Do Humans Imitate Robots?

An Investigation of Strategic Social Learning in Human-Robot Interaction

Debora Zanatto**
School of Computer Science,
University of Plymouth
Plymouth, UK
debora.zanatto@bristol.ac.uk

Massimiliano Patacchiola†
School of Computer Science,
University of Plymouth
Plymouth, UK
massimiliano.patacchiola@ed.ac.uk

Jeremy Goslin
School of Psychology,
University of Plymouth
Plymouth, UK
jeremy.goslin@plymouth.ac.uk

Serge Thill
Donders Institute for Brain,
Cognition, and Behaviour
Radboud University Nijmegen
Nijmegen, The Netherlands
s.thill@donders.ru.nl

Angelo Cangelosi
Department of Computer Science,
University of Manchester
Manchester, UK
angelo.cangelosi@manchester.ac.uk

ABSTRACT

Theories on social learning indicate that imitative choices are usually performed whenever copying the others' behaviour has no additional cost. Here, we extended such investigations of social learning to Human-Robot Interaction (HRI). Participants played the Economic Investment Game with a robot banker while observing another robot player also investing in the robot banker. By manipulating the robot banker payoff, three conditions of unfairness were created: (1) unfair payoff for the participants, (2) unfair payoff for the robot player and (3) unfair payoff for both. Results showed that when the payoff was low for the participants and high for the robot player, participants invested more money in the robot banker than when both parties received a low return. Also, for this specific condition, participants' investments increased further with a more interactive robot player (defined as demonstrating increased attention, congruent movements and speech). This suggests that social and cognitive human competencies can be used and transposed to non-human agents. Further, imitation can potentially be extended to HRI, with interactivity likely having a key role in increasing this effect.

*Current address: Department of Computer Science, University of Bristol, Bristol, United Kingdom

†Current address: School of Informatics, University of Edinburgh, Edinburgh, United Kingdom

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

HRI '20, March 23–26, 2020, Cambridge, United Kingdom

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-6746-2/20/03...\$15.00

<https://doi.org/10.1145/3319502.3374776>

CCS CONCEPTS

- **Human-centered computing** → **Empirical studies in HCI**;
- **Applied computing** → **Psychology**.

KEYWORDS

Imitation; Social Learning; Human-Robot Interaction; Strategy; Investment Game

ACM Reference Format:

Debora Zanatto, Massimiliano Patacchiola, Jeremy Goslin, Serge Thill, and Angelo Cangelosi. 2020. Do Humans Imitate Robots? An Investigation of Strategic Social Learning in Human-Robot Interaction. In *Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction (HRI '20)*, March 23–26, 2020, Cambridge, United Kingdom. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3319502.3374776>

1 INTRODUCTION

1.1 Human imitation of others' strategies

Human beings naturally imitate, from newborns rapidly learning by observing others [33], to adults subconsciously mimicking partners' postures and mannerisms [15, 16]. Indeed, observing other persons' choices and decisions and the associated outcomes offers a rich source of information that individuals can use to improve their behaviour without direct experience, and thus also without the risks it might entail [2, 3].

Imitative behaviour is usually considered a form of social learning [e.g. 2, 4, 24, 34]. In social interactions, imitation is performed when perceived as the most convenient way to arrive at a specific goal [24]. In other words, humans usually modify their beliefs when they receive information about their decision. This information might derive from their direct experience (e.g. applying past successful strategies) or the observation of others' behaviour when facing the same problem and their outcomes. Hence, imitating someone else would generally be less costly than autonomously learning a new behaviour; in fact, it is often costless. This type of social learning can thus be a useful instrument to improve our

ability to choose the right behaviour, to apply the right problem-solving strategy, to increase the probability of a positive outcome and give the human access to several socially accepted behaviours [38, 47, 53].

Importantly, however, merely copying others indistinctly does not always result in producing the correct behaviour [11]. Instead, this type of imitation needs to be applied selectively, based on whether and when it would be more beneficial than to learn through direct experience.

Human strategies to imitate others have been previously studied. It has for example been argued that imitation plays an important role in real-world decision-making because it is one of the ways that allow the decision-maker to economise on decision costs [39]. In a series of five experiments to study this, participants faced computerised decision situations in which they had to select the amounts of three different inputs. The possibility to imitate other players has been manipulated throughout the experiments by setting different costs to access other players choices, as well as their reliability. It was found that participants increased their tendency to imitate when making unfamiliar decisions when they believe their strategy could be substantially improved, and when they became unsatisfied with the outcomes of their strategy. Another study [50] found that imitation is associated with others' high payoff in repeated choice problems, which can likely be explained by the use of an 'imitate the best' approach: simply choose the strategy that performed best among the observed actions in the previous period. A further study [29] found that the factors that people consider upon deciding when to imitate a behaviour include the productiveness of the established behaviour, the cost of learning, and their feeling of uncertainty.

1.2 Imitation in interaction with artificial agents

Human behaviours and decisions thus depend to some degree on the choices made by others, for example when they choose to imitate the strategy of others, and past research has studied this in human-human interaction. Presently, however, we are entering an age in which intelligent artificial agents become increasingly prevalent in society. For example, robots are used in both educational [27] and therapeutic settings with children [9, 20, 43]. Although the robot is often not completely autonomous, it is typically teleoperated in a manner to give the impression of autonomy (the so-called Wizard-of-Oz paradigm). It is thus realistic that a robot behaviour might be another source to imitate if it appears beneficial for humans.

In this paper, we therefore study whether humans are willing to imitate robots as they imitate other humans. In addition to having consequences for future societal roles for artificial agents – from assistive and medical applications to military roles and autonomous driving [e.g. 8, 14, 37] – this might lead to a novel measure of robot acceptance: while studies in HRI have previously addressed imitation, they have done so almost exclusively with a focus on the robot capability to imitate humans [48], such as whether and when to imitate people [12]. Although this has resulted in robots with increased capabilities and functionalities, the measurement of their acceptance is still confined to collecting people's judgements and opinions [e.g. 5, 18, 52]. By contrast, the experiment presented here also investigates whether imitative behaviours can be an indirect

measure of robot acceptance: does the willingness to reproduce a robot behaviour indicate it is perceived to be trustworthy and successful? As the literature on implicit measurements of robot perception in HRI is still limited, our study thus also contributes to the investigation of non-verbal response to robots as a measure of trust and acceptance.

1.3 Study aims and hypotheses

To investigate both the willingness to imitate and the potential as a measure of robot acceptance, we used the Economic Investment Game, a methodology derived from game theory that allows to implicitly measure trust over time. In the classic version of the game [10], participants in the role of the 'trustor' are given \$10 and can decide how much of this money to invest in the 'trustee'. The amount invested by the trustor is then tripled and given to the trustee, who can decide how much to return to the trustor. The amount of money that the trustor sends is considered an implicit measure of trust [49]. The Economic Investment Game has also been previously applied to HRI, focusing on the Uncanny Valley [30], on similarities to human-human investment behaviours [36] and whether the vocal features of a robot can affect participants' trust [45]. Most recently, the Economic Investment Game has also been used to evaluate cooperation with a robot [54].

In this study, participants were asked to play the Economic Investment Game with a humanoid robot banker, while also observing a second humanoid robot playing the same game. Specifically, we used a modified iterative version of the game following [44–46, 54] since an iterative version would allow the participants to perceive either the banker attitude and the robot player strategy and thus facilitate the emergence of imitative behaviours from the participants.

The payoff of the robot banker was manipulated so that participants perceived different degrees of fairness toward the two players: the banker payoff was programmed either to be low for both players, only for the human player or only for the robot player. The robot player was itself programmed to invest consistently with the banker behaviour and to maximise its income (such as playing high when generously rewarded and low on the opposite), and thereby appearing strategically competent to participants. It was hypothesised that an imitative tendency would arise when exposed to other's success, as described by [50]. Specifically, we expected participants to follow the costless 'imitate the best' rule strategy and copy the robot player strategy whenever the robot player was generously rewarded by the robot banker.

The key-dependent measure of this study was the monetary amount invested by the participants, predicted to be modified by the observation of the robot player success. If participants decided to copy the robot player following the 'imitate the best' rule, then their investments should rise significantly despite the banker payoff. If imitation is present, it could be used as an implicit indication of acceptance as it might indicate that participants consider the robot player a reliable agent. This, in turn, could be further evidence of trust-based behaviours toward the robot player.

In addition, this potential imitation effect could also provide a measure of trust in the banker. As [6] stated 'with the term "imitation" we refer to situations in which actors facing a trust problem

base their decision upon the behaviour of other trustors in similar conditions. If several other trustors trust a certain trustee, their behaviour can be perceived as a signal that trust can be placed safely, even though it is unknown whether this trustee honours trust. In other words, individuals sometimes decide to trust somebody just because they see others doing so'. In line with this argument, if the robot player is trusting and investing in the banker, its behaviour can be a signal to participants that the robot banker can be trusted.

Further, we did not expect participants to imitate the robot player when both players received low returns, or when only the robot player received low returns. In the former condition, the strategy is identical for both, while, whenever the participants are rewarded, imitating a player that is not rewarded would be detrimental.

This type of study further needs to take into account the social competences of the robot to be imitated. There is evidence from the literature that a socially competent robot can further improve the quality of the interaction with humans [e.g. 13, 17, 31]. In particular, enhancing social cues in non-human agents seems to elicit more social responses, improve trust and engagement [e.g. 21, 22, 51]. As suggestions from socially competent robots are perceived as more credible, acceptable and usable, we assumed that a more interactive robot could also be a more credible source of imitation. We therefore compared participants' investment during the game while exposed to an interactive robot player (engaging in joint attention behaviour, performing congruent movements, and speaking) and an immobile one. We expected an effect of the interactive behaviour in increasing imitation only when the robot player was successfully rewarded by the banker while the returns for the participants were low. Since there is no strategic reason for participants to imitate the robot player in the other conditions, we expected no effect in these.

2 METHODS

2.1 Participants

Eighty-nine individuals (66 female and 23 male; mean age = 19.81 years, SD = 2.55 years) participated in the study. Participants were equally assigned to three experimental groups: Unfair-for-human (19 female and 10 male, age mean=20.10, SD=2.49), Unfair-for-robot (18 female and 12 male, age mean= 20.09, SD=3.37) and Unfair-for-both (21 female and 9 male, age mean= 19.23, SD=2.16). No significant age differences between the three groups have been found ($p_s > .050$).

Participants have been recruited from the University of Plymouth School of Psychology Participation Pool. To exclude previous experience with robots, all the participants that previously participated in HRI experiments have been excluded. This has been accomplished by listing all the previous experiments (including those conducted by other researchers) as exclusion criteria. All the people listed as participants in those previous experiments, could not sign up to our experiment. All participants were naïve as to the purpose of the investigation and gave informed written consent to participate in the study. All participants were native British English speakers, they were right-handed and reported normal hearing and no language or neurological impairments. The study was approved by the South West—Central Bristol Research Ethics Committee.

2.2 Procedure

2.2.1 Investment game. The participants were seated next to one NAO robot (called "robot player") and a second NAO robot (called "robot banker") was standing in front of them. A touchscreen monitor was placed between the two parties (Figure 1).

Figure 1: Imitative Investment Game setup.

Participants were instructed to choose an amount of money to invest in the robot banker, knowing that the robot banker would then decide how much to return to them. The participants and the robot player were engaging in independent games with the robot banker, and the goal for both was to earn as much money as possible. Each player started with a notional sum of 10 Experimental Currency Units (ECU) at the beginning of every round of the game. The number of rounds was not known to the participants.

Each round of the game proceeded as follows: first both the participants and the robot player made their choices on how much to invest in the banker by indicating a number on the respective side of the screen; second, the two amounts were separately tripled and sent to the robot banker, which then separately returned some of the money to both players by declaring the payoff.

The display was split into two equal sections (Figure 2), each with 11 numbered touchscreen buttons which participants could press on their side to indicate the amount of money they wanted to invest in each round. Once both players made their investment choice, the respective numbers lit up in green. Participants could therefore see the robot player investment in the other half of the screen. The display further showed how much money the two players were making. These displays were named "bank" and were accumulating all the money each player earned separately.

The total amount earned by the participants at the end of each game was converted to the local currency, at a rate of 30 ECU = 0.10, and paid at the end of the experiment.

The Economic Investment Game was taken from [10] and modified to suit experimental needs. In particular, the amount of

Figure 2: Imitative Investment Game display.

money returned by the robot banker has been manipulated to give the impression of a generous or mean attitude as in [45].

Participants played two blocks of 20 rounds. In the first familiarisation block, the robot banker was programmed to be generous and rewarded both players by returning 50% to 80% of the received amount at each round. For each player, the robot banker followed a separate script so that the players did not receive the same amount in any given round.

In the second block, the robot banker followed three different game conditions, namely unfair-for-human, unfair-for-robot and unfair-for-both. In the unfair-for-human condition, the robot banker returned 50% to 80% to the robot player, and 0% to 30% to the participants. In the unfair-for-robot condition, the robot banker instead returned 0% to 30% to the robot player and 50% to 80% to the participants. In the unfair-for-both condition, the robot banker returned 0% to 30% to both players. The robot player, meanwhile, was programmed to play optimally given the robot banker payoff. Specifically, it was programmed to invest a number between 6 and 10 ECUs in the unfair-for-human condition, and between 0 and 4 ECUs otherwise. The robot player investments, as well as the robot banker payoff, followed a fixed script. The banker payoff followed the same scripts used by [44–46, 54].

In terms of behaviour, the robot player arms and head movements were manipulated, thus creating an interactive and an immobile behaviour. In the interactive condition, the robot player started every round by looking at the screen, pointing at it when making a choice and then turning its attention to the participants in case they had not taken any decision yet; it then looked at the robot banker and declared its choice, waiting for the robot banker to return the money and finally redirected the gaze to the display. In the immobile condition, the robot player always stared at the screen and never spoke.

To avoid any confounding effect, the robot banker remained immobile throughout the whole experiment. Moreover, the participants and the robots were left alone in the room, to avoid any unintentional interference by the experimenter during the game. Nevertheless, the experimenter could still monitor the game from a window placed behind the participants.

Given that the first block was equal for all participants, statistical analyses were only performed on the investment choices during the second block. The experiment was counterbalanced in a 2 (behaviour: interactive or immobile) by 3 (condition: unfair-for-human, unfair-for-robot, unfair-for-both) between-subject design.

2.2.2 Questionnaires. Four questionnaires were used as secondary measures to the main game task. Three short scales measured likeability, trust, and credibility. The likeability questionnaire was based on [41] and was composed of 10 items looking at friendliness, approachability, levels of knowledge, similarity to oneself, and agreeableness. The trust scale was an adaptation of the Receptivity/Trust subscale from the Relational Communication Questionnaire and of the selection of trust items in the International Personality Item Pool [23]. It included six items measuring factors like sincerity, honesty and trust. Likeability and trust questionnaires were structured as Likert-type scales with a range of values between 1 and 7. The credibility scale was based on the Source Credibility Scale [32] and included 12 items relating to the perception of intelligence, selfishness, expertise, and competence. This was structured as a semantic differential scale with a range of values between 1 and 7. In addition, the 24 items of the Bartneck et al's Godspeed questionnaire [7] was used to measure a range of HRI factors (anthropomorphism, animacy, likeability, perceived intelligence and perceived safety). This was structured as a semantic differential scale with a range of values between 1 and 5. Participants completed the questionnaires about the robot player after the investment games ended. They were then debriefed, paid the show-up fee plus any money earned in the game, and they left.

3 RESULTS

3.1 Investment game results

A linear mixed-effects model was fitted to the data using backwards stepwise selection, using the "lmerTest" package with the R statistical software [28], with participants' investment as the dependent variable, condition (unfair-for-both, unfair-for-human, unfair-for-robot), behaviour (interactive, immobile) and game turn (the 20 rounds of the game) as independent variables, and participants' id as random factor. Post-hoc comparisons were assessed using t-tests and Bonferroni's correction was applied when needed. The overall effect size of the model was $r^2 = 0.63$. For each predictor, estimate, 95% C.I., t-value, p-value and Cohen's d are reported in Table 1.

A significant main effect of condition has been found ($\chi^2_{(2)} = 195.09, p < .001$). Participants' investment was higher in the unfair-for-robot condition (mean = 8.53, SD = 1.96) when compared with both unfair-for-human (mean = 4.96, SD = 3.65, $p < .001$) and unfair-for-both (mean = 2.85, SD = 2.75, $p < .001$). Moreover, in the unfair-for-human condition participants were investing more than in the unfair-for-both condition ($p < .001$).

Whether or not the behaviour of the robot player was interactive did not affect the investment ($p > .05$), but a main effect for the game turn emerged ($\chi^2_{(1)} = 54.03, p < .001$), showing a decrease of investment over time.

A significant two-way interaction between condition and behaviour has also been found ($\chi^2_{(2)} = 7.23, p = .027$) (Figure 3). The unfair-for-human condition showed higher investments with the interactive behaviour (mean = 5.85, SD = 3.48) than with the immobile one (mean = 4.06, SD = 3.61, $p < .001$). Post-hoc comparisons did not show any other significant effect ($p_s > .05$).

Finally, a significant interaction between condition and game turn ($\chi^2_{(2)} = 138.85, p < .001$) has been found (Figure 4). Post-hoc

Predictor	Estimates	C.I.	t	p	d
<i>Condition-Unfair-for-Human</i>	2.29	1.07 - 3.53	3.66	< .001	0.67
<i>Condition-Unfair-for-Robot</i>	2.71	1.48 - 3.94	4.31	< .001	0.79
<i>Behaviour-Immobile</i>	-0.26	-1.40 - -0.88	-0.45	.651	-0.09
<i>Game turn</i>	-0.17	-0.20 - -0.14	-10.65	< .001	-0.51
<i>Condition-Unfair-for-Human:Game turn</i>	0.05	0.01 - 0.10	2.41	.002	0.11
<i>Condition-Unfair-for-Robot:Game turn</i>	0.25	0.21 - 0.30	11.19	< .001	0.54
<i>Condition-Unfair-for-Human:Behaviour-Immobile</i>	-1.53	-3.14 - -0.08	-1.86	.065	-0.39
<i>Condition-Unfair-for-Robot:Behaviour-Immobile</i>	0.62	-0.99 - -2.23	0.75	.455	0.15

Table 1: Regression output. For each predictor, estimate, 95% confidence interval, t value, p value and Cohen's d are reported.

Figure 3: Participants' investment in the three different conditions (unfair-for-both, unfair-for-human and unfair-for-robot) with interactive and immobile robot player.

analyses revealed a decrease of investment over time for the unfair-for-human ($\chi^2_{(1)} = 32.13, p < .001$) and unfair-for-both ($\chi^2_{(1)} = 115.33, p < .001$) conditions, but the investment in the former was higher than in the latter; on the opposite, participants' investment was constantly increasing for the unfair-for-robot condition ($\chi^2_{(1)} = 70.7, p < .001$).

3.2 Questionnaires Results

For all the scales, a linear mixed-effects model was fitted to the data using backwards stepwise selection, with scale as the dependent

Figure 4: Participants' investment over time in the three different conditions (unfair-for-both, unfair-for-human, unfair-for-robot)

variable, condition (unfair-for-both, unfair-for-human, unfair-for-robot) and behaviour (interactive, immobile) as independent variables, and participants' id as random factor. Post-hoc comparisons were assessed using t-tests and Bonferroni's correction was applied when needed.

As reported in Table 2, the type of condition affected Likeability and Trust questionnaires, Godspeed Anthropomorphism, Likeability, Intelligence and Safety scales. Likeability, Godspeed Likeability and Safety had higher ratings for the unfair-for-robot condition ($p_s < .050$). Trust, Anthropomorphism and Intelligence were higher in the unfair-for-human condition ($p_s < .050$). Behaviour affected the Godspeed Animacy scale, where the interactive robot was preferred to the immobile one ($p < .001$).

Finally, there was a significant two-way interaction between condition and behaviour that showed higher preferences for the interactive robot player over the immobile one in the unfair-for-human condition for Likeability ($p = .003$), Credibility ($p = .003$) and Godspeed Animacy ($p = .004$) and Intelligence scales ($p = .045$), as shown in Figure 5.

	r^2	Condition	Behaviour	Condition by Behaviour
<i>Likeability</i>	0.27	$\chi^2_{(2)} = 6.27, p = .043$	n.s.	$\chi^2_{(2)} = 15.75, p < .001$
<i>Trust</i>	0.22	$\chi^2_{(2)} = 15.98, p < .001$	n.s.	n.s.
<i>Credibility</i>	0.18	n.s.	n.s.	$\chi^2_{(2)} = 16.09, p < .001$
Godspeed Questionnaires				
<i>Anthropomorphism</i>	0.24	$\chi^2_{(2)} = 28.28, p < .001$	n.s.	n.s.
<i>Animacy</i>	0.26	n.s.	$\chi^2_{(1)} = 13.00, p < .001$	$\chi^2_{(2)} = 6.03, p = .049$
<i>Likeability</i>	0.27	$\chi^2_{(2)} = 26.24, p < .001$	n.s.	n.s.
<i>Intelligence</i>	0.80	$\chi^2_{(2)} = 5.95, p = .050$	n.s.	$\chi^2_{(2)} = 16.69, p < .001$
<i>Safety</i>	0.19	$\chi^2_{(2)} = 6.19, p = .045$	n.s.	n.s.

Table 2: Questionnaires analysis results, regression output. For each questionnaire, the effect size of the model and the Chi-square for each significant effect are reported.

	Condition			Behaviour	
	Unfair for Both	Unfair for Human	Unfair for Robot	Interactive	Immobile
<i>Likeability</i>	3.75(0.89)	3.71(0.87)	4.15(0.75)	4.00(0.74)	3.73(0.96)
<i>Trust</i>	2.51(0.91)	3.45(1.14)	2.54(1.05)	2.84(1.13)	2.81(1.10)
<i>Credibility</i>	4.55(0.41)	4.71(0.64)	4.69(0.68)	4.71(0.57)	4.57(0.60)
Godspeed Questionnaires					
<i>Anthropomorphism</i>	2.08(0.55)	2.71(0.56)	2.02(0.57)	2.24(0.72)	2.28(0.54)
<i>Animacy</i>	2.71(0.52)	2.88(0.64)	2.95(0.56)	3.03(0.54)	2.63(0.55)
<i>Likeability</i>	3.57(0.38)	3.66(0.57)	4.14(0.44)	3.86(0.54)	3.71(0.52)
<i>Intelligence</i>	3.73(0.33)	3.79(0.38)	3.58(0.41)	3.69(0.37)	3.71(0.39)
<i>Safety</i>	3.92(0.34)	4.10(0.49)	4.20(0.48)	4.08(0.38)	4.08(0.53)

Table 3: Questionnaires ratings. For each questionnaire, participants' mean ratings and SD in each condition and behaviour types are reported.

4 DISCUSSION

In this paper, imitation has been studied in the context of HRI, in order to see whether humans can imitate robots behaviour in the same way they do with other humans. Participants were asked to invest in a robotic banker while observing another robot playing the same game. The banker's generosity in returning the money invested was manipulated, creating three different unfair scenarios (unfair-for-both, unfair-for-human, unfair-for-robot). Participants were expected to apply the 'imitate the best' rule and thus to imitate the robot player strategy whenever the payoff was high for the robot player.

The results reported here show that humans can strategically imitate robots. The amount of money that participants invested in the robot banker was significantly higher when the payoff for the robot player was high than when both parties received low returns.

Even though the optimal strategy, for the participants, would be to reduce the investment to the minimum in both conditions, the observation of the robot player's higher returns increased their investment in that condition. The questionnaires were also provided some support for the behavioural data: the robot player in the unfair-for-human condition was rated more trustworthy, anthropomorphic and intelligent.

Nevertheless, participants did not fully imitate the robot strategy. Although their investment was significantly higher in the unfair-for-human than the unfair-for-both condition, participants did not adapt their strategy to an optimal level and invested significantly less than the robot player ($t_{(733)} = 22.52, p < .001$). Moreover, participants' investment in the unfair-for-human condition gradually decreased towards the end of the game, almost reaching the same investment chosen in the unfair-for-both condition. This reduction of investment could be affected by the structure of the game.

Figure 5: Participants' ratings in the three different conditions (unfair-for-both, unfair-for-human and unfair-for-robot) and in the interactive and immobile behaviours, for Likeability, Credibility, Animacy and Intelligence scales.

Once participants realised that copying the robot player strategy would not increase their payoff, they might have applied a different strategy, thus reducing their investment.

These results are in line with previous literature on human and animal social learning [25, 26, 29, 35, 40], which found that productiveness, cost of learning and uncertainty are the main contributors in deciding whether to imitate a behaviour or not. Following someone's successful behaviour, like in the unfair-for-human condition, reduces uncertainty and increases the chances of obtaining a positive outcome. When instead the other's behaviour results in non-positive outcomes, imitation is not a suitable costless option. Here, the high investment observed in the unfair-for-robot condition showed that participants were able to selectively identify in which condition it was useful to imitate the robot player's strategy.

An alternative explanation of the imitative effect can be related to the participants' willingness to please the robot banker. In an attempt to persuade the banker to change its attitude and increase the payoff, participants might have decided to copy the robot player's strategy. Following this interpretation, the questionnaires ratings could be an a posteriori attempt to rationalise their strategical imitation. That is, participants trusted the robot player more to justify their imitative attitude. Considering these two alternative explanations, it would be advisable to explore further this triadic interaction and also investigate the role of the banker in affecting this strategical social learning effect.

Participants' performance next to a more or less interactive robot player has also been investigated. This was done under the premise that a more socially competent robot player would be perceived as a more valid source of imitation. Results showed that the type of behaviour had no main effect, but the two-way interaction between condition and behaviour indicated that participants increased their investment in the unfair-for-human condition whenever the robot player was more interactive. In conditions where participants were playing by their own rules, they were also not interested in the robot behaviour and performance (since this was not helpful to them). When participants noticed that a better payoff was possible (by observing the gains of the robot player), they not only looked to what the robot player was doing but also noticed whether its behaviour was more or less human-like.

It has been theorised that people tend to prefer a more human-like agent to reduce the frustration and the anxiety associated with the uncertainty of the interaction with the agent and establish a feeling of belonging [19]. Here, participants might thus have increased their tendency to imitate a more interactive robot player in response to the need for establishing a connection and reduce uncertainty. Moreover, as imitation is usually human-human based, it is possible that social learning became easier when paired a socially competent robot.

The questionnaires also support this explanation: the anthropomorphism scale reported higher ratings for the robot player in the

unfair-for-human, thus indicating that participants were actively engaged and focused on the robot player in that condition. Moreover, the two-way interaction between condition and behaviour confirmed the behavioural results: participants rated the interactive robot player more likeable, credible, animated and intelligent than the immobile version in the unfair-for-human condition only. A robot showing social competences and engagement capabilities thus increased the probability that participants accepted its strategy as a valuable and successful alternative. Nevertheless, these behavioural improvements in the robot behaviour were useless complements in the remaining conditions.

5 CONCLUSION

Taken together our results not only lay the foundation to the exploration of social learning between humans and robots, but they have also shown that this phenomenon can be used as a reliable implicit measure of the evolution of social processes in HRI. While research in HRI has predominantly focused on producing and designing robots capable of interacting with humans, our results demonstrate the social impact robots can have in the way humans interact with them. Moreover, evaluations of the effect of social robots have mostly been carried out using explicit measurements. An implicit evaluation such as demonstrated here could be instrumental in future research on how humans subconsciously perceive, treat, and accept robotic agents.

The present results demonstrate that humans can potentially apply social learning strategies to non-human agents and this disposition can be notably increased when the agent is showing social competences. This has important implications in the field of social cognition and offers a new experimental paradigm for investigating social learning. The limited capabilities of a robotic agent deliver an ideal experimental scenario for identifying under which conditions humans activate social processes. This, in turn, could lead to a clearer comprehension of these processes.

Although some experimental limitations of course exist, it appears that a robot is potentially capable of influencing a human decision just by showing the benefits that can be obtained by following its instructions. This offers a wide range of improvements in every field covered by HRI. As social learning can stimulate the production of socially desired behaviours from models, robots could become one of these models in helping to generate social change. This is true in particular for health application (e.g. eliminating risk factors such as poor diet, inactivity, improving adherence to a prescribed regimen, etc.) and education. For example, [1] observed that exposing children to a more creative social robotic peer improved their creative thinking. Further, exposure to a robot vulnerable behaviour can increase a human engagement with the robot, and also enhances tension reduction behaviours (explaining mistakes, laughing, consolation) in human teams [42].

Future studies, however, should further investigate the role of the banker. As a social learning effect has been found, it is not completely clear whether participants' increase of investment was derived by the mere observation of the robot player strategy or by any attempt to please the banker. Therefore, research should focus also on investigating participants' perception of the banker

through questionnaires. Moreover, it would be appropriate to expand this scenario by substituting a robot banker with a human banker. Would participants choose to imitate the robot player in front of a human banker? Would the robot player be perceived as an 'anthropomorphic outlier' in this kind of setting and reduce participants' imitative tendency? These types of question would lead to a better understanding of the role of social learning in HRI

6 ACKNOWLEDGMENTS

This material is based upon work supported by the Air Force Office of Scientific Research, Air Force Materiel Command, USAF under Award No, FA9550-19-1-7002 (THRIVE++).

REFERENCES

- [1] Safinah Ali, Tyler Moroso, and Cynthia Breazeal. 2019. Can Children Learn Creativity from a Social Robot?. In *Proceedings of the 2019 on Creativity and Cognition*. ACM, 359–368.
- [2] Albert Bandura. 1962. *Social learning through imitation*. Univer. Nebraska Press.
- [3] Albert Bandura. 1978. Social learning theory of aggression. *Journal of communication* 28, 3 (1978), 12–29.
- [4] Albert Bandura and Richard H Walters. 1963. Social learning and personality development. (1963).
- [5] Jaclyn Barnes, Maryam FakhrHosseini, Myounghoon Jeon, Chung-Hyuk Park, and Ayanna Howard. 2017. The influence of robot design on acceptance of social robots. In *2017 14th International Conference on Ubiquitous Robots and Ambient Intelligence (URAI)*. IEEE, 51–55.
- [6] Davide Barrera and Vincent Buskens. 2007. Imitation and learning under uncertainty: a vignette experiment. *International sociology* 22, 3 (2007), 367–396.
- [7] Christoph Bartneck, Dana Kulić, Elizabeth Croft, and Susana Zoghbi. 2008. Measurement Instruments for the Anthropomorphism, Animacy, Likeability, Perceived Intelligence, and Perceived Safety of Robots. *International Journal of Social Robotics* 1, 1 (2008), 71–81. <https://doi.org/10.1007/s12369-008-0001-3>
- [8] Chandrayee Basu and Mukesh Singhal. 2016. Trust dynamics in human autonomous vehicle interaction: a review of trust models. In *2016 AAAI Spring Symposium Series*.
- [9] Tony Belpaeme and Anthony Morse. 2012. Word and category learning in a continuous semantic domain: Comparing cross-situational and interactive learning. *Advances in Complex Systems* 15, 03n04 (2012), 1250031.
- [10] Joyce Berg, John Dickhaut, and Kevin McCabe. 1995. Trust, reciprocity, and social history. *Games and economic behavior* 10, 1 (1995), 122–142.
- [11] Robert Boyd and Peter J Richerson. 1988. *Culture and the evolutionary process*. University of Chicago press.
- [12] Cynthia Breazeal and Brian Scassellati. 2002. Robots that imitate humans. *Trends in cognitive sciences* 6, 11 (2002), 481–487.
- [13] Cynthia Breazeal, Atsuo Takanishi, and Tetsunori Kobayashi. 2008. Social robots that interact with people. *Springer handbook of robotics* (2008), 1349–1369.
- [14] Andrew Brodie and Nikhil Vasdev. 2018. The future of robotic surgery. *The Annals of The Royal College of Surgeons of England* 100, Supplement 7 (2018), 4–13.
- [15] Tanya L Chartrand and John A Bargh. 1999. The chameleon effect: the perception-behavior link and social interaction. *Journal of personality and social psychology* 76, 6 (1999), 893.
- [16] Tanya L Chartrand, Clara Michelle Cheng, and Valerie E Jefferis. 2002. You're just a chameleon: The automatic nature and social significance of mimicry. *Natura automatyzmow (Nature of Automaticity)* (2002), 19–24.
- [17] Kerstin Dautenhahn. 2007. Socially intelligent robots: dimensions of human-robot interaction. *Philosophical Transactions of the Royal Society of London B: Biological Sciences* 362, 1480 (2007), 679–704.
- [18] Maartje MA De Graaf and Somaya Ben Allouch. 2013. Exploring influencing variables for the acceptance of social robots. *Robotics and Autonomous Systems* 61, 12 (2013), 1476–1486.
- [19] Nicholas Epley, Adam Waytz, and John T. Cacioppo. 2007. On seeing human: a three-factor theory of anthropomorphism. *Psychological Review* 114, 4 (2007), 864–886.
- [20] Pablo G Esteban, Paul Baxter, Tony Belpaeme, Erik Billing, Haibin Cai, Hoang-Long Cao, Mark Coeckelbergh, Cristina Costescu, Daniel David, Albert De Beir, et al. 2017. How to build a supervised autonomous system for robot-enhanced therapy for children with autism spectrum disorder. *Paladyn, Journal of Behavioral Robotics* 8, 1 (2017), 18–38.
- [21] Kerstin Fischer, Lars Christian Jensen, Franziska Kirstein, Sebastian Stabinger, Özgür Erkent, Dadihichi Shukla, and Justus Piater. 2015. The effects of social

- gaze in human-robot collaborative assembly. In *International Conference on Social Robotics*. Springer, 204–213.
- [22] Aimi Shazwani Ghazali, Jaap Ham, Emilia Barakova, and Panos Markopoulos. 2019. Assessing the effect of persuasive robots interactive social cues on users' psychological reactance, liking, trusting beliefs and compliance. *Advanced Robotics* 33, 7–8 (2019), 325–337.
- [23] Lewis R. Goldberg, John A. Johnson, Herbert W. Eber, Robert Hogan, Michael C. Ashton, C. Robert Cloninger, and Harrison G. Gough. 2006. The international personality item pool and the future of public-domain personality measures. *Journal of Research in personality* 40, 1 (2006), 84–96.
- [24] Peter Hedström, Richard Swedberg, Gudmund Hernes, et al. 1998. *Social mechanisms: An analytical approach to social theory*. Cambridge University Press.
- [25] William Hoppitt and Kevin N Laland. 2013. *Social learning: an introduction to mechanisms, methods, and models*. Princeton University Press.
- [26] Rachel L Kendal, Neeltje J Boogert, Luke Rendell, Kevin N Laland, Mike Webster, and Patricia L Jones. 2018. Social learning strategies: Bridge-building between fields. *Trends in cognitive sciences* 22, 7 (2018), 651–665.
- [27] James Kennedy, Paul Baxter, and Tony Belpaeme. 2017. The impact of robot tutor nonverbal social behavior on child learning. *Frontiers in ICT* 4 (2017), 6.
- [28] Alexandra Kuznetsova, Per B Brockhoff, and Rune Haubo Bojesen Christensen. 2017. lmerTest package: tests in linear mixed effects models. *Journal of Statistical Software* 82, 13 (2017).
- [29] Kevin N Laland. 2004. Social learning strategies. *Animal Learning & Behavior* 32, 1 (2004), 4–14.
- [30] Maya B Mathur and David B Reichling. 2016. Navigating a social world with robot partners: A quantitative cartography of the Uncanny Valley. *Cognition* 146 (2016), 22–32.
- [31] Nikolaos Mavridis. 2015. A review of verbal and non-verbal human-robot interactive communication. *Robotics and Autonomous Systems* 63 (2015), 22–35.
- [32] James C. McCroskey and Thomas J. Young. 1981. Ethos and credibility: The construct and its measurement after three decades. *Communication Studies* 32, 1 (1981), 24–34.
- [33] Andrew N Meltzoff and Peter J Marshall. 2018. Human infant imitation as a social survival circuit. *Current Opinion in Behavioral Sciences* 24 (2018), 130–136.
- [34] Neal Elgar Miller and John Dollard. 1941. Social learning and imitation. (1941).
- [35] TJH Morgan, LE Rendell, Micael Ehn, W Hoppitt, and KN Laland. 2012. The evolutionary basis of human social learning. *Proceedings of the Royal Society of London B: Biological Sciences* 279, 1729 (2012), 653–662.
- [36] Roberta C Ramos Mota, Daniel J Rea, Anna Le Tran, James E Young, Ehud Sharlin, and Mario C Sousa. 2016. Playing the 'trust game' with robots: Social strategies and experiences. In *2016 25th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*. IEEE, 519–524.
- [37] Robin R Murphy. 2004. Human-robot interaction in rescue robotics. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)* 34, 2 (2004), 138–153.
- [38] Andreas Olsson and Elizabeth A Phelps. 2007. Social learning of fear. *Nature neuroscience* 10, 9 (2007), 1095.
- [39] Mark Pingle. 1995. Imitation versus rationality: An experimental perspective on decision making. *The Journal of Socio-Economics* 24, 2 (1995), 281–315.
- [40] Luke Rendell, Robert Boyd, Daniel Cownden, Marquist Enquist, Kimmo Eriksson, Marc W Feldman, Laurel Fogarty, Stefano Ghirlanda, Timothy Lillicrap, and Kevin N Laland. 2010. Why copy others? Insights from the social learning strategies tournament. *Science* 328, 5975 (2010), 208–213.
- [41] Stephen Reysen. 2005. Construction of a new scale: The Reysen likability scale. *Social Behavior and Personality: an international journal* 33, 2 (2005), 201–208.
- [42] Sarah Strohkorb Sebo, Margaret Traeger, Malte Jung, and Brian Scassellati. 2018. The ripple effects of vulnerability: The effects of a robot's vulnerable behavior on trust in human-robot teams. In *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*. ACM, 178–186.
- [43] Serge Thill, Cristina A Pop, Tony Belpaeme, Tom Ziemke, and Bram Vanderborght. 2012. Robot-assisted therapy for autism spectrum disorders with (partially) autonomous control: Challenges and outlook. *Paladyn* 3, 4 (2012), 209–217.
- [44] Ilaria Torre, Jeremy Goslin, and Laurence White. 2015. Investing in accents: How does experience mediate trust attributions to different voices?. In *ICPhS*.
- [45] Ilaria Torre, Jeremy Goslin, Laurence White, and Debora Zanatto. 2018. Trust in artificial voices: A congruency effect of first impressions and behavioural experience. In *Proceedings of the Technology, Mind, and Society*. ACM, 40.
- [46] Ilaria Torre, Laurence White, and Jeremy Goslin. 2016. Behavioural mediation of prosodic cues to implicit judgements of trustworthiness. In *Proceedings of Speech Prosody*. 816–820.
- [47] Ruud Van Den Bos, Jolle Jolles, and Judith Homberg. 2013. Social modulation of decision-making: a cross-species review. *Frontiers in Human Neuroscience* 7 (2013), 301.
- [48] Dieter Vanderelst and Alan Winfield. 2017. Rational imitation for robots: the cost difference model. *Adaptive Behavior* 25, 2 (2017), 60–71.
- [49] Mascha Van't Wout and Alan G. Sanfey. 2008. Friend or foe: The effect of implicit trustworthiness judgments in social decision-making. *Cognition* 108, 3 (2008), 796–803.
- [50] Fernando Vega-Redondo. 1997. The evolution of Walrasian behavior. *Econometrica* 65, 2 (1997), 375–384.
- [51] Frank MF Verberne, Jaap Ham, Aditya Ponnada, and Cees JH Midden. 2013. Trusting digital chameleons: The effect of mimicry by a virtual social agent on user trust. In *International Conference on Persuasive Technology*. Springer, 234–245.
- [52] Katie Winkle, Praminda Caleb-Solly, Ailie Turton, and Paul Bremner. 2018. Social robots for engagement in rehabilitative therapies: Design implications from a study with therapists. In *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*. ACM, 289–297.
- [53] Thomas N Wisdom, Xianfeng Song, and Robert L Goldstone. 2013. Social learning strategies in networked groups. *Cognitive Science* 37, 8 (2013), 1383–1425.
- [54] Debora Zanatto, Massimiliano Patacchiola, Jeremy Goslin, and Angelo Cangelosi. 2019. Investigating cooperation with robotic peers. *PLoS one* 14, 11 (2019), e0225028.