\tau_4, \mathbf{l}_r) \end{bmatrix} \quad (1)$$

$$\boldsymbol{\rho}_p = \begin{bmatrix} \delta(\tau_1, \mathbf{l}_r)^{-1} + \delta(\tau_1, \mathbf{l}_p)^{-1} \\ \delta(\tau_2, \mathbf{l}_r)^{-1} + \delta(\tau_2, \mathbf{l}_p)^{-1} \\ \delta(\tau_3, \mathbf{l}_r)^{-1} + \delta(\tau_3, \mathbf{l}_p)^{-1} \\ \delta(\tau_4, \mathbf{l}_r)^{-1} + \delta(\tau_4, \mathbf{l}_p)^{-1} \end{bmatrix} \quad (2)$$

Note that $\boldsymbol{\rho}_r \in \mathbb{R}^4$ and $\boldsymbol{\rho}_p \in \mathbb{R}^4$, since the cardinality $n(\mathcal{T})$ of the set of towers \mathcal{T} is equal to four.

The robot's target tower τ_r^* is selected by detecting the maximum value in the array $\boldsymbol{\rho}_r$:

$$\tau_r^* = \arg \max_i \boldsymbol{\rho}_r^{(i)}, \quad (3)$$

where i refers to the vector component index. In other terms, the robot would aim at the tower that is closer to it and farther away from the player.

The robot calculates in an analogous way also the expected target preference for the player, assuming the player is a rational agent that would select the best move.

$$\tau_p^* = \arg \max_i \boldsymbol{\rho}_p^{(i)} \quad (4)$$

To decide when it is the case to deceive, a set of outcome matrices that model the in-game interaction have been defined. Such matrices quantify the payoff, also referred to as *utility*, associated to each player's action. As actions,

the possibility for the players to move towards any of the four towers are considered.

Notice that, despite the fact that the presented implementation of the two outcome matrices reflects the specific goals of the players in this game, similar matrices can be defined for any conflicting situation, not only for games, following the approach typical of Game Theory. The basic aspects of the proposed approach are not tied to the specific shape of the outcome matrices, and could be applied to any other conflicting situation.

We report below how the outcome matrices are computed in this specific game, as an example of the general approach. These matrices include, respectively on rows and columns, the possible actions each of the two players can take, and the corresponding values are the payoffs for the respective choices.

As in any conflicting game, each of the two players gains benefit from the loss of the other. For this reason, the robot earns more when itself and the opponent choose a different target. From a matrix point of view, this is seen as lower values on the diagonal, because this area represents when the two opponents choose the same action. On the other hand, the player's outcome matrix is built considering that the human player earns more when her/his choice is the same as the robot's one. For this reason, the values on the diagonal of the player's outcome matrix are greater than the non-diagonal ones.

The elements of the robot's outcome matrix \mathbf{O}_r are periodically computed as described below.

– *Non-diagonal Elements:*

$$\gamma_r \cdot \frac{\rho_r^{(i)}}{\sum_i \rho_r^{(i)}} \cdot \frac{\delta(\tau_i, \mathbf{l}_p)}{\sum_i \delta(\tau_i, \mathbf{l}_p)}, \quad (5)$$

where $\rho_r^{(i)}$ refers to the component i of the vector ρ_r , and γ_r is a tuning constant.

– *Diagonal Elements:*

$$\delta(\tau_i, \mathbf{l}_r) \quad (6)$$

In \mathbf{O}_r , γ_r is used to weight the non-diagonal values, produced by equation 5, while in the player's matrix, another tuning constant has been used for the diagonal values. Tuning the values of the outcome matrices is needed because the two players tend to win in two different situations, corresponding to opposite actions: the robot would like to run "alone" to the tower, while the players prefer to have the robot next to them, so that the robot cannot run and take another tower. Giving the right values to the tuning constants is necessary to balance the interdependence-correspondence values.

The second term in equation 5 is used to give a weight to the payoff: the more an individual is interested in a given target (supposed that this is a consequence of rational/maximizer thinking), the higher the payoff that is obtained when he can achieve it. The third and last term in equation 5 expresses how much advantage (spatial distance) the player has on the other one.

		ROBOT			
		T1	T2	T3	T4
PLAYER	T1	8 / 1.4	0 / 4	9 / 2	8 / 0.3
	T2	4 / 2	9 / 3	8 / 1	7 / 0.5
	T3	6 / 9	2 / 4	8 / 3	8 / 2
	T4	2 / 5	3 / 3	5 / 6	7 / 4.2

Fig. 2: Example of outcome matrix. The columns represent the robot's payoff to attack one of the four towers. The rows represent the player's payoff. Each element is a couple of values, respectively the player's payoff (below) and the robot's payoff (above).

Similarly, the human player's outcome matrix \mathbf{O}_p is computed as described below.

– *Non-diagonal Elements:*

$$\frac{\rho_p^{(i)}}{\sum_i \rho_p^{(i)}} \cdot \frac{\delta(\tau_i, l_r)}{\sum_i \delta(\tau_i, l_r)} \quad (7)$$

– *Diagonal Elements:*

$$\gamma_p \cdot \delta(\tau_i, l_r)^{-1} \quad (8)$$

An example of outcome matrix is reported in Figure 2

Using these equations, the utility of both the players is periodically recomputed and then incorporated into the interdependence-correspondence space [45], as described below.

For our purposes, the *interdependence* dimension represents how much influential the outcome of a player is on the other one, while *correspondence* quantifies how much conflict exists between the payoffs of the actions selected by the two players. These two values are calculated by considering three aspects: the variation of the robot's outcome matrix resulting from its own decisions, the variation of the player's outcome matrix resulting from the player's decisions, the variance in outcome matrices resulting from both, joint, interactive decisions.

Algorithm 1 computes *interdependence*, where the output $\alpha \in [-1; 0]$ represents the level of interdependence of the considered situation at that specific time, and Δ is the maximum variation the player can obtain with respect to the current situation on the robot's outcome matrix, by taking the given action, i.e. deciding to go to tower i . The variable K represents the range, used for normalization, and $n(\mathcal{T})$ the number of towers.

```

input : Robot's outcome matrix  $\mathbf{O}_{n(\mathcal{T}) \times n(\mathcal{T})}$ 
output:  $\alpha$  (Interdependence value)
for  $i \in \text{range}(1, n(\mathcal{T}))$  do
   $\Delta = |\max(\mathbf{O}_r(:, i)) - \min(\mathbf{O}_r(:, i))|$ 
   $K = \max(\mathbf{O}_r(:, i)) + \min(\mathbf{O}_r(:, i))$ 
   $\alpha += \frac{\rho_r^{(i)}}{\sum_i \rho_r^{(i)}} \cdot \frac{\Delta}{K}$ 
end

```

Algorithm 1: Interdependence algorithm.

Algorithm 2 computes the *correspondence* value, represented by $\beta \in [0; 1]$.

```

input : Outcome matrices  $\mathbf{O}_r, \mathbf{O}_p$ 
output:  $\beta$  (Correspondence value)
Avg(outcome  $\forall$  Robot's action)
Avg(outcome  $\forall$  Player's action)
for  $i \in \text{range}(1, n(\mathcal{T}))$  do
   $amr = \text{tower index that maximizes robot's outcome}$ 
   $amp = \text{tower index that maximizes player's outcome}$ 
   $\Delta_r = \mathbf{O}_r(amr, i) - \mathbf{O}_r(amp, i)$ 
   $\Delta_p = \mathbf{O}_p(amr, i) - \mathbf{O}_p(amp, i)$ 
   $K_r = \mathbf{O}_r(amr, i) + \mathbf{O}_r(amp, i)$ 
   $K_p = \mathbf{O}_p(amr, i) + \mathbf{O}_p(amp, i)$ 
   $\beta += 0.5 \times \frac{Avg_r}{\sum Avg_r} \times \frac{\Delta_r}{K_r} \times \frac{\Delta_p}{K_p}$ 
end
for  $i \in \text{range}(1, n(\mathcal{T}))$  do
   $amr = \text{tower index that maximizes robot's outcome}$ 
   $amp = \text{tower index that maximizes player's outcome}$ 
   $\Delta_r = \mathbf{O}_r(i, amr) - \mathbf{O}_r(i, amp)$ 
   $\Delta_p = \mathbf{O}_p(i, amr) - \mathbf{O}_p(i, amp)$ 
   $K_r = \mathbf{O}_r(i, amr) + \mathbf{O}_r(i, amp)$ 
   $K_p = \mathbf{O}_p(i, amr) + \mathbf{O}_p(i, amp)$ 
   $\beta += 0.5 \times \frac{Avg_p}{\sum Avg_p} \times \frac{\Delta_r}{K_r} \times \frac{\Delta_p}{K_p}$ 
end

```

Algorithm 2: Correspondence algorithm.

The procedure works by going through every action the players can take. For example, if the robot considers taking action 1 (i.e., go to tower 1), it temporarily ignores all the other columns of the outcome matrix (corresponding to its other possible actions) and checks what are the indexes (corresponding to the actions of the player) that maximize both its outcome and the player's one.

It then calculates the difference, Δ , between the outcomes corresponding to the two indexes for both robot and player. These Δ are then normalized with respect to the best outcomes respectively for the robot and the player, and then multiplied by a factor that expresses how likely, in terms of the average of outcomes, the action would be taken.

Once an extremely dependent and conflicting situation is detected, deception is triggered: in our scenario, we identified the critical condition to trigger deception to be $\alpha > 0.6$ and $\beta < -0.6$. These values have been empirically tuned as the most suitable for average players.

When deception is decided, the algorithm calculates the fake target to be communicated to the player through an appropriate action. The fake target is chosen by the following steps: the robot calculates the two actions that would provide itself the highest reward, and calculates which of the two maximizes the player's payoff. This last step is aimed at identifying the fake target that would motivate more the player to aim at it, so to be deceived. The robot will select as its real target the other one of the selected two. The control flow of the complete algorithm is detailed in Figure 3.




Fig. 3: The procedure to select deception.

4.2 Communicating false goals

In general, deception should be implemented by considering the possible goals the human subject could attribute to the robotic player, and the ways to communicate the false goal. It is part of the deception process to delay as much as possible the discovery of the true goal. A posteriori, the human player should perceive to have been deceived, so to attribute to the robot some intention to deceive.

Having presented how to trigger deception, in this section we present the two different methods developed to implement deceptive behaviors in the test setting: the first keeping the speed of the robot constant and changing the trajectory at a late moment, the last reaching a point and then suddenly changing speed and direction. The aim for these behaviors is to communicate false targets, i.e., false attempts to attack a tower.

We decided to implement two different deception behaviors, both to explore what could actually be perceived as a deception by the player, and, possibly, whether diversity of trajectories could increase the appeal of the game by making the robot movements less predictable. These can be taken as examples,

specific to this setting, of possible implementations of ways to communicate a fake goal.

4.2.1 Static trajectory approach

In this first method, the trajectories are generated when the navigation node receives the fake and the real targets. Based on the actual position of the robot, this algorithm calculates the series of points the robot has to reach in order to communicate the deception.

The algorithm can select one of two different types of trajectories. The trajectories are implemented by deciding a series of points to be followed. These points are spatially distributed along the desired trajectory. A visual description of both trajectories is presented in Figure 4.

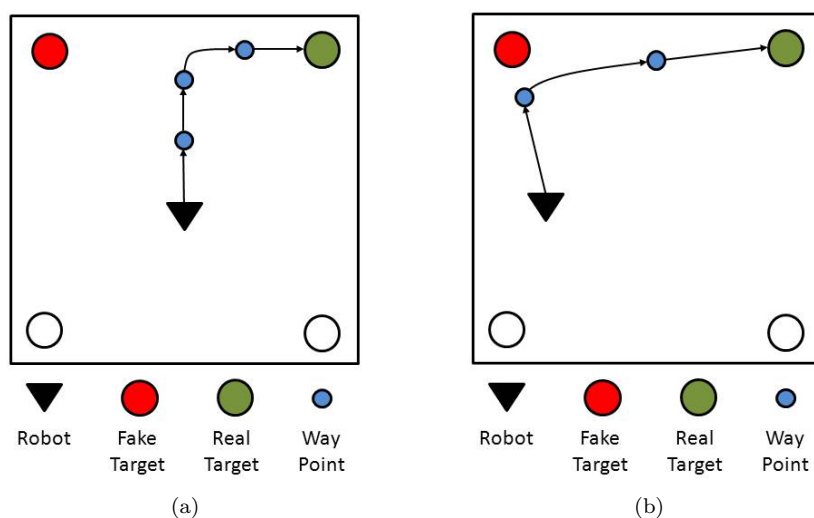


Fig. 4: The two types of deception when using the static trajectory approach. a) Moving forward and then changing the direction; b) Moving toward the false target and then changing direction.

The first one is activated when the robot is close to a point in the middle between the two targets. The robot will then move forward on the mid line, without letting the player know which tower it is going to aim to, leaving 50% of chance of guessing. Only when “close”, within a predefined threshold, to the midpoint of a virtual line connecting the two towers (true and fake target, respectively), the robot takes a straight trajectory towards the true target. The second type of deception aims to communicate from the beginning the fake target in order to induce the player to run to a particular tower, thus taking spatial advantage to reach the real target.

In this proposal, deception is implemented as a complete path to follow, which starts when the robot receives the deception information (true and false targets) and lasts until the robot hits the tower. Since no look-ahead for determining the player position in the future is considered, it may happen that the player understands the deception and can block the robot, so preventing it to finish the deception procedure coherently. For this reason, the algorithm calculates the time required to produce the full deceptive path and, when something goes wrong (for example, the player blocks the robot), the estimated time expires, and the deception procedure is aborted. Once it happens, the main controller selects a real target with the usual method, until a new need for deception is detected.

4.2.2 Dynamic steering behavior approach

The second approach we propose to implement a deceiving trajectory is based on steering behaviors [36]. The paper by Reynolds proposes a force-based approach to guide an actor in a life-like and improvisational manner. Given a target, it will generate a force, either attractive or repulsive, based on its position with respect to the robot: by applying this force on the robot, it will be driven either towards or away from the target following a smooth path.

Being our robot holonomic, it has been represented as a point mass to calculate the results of the application of the forces generated by the steering behavior. This approach gives the possibility to dynamically change the robot response to forces by changing the kinematic properties of the representation during the calculation process.

Instead of planning the complete trajectory, only the point where we want to change the motion parameters of the robot and reveal the true target is computed, and set as a temporary target. The steering behavior framework drives the robot to reach this target following a trajectory depending on its initial velocity and position. The robot's dexterity is dynamically increased when it finally aims at its real target: when the way point is reached the virtual mass of the robot is lowered and the virtual force applied to it is increased, along with a slight increase of its maximum velocity, then the real target is set as the target to be reached. This procedure generates a sharp turn and an acceleration of the robot towards the real goal. A visualization of the parameter update effect is reported in Figure 5.

5 Evaluation

In this section, an evaluation of the proposed strategy in the mentioned game scenario is presented. The experimental setup is first introduced, stating the hypotheses and presenting the post-match survey used to evaluate them. The analysis and related discussion follow next.




Fig. 5: Once reached the way point the robot virtual mass is reduced and the virtual force exerted on is increased: this will produce an increment in robot velocity as well as a smooth bend in the real target direction.

5.1 Experimental setup

5.1.1 Hypotheses

The performed trials were aimed at evaluating the following hypotheses.

- Hyp1: The subjects consider the robot as a rational agent that aims at winning.
- Hyp2: Deception improves the amusement.
- Hyp3: Both the trajectory approaches can create a recognizable level of deception.
- Hyp4: The dynamic steering approach is more appealing than the static one.

5.1.2 Independent variables

Two parameters could be considered to control the general game difficulty: maximum speed and maximum acceleration of the robot. They have been fixed, respectively, to $0.7m/sec$ and $0.5m/sec^2$ for all the trials, so that they cannot influence the evaluation of the effect of deception. These specific values have been selected mainly to guarantee safety of the player.

The robot's behavior is described by a nominal variable, factor *Behavior*, with three levels: *basic*, for the behavior without any deception, *static*, and *dynamic*, corresponding to the two different implementations of deceptive behaviors.

5.2 Participants

We recruited 78 participants among the visitors of a science fair and students. Most of the experiments were conducted with young people, spanning from 18 to 33 years old. We aimed at allocating the number of people evenly among the three different versions of the game.

One third of subjects were asked to play a version of the game where the deceptive mechanism was disabled: in the next paragraphs this sample will be labelled as *Basic*. The remaining two thirds of subjects were assigned to the other two levels of the *Behavior* factor. Therefore, these samples will be respectively labelled as *Static* and *Dynamic*. We randomly assigned the participants to batches, and none of them was made aware of this allocation. Moreover, they were not informed about the possibility of deceptive behavior, although subjects might have seen other players from the same batch playing before them.

The distribution of subjects is described in table 1.

Table 1: Demographic characteristics of the recruited sample, stratified by deception strategy.

Version	Age Mean	Age variance	Gender		Total
			Male	Female	
Basic	23.1	3.2	23	3	26
Static	23.8	4.9	17	9	26
Dynamic	21.9	7.6	18	8	26

A brief about the game rules was given to each participant before the game, and then she/he was introduced in the field, facing the robot at the center of the area; the game was started when ready. Each game lasted about 90 seconds, with a large variability, from even less than one minute to more than two minutes.

After each match, a questionnaire was self-administered to each player. It included a demographic section about the subject and thirteen statements about the game. In this paper, we present results about only three of them, relevant to evaluate the specific hypotheses that we focus on. In particular, here we consider the level of agreement on a Likert scale from “Strongly disagree” to “Strongly agree” to the following statements: “The robot wanted to win”, “I had fun”, “The robot did some feints”. We used the other questions to evaluate features of the game design, ranging from its duration to its safety, and we will report about them in other papers. Notice that, given the high dynamicity of the game, we could not evaluate fun with psycho-physiological measures, such as heart-rate variability or skin conductance, which are also affected, in much larger measure, by muscular activity. Moreover, it was not possible to

analyze emotional expressions on the face of the subjects, which could hardly be detected during the play activity, given that they were running.

5.3 Statistical tools

In order to evaluate our hypotheses, the first step is to evaluate whether the three samples come actually from the same population. Addressing the warnings raised by Kruskal and Wallis themselves in their seminal paper [25] with respect to the presence of tied data and the difference in the empirical distributions of our results, we used a dense ranking approach (as implemented in the R [35] package `dp1yr` [47]) and performed a k-sample permutational test as implemented in the R package `coin` [20]. Thus, we tested whether the distribution of the level of agreement expressed by participants differs among the three deception groups using a χ^2 statistic, the Asymptotic General Independence Test, hereafter χ^2 *AGIT*. For the post-hoc analysis, we also performed permutational pairwise comparisons on ranked data using a \mathcal{Z} test statistic (Asymptotic General Independence Test, hereafter \mathcal{Z} *AGIT*) as implemented in the R package `rcompanion` [28]. We also used the False Discovery Rate approach in order to adjust p-values for multiple comparisons. The level of significance was set to $\alpha=0.05$.

To assess the effect size for the three groups analysis, we adopted the Freeman's θ [5] as implemented in the R package `rcompanion`. This is a measure of association which ranges from 0 to 1 and enables to assess the strength of the relation between a nominal character (in our case the deception strategy, three levels: *Basic*, *Static*, and *Dynamic*) and an ordinal one (here, the level of agreement, five levels: *strongly disagree*, *disagree*, *indifferent*, *agree*, *strongly agree*). Beside the point estimation, we report also 95% confidence interval for θ , as computed using a bootstrap procedure (R package `boot` [11]).

We also present a descriptive analysis considering joint frequencies related to selected item pairs in the administered questionnaire.

5.4 Experimental results

5.4.1 Win shares

In figure 6, we report the win rate for the games played. As we can see, the *Basic* version did not offer an appropriate level of challenge, as most of the players won. On the other hand, the *Dynamic* version resulted in an unbalanced experience in favour of the robot, as most people could not match its ability. The *Static* version succeeded in providing an even game, with an equal share of wins between the robot and the players, which is one of the main goals for a game in order to keep the player engaged, enjoying playing.

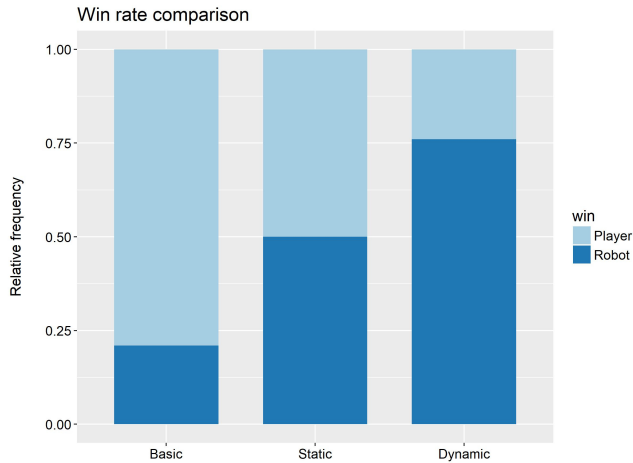


Fig. 6: Distribution of wins for each version of the game⁴.

5.4.2 Evaluating robot's rationality

In order to evaluate whether the participants perceived the robot as a rational agent aiming at winning the game (hypothesis *Hyp1*), it is possible to consider the agreement to the statement *The robot wanted to win* for the subjects facing the three different situations, reported in figure 7.

The result of the χ^2 *AGIT* gives $\chi^2_2 = 7.5257$, p-value = 0.02322, thus we can reject the null hypothesis that the three samples come from the same population.

When we consider the effect size, we find Freeman's $\theta = 0.238$ (95%ci = 0.0876 - 0.444) This is not very high, and the confidence interval is quite wide, so the association between membership to one the three groups and level of agreement expressed to the statement *The robot wanted to win* is not particularly strong.

It seems that most of the subjects agreed with the statement *The robot wanted to win* for both versions implementing deception, while for the *basic* version some subjects disagreed. A possible cause for this might be the fact that this version did not offer an adequate level of competition, possibly resulting in a game too easy for some people.

The statistical analysis about the pairwise differences among the groups, as reported in Table 2, shows that there is a statistically significant difference between the group facing *static* deception and the one facing the *basic* version, while the difference between *dynamic* and *basic* is not statistically significant, although the adjusted p-value is close to 0.05. Moreover, the difference between the two approaches to deception is not statistically significant.

⁴ This and the following plots have been realized with ggplot [46], and the R packages: grid [35], fontcm [13] and extrafont [12, 28]

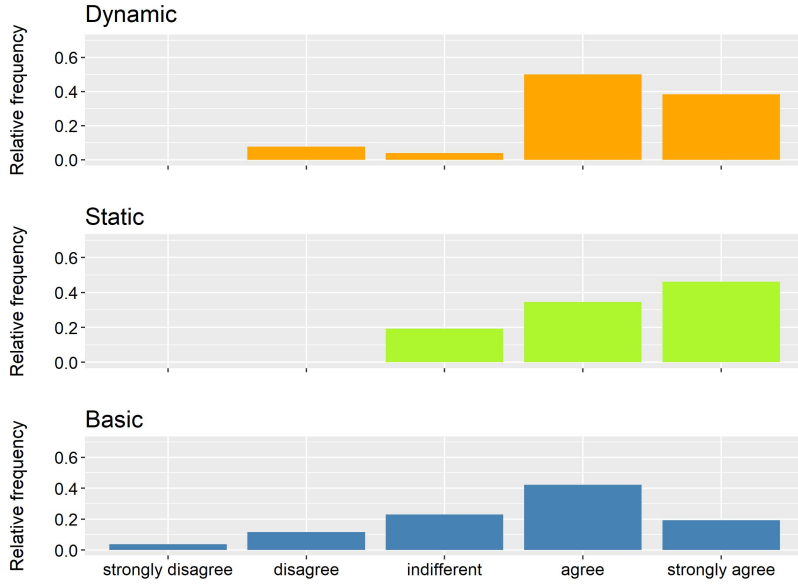


Fig. 7: Agreement level to the statement *The robot wanted to win* for the three groups.

The robot wanted to win		
	Z AGIT	Adjusted p-value
Dynamic vs Static	-0.3436	0.7312
Basic vs Static	2.09	0.04794
Basic vs Dynamic	2.409	0.05493

Table 2: Results for pairwise comparison on the agreement to the statement *The robot wanted to win*, between the *Basic*, *Static*, and *Dynamic* versions, respectively.

5.4.3 Evaluating the impact of deception on amusement

To evaluate hypothesis *Hyp2*, it is possible to compare the distributions of the agreement level to the statement *I had fun* of subjects playing the *Basic* version, i.e. without deception, with the same for subjects playing either the *Static* or *Dynamic* versions of the game, reported in Figure 8.

The result of the χ^2 AGIT gives $\chi^2_2 = 9.0854$, p-value = 0.01064, thus the three samples come from different populations.

When we consider the effect size, we find Freeman's $\theta = 0.22$ (95%ci = 0.0611 - 0.438). This is again not very high, and the confidence interval is quite wide, so the association between membership to one the three groups and level of agreement expressed to the statement *I had fun* is not particularly strong.




Fig. 8: Agreement level to the statement *I had fun* for the three groups.

The statistical analysis about the pairwise differences among the groups, as reported in Table 3, shows that there is a statistically significant difference between each of the groups with deception and the *basic* version, while the difference between *static* and *dynamic* is not statistically significant.

I had fun		
	Z_{AGIT}	Adjusted p-value
Dynamic vs Static	-0.2425	0.80840
Basic vs Static	2.266	0.03519
Basic vs Dynamic	2.445	0.03519

Table 3: Results for pairwise comparison on the agreement to the statement “I had fun”, between the *Basic*, *Static*, and *Dynamic* versions, respectively.

Results reported in Table 3 indicate that the groups receiving the two deceptive behaviors show a similar distribution for the level of agreement to the statement “I had fun”, whereas they both differ from the group receiving the basic behavior.

5.4.4 Evaluating the recognition of a deceiving strategy

In order to evaluate hypothesis *Hyp3*, that is to test whether our approach was able to provide a recognizable level of deception, we performed the statis-

tical analysis comparing the agreement level of subjects that played with the *Basic* version and subjects that played either the *Static* or *Dynamic* version on the *The robot did some feints* statement, whose distributions are reported in figure 9.

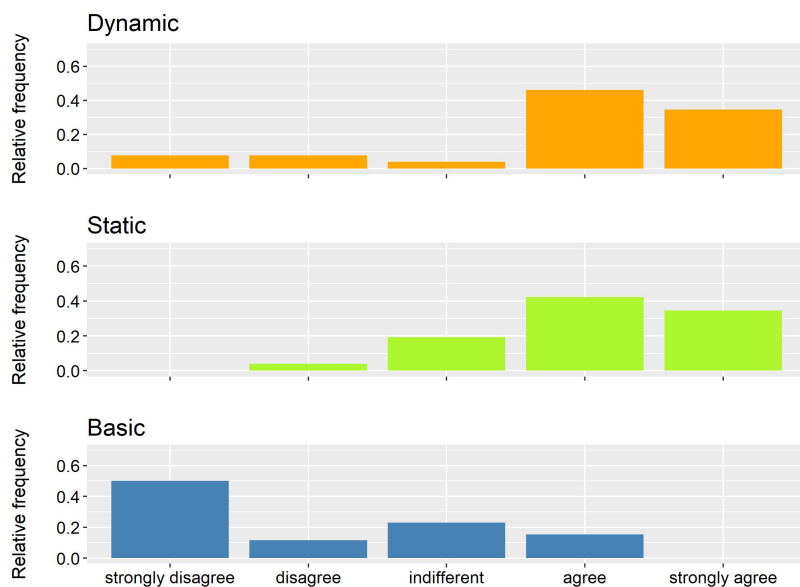


Fig. 9: Agreement level to the statement *The robot did some feints* for the three groups.

The result of the χ^2 *AGIT* gives $\chi^2_2 = 33.144$, p-value = 6.351×10^{-8} , thus the three samples come from different populations.

When we consider the effect size, we find Freeman's $\theta = 0.509$ (95%ci = 0.418 - 0.6830). This is higher than that in the previous cases, and the confidence interval is not as wide, so the association between membership to one of the three groups and level of agreement expressed to the statement "*The robot did some feints*" can be considered as stronger than in the previous cases, although still not too high. The results presented in table 4 show that, for both deceptive approaches, it is possible to reject the null hypothesis that the two groups, the one playing with deception and the other playing without it, provide the same perception of deceiving behavior, while the difference between the two types of deception w.r.t. this statement is not significant. In both cases, deception was perceived.

The robot did some feints		
	Z AGIT	Adjusted p-value
Dynamic vs Static	-0.5391	0.5898
Basic vs Static	4.487	1.085×10^{-5}
Basic vs Dynamic	5.078	1.142×10^{-6}

Table 4: Results for pairwise comparison on the agreement to the statement *The robot did some feints*, between the *Basic*, *Static*, and *Dynamic* versions, respectively.

5.4.5 Comparing Dynamic trajectory with Static

To validate Hypothesis *Hyp4* (*Is the dynamic approach better than the static one*), we considered agreement levels to the statements *I had fun* and *The robot did some feints*. As mentioned above, there is no significant difference between the two approaches, both in the fun level and in deception perception, so Hypothesis 4 has to be rejected.

5.5 Joint frequencies

We report in Figure 10 bubble plots of the joint frequencies of the agreement level of pairs of sentences, and discuss the interactions between the level of agreement between two statements. In the plots, the size of each bubble is proportional to the absolute joint frequency, reported in its center. Absolute frequencies, instead than relative frequencies, are reported to improve plot readability, given that each group has the same sample size, namely $N=26$.

From Figure 10 it is possible to notice that for subjects facing one or the other deception behaviors (on the two top rows), the ones that tended to recognize the presence of deception also declared to have had more fun. Moreover, these subjects perceived that the robot aimed at winning, with higher degree with respect to subjects that faced the behavior where deception was not present (lower row).

From Figure 11, focusing on the first column, it is possible to observe that subjects not facing a deceiving robot tended to expect to have more fun if deception was present, but also people facing, and detecting, static deception had liked to have some more. Considering the second column of Figure 11 we may notice that subjects facing a non-deceptive robot tended to recognize what the robot was about to do more than the other subjects, which tended to distribute their agreement around the middle range.

6 Discussion

Following the proposed approach to decide when and how produce deceptive behaviors in an interactive game, we designed a game that obtained a good acceptance by all the participants, most of which queued for participating and




Fig. 10: Joint frequencies of level of agreement to the sentences *I had fun*, *The robot wanted to win* *The robot did some feints* for the three groups. The color code is the same used in the previous figures: blue for *Basic*, green for *Static* and orange for *Dynamic*.

were pleased after the game. Most of the participants perceived the robot as a rational agent aiming at winning, as reported in Figure 7, although for some of them it was not offering an adequate level of competition, possibly due to relatively low ability, mostly for the *basic* version.

Although a kind of Hawthorne effect [22] could be expected in all the subjects, since they knew that they were participating to a test concerning their playing activity, hypothesis *Hyp1* is satisfied. In this paper, we are mostly interested in the difference between subjects facing the different behaviors of the robot to evaluate the effect of deception, and all the subjects might have suffered in similar measure from the Hawthorne effect, which can be considered as irrelevant for what reported this paper.

From the statistical analysis, it appears that the agreement to the statement *The robot wanted to win* is different with statistical significance only between the group facing the *static* deception and the *basic* behavior. This may seem to attribute to the static approach a preference a little bit higher with respect to the *dynamic* approach.




Fig. 11: Joint frequencies of level of agreement to the sentences *I would enjoy more if it did some feints*, *It was easy to understand what the robot would have done*, and *The robot did some feints* for the three groups.

Deception had a positive impact on the amusement level, with a statistically significant difference between both the deceptive behaviors and the basic one. We can state that hypothesis *Hyp2* is satisfied, although the magnitude of the improvement due to deception was possibly limited by the quite engaging nature of the game even without it.

Statistical results support well the hypothesis *Hyp3*, since both the deceiving behaviors were recognized as producing feints, differently from the basic behavior.

Hypothesis *Hyp4* that the dynamic trajectory would be a better characterization of deception, so producing a more enjoyable interaction, has to be rejected since for the statement *I had fun* there is no significant difference between the two types of deceptive behaviors. This may be due to their very basic nature. Actually, they implement different trajectories, but both can be classified by the subjects as belonging to the same conceptual category, i.e. as producing feints, and they played a similar role for producing fun.

6.1 Limitations

Among the limiting conditions of this work is the lack of a player model that can be used to tailor to the specific player the deceptive motion on line. For instance, Machine Learning algorithms can be used to model the player's actions and try to maximize the level of surprise [3] regarding tower attack.

Another limitation comes from constraints that prevent quick changes in motion direction, imposed for safety reasons, which, in turn, may produce negative impacts on the perception of deceptive movements.

7 Conclusion and Future Works

This paper shows as deception can play a relevant role in supporting the perception of a robot as a rational agent, at least in a setting, like a physically interactive game, where deception is expected as a feature characteristic of playful interaction.

We implemented a lively, enjoyable, physically interactive robogame, where the robot was perceived by the majority of the participants as a rational agent aiming at winning the game, even if its capabilities were strongly reduced with respect to its possible top performance. This is an example where effective human-robot interaction is obtained by implementing robots that are not behaving optimally, in the sense of minimizing some performance parameter, but that can match the expectation of interacting people to attribute rationality to the robot companion, which is one of the aims of human-robot interaction.

This result has been obtained thanks to two different deception algorithms, together with the system that triggers them by analyzing the game situation. These deceptive trajectories are not optimal to win, but obtain the desired effect of communicating a fake target to the player. The triggering mechanism activates deception only when this is functional to try to make the game even, which is another important aspect to ensure interest and engagement in the player [15].

The mechanism to decide whether and when applying a deception is based on general models, derived from Game Theory, that could be implemented for any competitive game, and, in general, for any interactive situation where it is possible to estimate the values of possible goals for the interacting agents, and the utility of a, possibly deceptive, action. In our case, the utility was oriented to produce similar rewards for the two interacting agents, in other cases the same framework could be adopted to obtain other effects, e.g., to maximize the reward of one of the two agents. Moreover, this model can be applied in any situation where deception is expected by a rational agent to support the perception of intelligent behavior.

In the next future, we plan to finalize a system to adapt the behavior and strategy of the robot to the perceived behavior, ability, and strategies of the human players, in real time, so to optimize their satisfaction and fun.

The need for this adaptation is highlighted by the lack of competitiveness some people reported in the questionnaire, mainly for the *basic* version of the game, and also by the different results obtained by the two implementations of deceptive strategies, selected by sound algorithms, but then run in open loop. By performing an online performance analysis it will be possible to adapt the parameters of the robot so to aim at playing in an even situation, with or without deception. Moreover, by performing a detailed analysis of the player's behavior and single actions, it would be possible to adapt the detailed way of playing and specific actions of the robot, so to give it a sort of character that would even more support the perception of the robot as a rational agent, personalized on the specific user.

Acknowledgements We would like to thank Roberta Carabalona for the statistical analysis of the collected data.

Compliance with Ethical Standards

The authors declare that this paper complies with the ethical standards of this Journal.

Funding

This work was partially supported by the *Brazilian National Council for Scientific and Technical Development* (CNPq) under the *Science Without Border* scholarship program number 203677/2014-5.

Conflict of Interest

The authors declare that they have no conflict of interest.

References

1. Hobby Products Report 2018 - Toys and games. Statista (2018)
2. World Service Robots 2018. International Federation of Robotics (2018)
3. Baldi, P.F., Itti, L.: Of bits and wows: A bayesian theory of surprise with applications to attention. *Neural Networks* **23**(5), 649–666 (2010)
4. 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 (2009)
5. Berry, K.J., Johnston, J.E., Mielke Jr, P.W.: A primer of permutation statistical methods. Springer **8**(490), 978–3 (2019)
6. Besio, S.: The need for play for the sake of play. In: S. Besio, D. Bulgarelli, V. Stancheva-Popkostadinova (eds.) *Play development in children with disabilities*, pp. 9–52. De Gruyter Publishers, Berlin, D (2018)
7. Besio, S., Caprino, F., Laudanna, E.: Profiling robot-mediated play for children with disabilities through icf-cy: the example of the european project iromec. In: *International Conference on Computers for Handicapped Persons*, pp. 545–552. Springer (2008)

8. Bonarini, A., Garzotto, F., Gelsomini, M., Romero, M., Clasadonte, F., Çelebi Yilmaz, A.N.: A huggable, mobile robot for developmental disorder interventions in a multimodal interaction space. In: *Proceedings of the 25th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN 2016)*, pp. 823–830. IEEE Computer Press, New York, NY, USA (2016)
9. Brooks, A.G., Gray, J., Hoffman, G., Lockerd, A., Lee, H., Breazeal, C.: Robot’s play: interactive games with sociable machines. *Computers in Entertainment (CIE)* **2**(3), 10–10 (2004)
10. Bruner, J.S., Jolly, A., Sylva, K.: *Play: Its role in development and evolution* (1976)
11. Canty, A., Ripley, B.D.: *boot: Bootstrap R (S-Plus) Functions* (2017). R package version 1.3-19
12. Chang, W.: *extrafont: Tools for using fonts* (2014). URL <https://CRAN.R-project.org/package=extrafont>. R package version 0.17
13. Chang, W., Kryukov, A., Murrell, P.: *fontcm: Computer Modern font for use with extrafont package* (2014). URL <https://CRAN.R-project.org/package=fontcm>. R package version 1.1
14. Cook, T.: <https://www.youtube.com/watch?v=Qk51skVkyg>. Accessed: 2019-05-19
15. Csikszentmihalyi, M.: *Finding flow: The psychology of engagement with everyday life*. Basic Books (1997)
16. Dragan, A.D., Holladay, R.M., Srinivasa, S.S.: An analysis of deceptive robot motion. In: *Robotics: science and systems*, p. 10. Citeseer (2014)
17. Fujita, M., Kageyama, K.: An open architecture for robot entertainment. In: *Proceedings of the first international conference on Autonomous agents*, pp. 435–442. ACM (1997)
18. Fukuda, H., Ueda, K.: Interaction with a moving object affects ones perception of its animacy. *International Journal of Social Robotics* **2**(2), 187–193 (2010)
19. Hoffman, G.: Anki, jibo, and kuri: What we can learn from social robots that didn’t make it. *IEEE Spectrum*, <https://spectrum.ieee.org/automaton/robotics/home-robots/anki-jibo-and-kuri-what-we-can-learn-from-social-robotics-failures?> (2019)
20. Hothorn, T., Hornik, K., van de Wiel, M.A., Zeileis, A.: Implementing a class of permutation tests: The coin package. *Journal of Statistical Software* **28**(8), 1–23 (2008). URL <http://www.jstatsoft.org/v28/i08/>
21. Ishak, N.I., Yusof, H.M., Naim Sidek, S., Rusli, N.: Robot selection in robotic intervention for asd children. In: *2018 IEEE-EMBS Conference on Biomedical Engineering and Sciences (IECBES)*, pp. 156–160. IEEE (2018)
22. Jones, S.R.: Was there a hawthorne effect? *American Journal of sociology* **98**(3), 451–468 (1992)
23. Kelley, H.H., Thibaut, J.W.: *Interpersonal Relations: A Theory of Interdependence*. New York: John Wiley Sons (1978)
24. Klein, T., Gelderblom, G.J., de Witte, L., Vanstipelen, S.: Evaluation of short term effects of the iromec robotic toy for children with developmental disabilities. In: *2011 IEEE international conference on rehabilitation robotics*, pp. 1–5. IEEE (2011)
25. Kruskal, W.H., Wallis, W.A.: Use of ranks in one-criterion variance analysis. *Journal of the American statistical Association* **47**(260), 583–621 (1952)
26. Lund, H.H., Klitbo, T., Jessen, C.: Playware technology for physically activating play. *Artificial life and Robotics* **9**(4), 165–174 (2005)
27. Lupetti, M.L., Piumatti, G., Rossetto, F.: Phygital play hri in a new gaming scenario. In: *2015 7th International Conference on Intelligent Technologies for Interactive Entertainment (INTEENTAIN)*, pp. 17–21. IEEE (2015)
28. Mangiafico, S.: *rcompanion: Functions to Support Extension Education Program Evaluation* (2017). URL <https://CRAN.R-project.org/package=rcompanion>. R package version 1.9.1
29. Martinoia, D., Calandriello, D., Bonarini, A.: Physically Interactive Robogames: Definition and design guidelines. *Robotics and Autonomous Systems* **61**(8), 739–748 (2013)
30. Michotte, A.: *The perception of causality*. Routledge (2017)
31. Oliveira, E., Orr, D., Nascimento, T., Bonarini, A.: Activity Recognition in a Physical Interactive RoboGame. In: *2017 Joint {IEEE} International Conference on Development and Learning and Epigenetic Robotics, ICDL-EpiRob 2017, Lisbon, Portugal, September 18-21, 2017*, pp. 1–6. Lisbon (2017)

32. Oliveira, E., Orr, D., Nascimento, T., Bonarini, A.: Modeling Player Activity in a Physical Interactive Robot Game Scenario. In: Proceedings of the 5th International Conference on Human Agent Interaction, HAI '17, pp. 411–414. ACM, New York, NY, USA (2017). DOI 10.1145/3125739.3132608. URL <http://doi.acm.org/10.1145/3125739.3132608>
33. Oliveira, E.L.S., Orrú, D., Morreale, L., Nascimento, T.P., Bonarini, A.: Learning and mining player motion profiles in physically interactive robogames. *Future Internet* **10**(3) (2018). DOI 10.3390/fi10030022. URL <http://www.mdpi.com/1999-5903/10/3/22>
34. Piaget, J.: Play, dreams and imitation in childhood, vol. 25. Routledge (2013)
35. R Core Team: R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria (2017). URL <https://www.R-project.org/>
36. Reynolds, C.: Steering behaviors for autonomous characters (1999)
37. Salter, T., Davey, N., Michaud, F.: Designing & developing queball, a robotic device for autism therapy. In: The 23rd IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN 2014), pp. 574–579. IEEE (2014)
38. Shim, J., Arkin, R.C.: A taxonomy of robot deception and its benefits in hri. In: Systems, Man, and Cybernetics (SMC), 2013 IEEE International Conference on, pp. 2328–2335. IEEE (2013)
39. Short, E., Hart, J., Vu, M., Scassellati, B.: No fair!! an interaction with a cheating robot. In: Human-Robot Interaction (HRI), 2010 5th ACM/IEEE International Conference on, pp. 219–226. IEEE (2010)
40. Terada, K., Ito, A.: Can a robot deceive humans? In: Proceedings of the 5th ACM/IEEE international conference on Human-robot interaction, pp. 191–192. IEEE Press (2010)
41. Vázquez, M., May, A., Steinfeld, A., Chen, W.H.: A deceptive robot referee in a multi-player gaming environment. In: Collaboration Technologies and Systems (CTS), 2011 International Conference on, pp. 204–211. IEEE (2011)
42. Vygotsky, L.S.: Play and its role in the mental development of the child. *Soviet psychology* **5**(3), 6–18 (1967)
43. Wagner, A.R., Arkin, R.C.: Analyzing social situations for human–robot interaction. *Interaction Studies* **9**(2), 277–300 (2008)
44. Wagner, A.R., Arkin, R.C.: Robot deception: recognizing when a robot should deceive. In: Computational Intelligence in Robotics and Automation (CIRA), 2009 IEEE International Symposium on, pp. 46–54. IEEE (2009)
45. Wagner, A.R., Arkin, R.C.: Acting deceptively: Providing robots with the capacity for deception. *International Journal of Social Robotics* **3**(1), 5–26 (2011)
46. Wickham, H.: *ggplot2: Elegant Graphics for Data Analysis*. Springer-Verlag New York (2009). URL <http://ggplot2.org>
47. Wickham, H., François, R., Henry, L., Müller, K.: *dplyr: A Grammar of Data Manipulation* (2019). URL <https://CRAN.R-project.org/package=dplyr>. R package version 0.8.0.1
48. Yannakakis, G.N.: How to model and augment player satisfaction: a review. In: First Workshop on Child, Computer and Interaction (2008)
49. Yannakakis, G.N., Hallam, J.: Real-time game adaptation for optimizing player satisfaction. *IEEE Transactions on Computational Intelligence and AI in Games* **1**(2), 121–133 (2009)
50. Yannakakis, G.N., Hallam, J., Lund, H.H.: Entertainment capture through heart rate activity in physical interactive playgrounds. *User Modeling and User-Adapted Interaction* **18**(1), 207–243 (2008). DOI 10.1007/s11257-007-9036-7
51. Yannakakis, G.N., Spronck, P., Loiacono, D., André, E.: Player modeling. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik (2013)