Fitting the Room: Social Motivations for Context-Aware Agents

Diogo Rato
diogo.rato@tecnico.ulisboa.pt
GAIPS, INESC-ID & Instituto Superior Técnico, Universidade de Lisboa
Portugal

Marta Couto
marta.couto@gaips.inesc-id.pt
GAIPS, INESC-ID
Portugal

Rui Prada
rui.prada@tecnico.ulisboa.pt
GAIPS, INESC-ID & Instituto Superior Técnico, Universidade de Lisboa
Portugal

ABSTRACT

Social agents should exhibit socially adequate behavior to fit the context they meet. Fitting the context is particular relevant for interactive agents that interact and are being observed by people. Hence, the perceptions of people of such social capabilities are an important concern. Exhibiting socially adequate behavior can more easily be identifiable when in the presence of other social actors. However, even alone, one's ability to adjust to the context might be socially motivated and interpreted as such. Similarly, intelligent agents may be identified as social beings when acting alone. Moreover, social context is triggered in different ways. In this study, we explore if adaptation to the physical surroundings (e.g., the agent's location) is enough to shape the perceptions of people observing the agent. We contribute to the study of situated cognition's role in interpreting an autonomous agent's behavior. In particular, we explore the impact of behavior changes grounded on the location as a contextual cue on the motivation ascribed by an observer to the agent's behavior.

We implemented a virtual scenario with multiple contexts and one simple character employing a computational model called Cognitive Social Frames that supports behavior change to context. We conducted a user study (n=92) to assess if an observer's perceptions of intention and motivation are affected by an agent's capability to adapt to different contexts. Our findings suggest that (a) despite no other agents being present, participants ascribe social motivations to the agent's adaptive behavior, (b) such attributions are independent of visual cues, and (c) even without any pre-established norms, agents that consistently adjust their behavior to the physical context are perceived as more social.

CCS CONCEPTS

• **Human-centered computing** → *Empirical studies in HCI*; *User studies*; • **Computing methodologies** → Agent / discrete models.

KEYWORDS

User Studies, Social Context, Social Motivations, Virtual Agents

HAI '21, November 9–11, 2021, Virtual Event, Japan

ACM ISBN 978-1-4503-8620-3/21/11.
https://doi.org/10.1145/3472307.3484185

ACM Reference Format:
Diogo Rato, Marta Couto, and Rui Prada. 2021. Fitting the Room: Social Motivations for Context-Aware Agents. In *Proceedings of the 9th International Conference on Human-Agent Interaction (HAI '21), November 9–11, 2021, Virtual Event, Japan.* ACM, New York, NY, USA, 8 pages. https://doi.org/10.1145/3472307.3484185

1 INTRODUCTION

Social context corresponds to the "specific circumstance or general environment that serves as a social framework for individual or interpersonal behavior. This context frequently influences, at least to some degree, the actions and feelings that occur within it" [43]. Some psychologists have noted the influence of the physical surroundings on the human social context. For example, in 1913, in his book, "On beginning the treatment'" Freud describes the therapeutic setting for psychoanalysis including, amongst other things, the importance of the position of the patient (lying down) and the disposition of the office, including the couch position and the therapist's chair [16]. Kurt Lewin describes the importance of life-space in deriving a person's behavior [27], and the theory of affordances (e.g., [17]) talks about the possibilities offered by the environment. However, when it comes to virtual agents, the physical context seems largely forgotten.

The physical surroundings give us an impression of the intended behavior. For example, in many courtrooms, the judge is standing higher than the defendants. For formal work presentations, we might have the entire audience facing the presenter, or if we desire a less traditional setting, we can rearrange the seats in a circle. When placed in different settings, one adjusts their actions based on the existing context. Showing adaptive behavior determined by the agent's surroundings can suggest intention and, to an extent, socially intelligent behavior [44]. Moreover, cognitive processes reflect these adaptations. Recognizing and adapting to distinct contexts reshapes one's internal mechanisms to fit the environment, hence situating that individual's cognition [9, 39].

When deciding what line of action to pursue, people consider meaningful and noticeable aspects of their surroundings [8]. Several motivators can influence people's interpretation and construct of the context including the physical setting, nearby social actors [5], cultural values and social norms [24], or even internal goals [20]. However, the literature reveals a lack of consensus on the definition of context. The lack of agreement on context's conceptualization extends to its categories. Categorization schemes are either too general or incomplete [2]. Still, among the several categories of context, the importance of the spatial characteristics of the surroundings is mentioned by several researchers, even though the terminology is not consistent (e.g. "locations" in [46] and "place" in [11]).

To endow intelligent agents with the necessary capabilities to live alongside humans, computational models for social agents should allow them to fit their context and match the surrounding social actors' expectations. Thus, it is necessary to include the mechanisms required to support and display social awareness in the artificial agent's computational models while shaping their internal processes to fit the social context [12] better. Researchers have proposed some models that attempt to establish a strong foundation of sociality embedded across all the artificial agent's cognition [38][13][31][36]. Nonetheless, to deploy autonomous agents with such socially focused models into the wild, researchers must understand how an observer perceives and acknowledges an agent's actions in multiple contexts, including its social motivations.

Although several researchers recognize the importance of spatial characteristics for agent cognition, there is a lack of empirical studies supporting these claims. We intend to explore this human sensitivity to behavior changes, and its motivations, in virtual agents when this adjustment is based on the physical context. We conducted a user study to investigate if an observer identifies social motivations when an agent exhibits behavior changes grounded on their location to create context-aware agents. We explore the situatedness of the agent's behavior regarding a physical context defined by rooms with clear boundaries and, optionally, marked with distinctive symbols on the floor. Using minimalist representations of social agents and decoupling the physical environment of potentially acquired social rules enables the observer with a new frame of reference. In this study, we attempt to validate the following two hypotheses:

Hypothesis 1. *When the behavioral changes are based on the physical context, participants will attribute higher values for social motivations to the agent than when the behavioral changes are random.*

Hypothesis 2. *When the physical environments have a visual cue that distinguishes them, participants will attribute higher values to context-related motivations than when no visual cue exists.*

In this paper, we refer to agents social motivations to define the causes that lead to behavioral changes in order to accommodate a rule related to the agent's context. If people perceive the agent's adaptation to different physical contexts as socially motivated, researchers and designers may focus on endowing virtual and robotic agents with the mechanisms to adapt their behavior based on their surroundings, to make the agents appear more social. For instance, when an human is observing the actions of an autonomous robot that consistently adjusts its behavior to the physical context, the agent might be perceived as more social when compared to one that does not coherently adapt its behavior.

In the following section, we review other empirical works exploring the perception of social behavior based on authored actions and intelligent mechanisms. Afterwards, we present a scenario inspired by the apparent behavior experiment [19] that uses artificial agents implementing the Cognitive Social Frames model [36]. We then elaborate on the experimental study we conducted to verify both hypotheses and discuss the implications of our findings in the development of computational models for social agents. The findings of this study shed new light on what it means for an intelligent agent to display social behavior based on its physical context, namely the relevance of adaptive mechanisms towards supporting socially situated agents.

2 RELATED WORK

Heider and Simmel's pioneered research on the perception of human-like behavior on virtual entities in their experimental study of apparent behavior [19]. Before any computational graphics started to be rendered, the authors relied on an animated film with simple shapes (circles and triangles) to conduct an experimental study about an observer's perception of behavior. Their results showed that participants ascribed intentions and motives to the shapes' movements, suggesting that humans attributed anthropomorphic qualities to non-human agents. Other social scientists have also contributed to the understanding of perceived behavior: from people's attribution of reasons [41] and intentionality [35] to individual and group behavior, to folk notions of belief, desire, and intentions used to predict and understand human behavior [23], there have been several theoretical frameworks to explain behavior attribution [29].

Gong explored how anthropomorphic qualities of virtual characters elicit more social responses from people [18]. Based on several facial images with distinct anthropomorphism levels, the author conducted an experimental study in which he asked participants to decide on several social dilemma games. The results suggest that increasingly more human-like computer representations elicited more social responses from people. Also, using functional magnetic resonance imaging (fMRI), Krach *et al.* [25] studied how different degrees of robots' anthropomorphism may affect a player's perception of its partners in a human-robot game scenario. Their findings suggest that the perceived human-likeness of a robot linearly increases a human capability to model other's "minds", hence promoting the attribution of their intentions and motives. These works suggest that anthropomorphism qualities in artificial agents evoke cognitive capabilities associated with social intelligence. However, we hypothesize that it is possible to attribute social motivations to the behaviors of non-humanoid abstract shapes, like the ones in [19].

Sturgeon *et al.* conducted an experimental study aimed at exploring the impact of theory of mind in the perceived social intelligence of robots [40]. The results suggest that the capability to adapt to another human's actions contributes to the perception of social intelligence. Also, other researchers claim that agents, namely robots, are useful to study social behavior when properties of the physical environment are expected to influence one's social cognition [33].

To the best of our knowledge, there is no empirical study relying on artificial agents to investigate the effect of the physical surroundings on social intelligence perception. Other researchers also identify the scarcity of empirical work to verify assumptions and test predictions about social influence [15]. With the increasing interest in explainable AI and the deployment of autonomous agents in the wild, our community is experiencing a valuable endeavor to support new contributions with findings from human science research works [32] while creating intelligible and interpretable systems [1]. We argue that the interpretability and explainability of the AI behavior should be grounded on the context it sits as well.

3 SCENARIO

Artificial agents are capable of evoking human socio-cognitive capabilities [45]. Relying on artificial agents to study human's social nature has been identified as an adequate compromise between the ecological validity of the findings and its experimental control and has found increased usage when studying the mechanism of socio-cognition in virtual characters [7] or robotic agents [6].

To verify our hypotheses, we created a virtual world with three different rooms and a cube-shaped character. The scenario's minimalist and abstract nature was inspired by Heider and Simmel's apparent behavior experiment [19]. Also, while measuring social motivators, the observer's cultural and societal influences may bring forward expectations about adequate behavior. These preconceptions can present an inconvenient and troublesome bias to the study of the impact of adaptive behavior and multiple motivators on the agent's perceived social behavior. To avoid these preconceptions, we conducted an experimental study that relied on non-humanoid agents in virtual worlds to mitigate external influences.

In our scenario, a colored cube-shaped character moves between rooms in a predefined path. The artificial agent can change between three colors (pink, yellow, or blue) based on a policy P that defines the color selection. For the scope of this scenario, the agent can either use a context-based policy or a random policy. When following the first, an appropriate color is selected for each context, *i.e.* each room has a specific color assigned (e.g. the character always changes to yellow on a particular room). When following the random policy, the color selection is determined by a controlled sequence obtained through randomization.

When entering a new room, the character chooses a new color based on its policy P. To reflect this decision and the amplitude of choices, regardless of the policy, the agent performs a fast animation showing all three possible colors, and then adopts the color determined by the policy. Figure 1 shows the base scenario, where the cube agent is exhibiting a pink color.

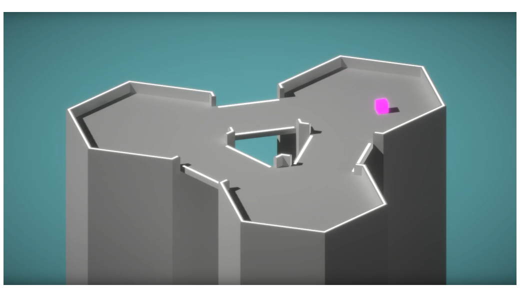

Figure 1: Base scenario with three rooms with no symbols. The character represented by the pink cube is able to move between rooms and adopt different colors.

The scenario allows the inclusion of visual cues - symbols - in the rooms' floor, either in all rooms or in none. This allows to study if the presence of distinctive visual cues in different contexts influences the identification of motivations. Rather than using colors or shapes that might reassemble the agent's physical characteristics, we used as symbols the characters **#**, *****, and **&**. Figure 2 shows two possible variations of the scenario.

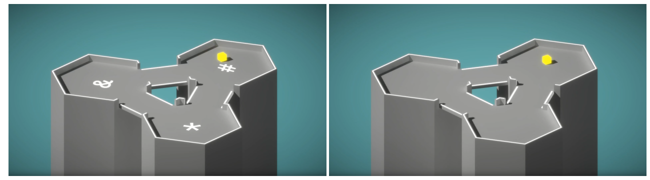

(a) With Symbol **(b) Without Symbol**

Figure 2: Two variations of the base scenario: (a) each room floor has a different symbol that stays visible and (b) there are no symbols in any room.

As Heider and Simmel identified in their experiment [19], the use of non-humanoid figures decouples any preconceptions regarding adequate behavior motivated by human's daily observations, establishing a clear framework to understand and explore an external observer's interpretation of virtual agents' actions.

4 EXPERIMENT

The goal of the experiment is to study the effect that an agent's behavior adaptation has on the perception of an external observer.

The adaptation is represented by the behavioral changes based on the physical context of the agent. Namely, we want to explore if different policies based on location (e.g. the room where the agent is) impact the motivators ascribed to an agent's behavior. The main goal of this experiment is to study the effect of an agent's behavior changes (context-based vs. random) on the attributions made by external observers. In this study, the behavior changes are either context-based or random. Context-based changes mimic behavioral adjustments made according to the agent's surroundings, whereas random behavioral changes occur just like the name indicates, at random. We want to explore if humans are sensitive to these changes and if that leads them to attribute higher social motivations to the agents whose behavioral changes are context-based compared to the random policy agents. We consider the context-based agents to be adapting their behavior to their surroundings and random policy agents as not adapting to the surroundings.

Additionally, we want to study if the presence of an explicit symbol strengthens the behavior's attribution to the spatial motivators (physical elements in the environment). By introducing visual cues to each room, we expect to strengthen the differences among them, help participants assign a distinct value to each physical space, thus highlighting the salience of distinct behaviors.

We used Amazon's Mechanical Turk platform to conduct the study and recruit participants, and the framework jsPsych [10] to create and execute the experiment.

4.1 Participants

We employed 136 participants from Mechanical Turk and, prior to data analysis, we eliminated 44 (32.4%) who failed the attention checks. As such, our analysis uses a sample of 92 participants: 47 were women and 45 were men and their ages ranged from 23 to 64 with a mean age of 38.46 (SD = 8.87). Each participant performed a single Mechanical Turk Human Intelligence Task (HIT) with an estimated duration of 10 minutes and as rewarded at a rate of 6 US dollars per hour (e.g. 1 US dollar per task). To ensure the participants were qualified for the experiment, we defined the following inclusion criteria: participants had to be from an English-speaking country (Canada, UK, or the US), have a high HIT approval rate (at least 97%), and have already performed a minimum of 5,000 HITs.

4.2 Materials

Every participant saw two videos of a cube exploring three rooms (see Figure 1), fill an attention check with two items, performed a distracting task, and complete the Social Motivation of Intelligent Agents Scale (SMIAS). We created four videos [1] of an agent moving around multiple rooms and changing colors whenever it enters a room. Each video has a duration of 84 seconds and the agent visits each one of the three rooms twice in a clockwise order starting from the top right room. Figure 3 illustrates the different steps that define the cube's path when entering a new room.

path.png

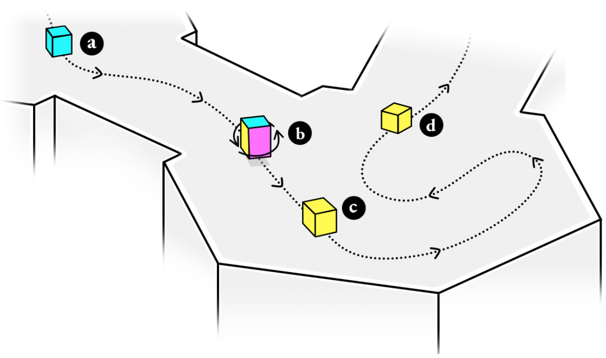

Figure 3: Path and animation of the agent inside each room: (a) the cube exits one room with a particular color - blue, (b) the cube enters a new color and performs the changing color animation, (c) the cube selects one color - yellow - after the changing color animation is concluded, and (d) after randomly walking around the room, the cube moves to the next room

To ensure participants were watching the videos (the attention check), we asked them two simple questions (Question 1: "How many rooms existed?"; Question 2: "What colors did the agent use?"). After the first video, attention check and scale, participants were asked to do a simple memory task that consisted of remembering 3 sequences of numbers. The goal of this task was to provide some buffer between the two portions of the experiment.

4.2.1 Social Motivation of Intelligent Agents Scale. We searched the literature for a scale that would measure the agency and motivation of intelligent agents, or that we could use for that purpose. We considered measures used in social robotics, many of which seem to focus on negative attitudes towards robots (for a review see [26]). Two frequently used measures in social robotics include the Godspeed questionnaire ([3]) and RoSAS ([4]). The Godspeed questionnaire ([3]) is a semantic differential scale where participants must rate robots' attributes from 1 to 5. For example "Fake (1-2-3-4-5) Natural", where one means the participant perceives the robot as fake, and five means the participant perceives the robot as being natural. The questionnaire has three factors: Anthropomorphism, animacy, and likeability, none of which seems to give an idea of motivation, agency and following social norms. Another frequently used questionnaire is RoSAS ([4]) which was built from the Goodspeed items and presents three factors: Warmth, competence and discomfort. Again this scale provides an insight about how people perceive interactions with social robots, and even how they feel about them still, they do not seem to give answers regarding motivation for action and agency. In the case of measures that focus on agency, such as self-reported scales (e.g. [42]) or perceived agency questionnaires (e.g. [22, 34]), these instruments provide little insight about the motivation for one's agency.

Due to the absence of measures that would serve our goals, we developed a scale for the scope of our study. The original scale included a list of 36 items initially organized in three groups: agency regarding the new color (e.g., "the cube chooses the color it wants"); agency about changing the color (e.g., "the cube decides when to change colors"); and motivation (e.g., "the cube decides to follow a rule about when to change colors"). For each group, there was an even number of options conveying full agency (e.g., "the cube chooses the color it wants"), 'mixed agency' (e.g., "the cube decides when to change colors"), and no agency (e.g., "the cube does not decide its color"). Towards using the scale in other scenarios, the wording of each item was planned to accommodate other agents and target behaviors as needed. For instance, the item "The cube knows what color to choose for each room" can be changed to "The robot knows what movement to choose for each game", where the *robot* is the agent, *movement* is the action and *game* is the context.

We also proposed the items with potential redundancy among them. Redundancy and length are important to conduct a factorial analysis to determine different intrinsic factors (for a review see [30]). We performed a factorial analysis using Winsteps [28] to reduce the dimension of our scale. The sampling adequacy was verified using Kaiser-Meyer-Olkin measure with a KMO= 0.81 ('great' according to [21]). Bartlett's test of sphericity $X\,2(153) = 969.0, p < 0.0001$ indicating that the correlation between items was large enough for a PCA. Four components had eigenvalue over Kaiser's criterion of 1 covering 69% of the variance; which correspond to the 4 factors selected for the scale. The items that cluster on the same components suggest that the first represents Social Motivation, the second Spatial Motivations, the third the agent's Awareness and the last represents its Agency. The resulting scale included 18 items grouped into 4 factors. The mappings between the factors and the scale items are described in Appendix A.

4.3 Design and Procedure

The experiment used a mixed design 2 (Policy: Context-based; Random) x 2 (Symbol: with symbol; without symbol), where policy was the within-subjects factor and Symbol, the between subjects factor. All participants followed the same procedure. The experiment started with a page with instructions, followed by a page asking for information about their age and gender. They were randomly

[1] https://youtu.be/f2-d6FYAlTg

assigned to one of two groups (symbol vs no symbol) and the order of the two videos (context-based or random) was counterbalanced. After the first video, subjects responded to the attention check and then to the items of the Social Motivation of Intelligent Agents Scale. After completing the scale, they were asked to perform the memory task, and afterward, they saw the second video, the attention check, and responded to the Social Motivation of Intelligent Agents Scale again. Once they finished, they had an open-ended question asking: "Using your own words, please try to explain the main characters behavior on both videos".

4.4 Results

We carried out a 2 (Symbol: Symbol; No Symbol) x 2 (Policy: Context-based; Random) x 4 (SMIAS: Social; Awareness; Spatial; Agency) mixed ANOVA with Symbol as a between factor and policy and SMIAS factors as within-subjects factors.

The ANOVA shows no main effect of symbol ($F < 1$), and no interactions with this variable (all F's < 1). There is a main effect of policy ($F(1, 90) = 43.68, p \leq .001, \eta_p^2 = .33$), a main effect of social motivation of intelligent agents ($F(1.9, 270) = 47.12, p \leq .001, \eta_p^2 = .14$) and the two variables interact ($F(1.7, 270) = 31.7, p \leq .001, \eta_p^2 = .26$). This interaction was decomposed using simple effects contrasts with a Bonferroni correction and the contrasts show that participants rated the behavior of the cubes as being significantly more social when the policy was context based, $M = 3.49; SE = .14$, than when the policy of the agent was random, $M = 2.48; SE = .15$; $F(1, 90) = 31.8, p \leq .001, \eta_p^2 = .26$, as Figure 4 shows. The behavior of the cubes was also rated as reflecting higher awareness when the policy was context-based, $M = 4.43; SE = .11$, than on the random policy condition,$M = 3.44; SE = .12$; $F(1, 90) = 52.2, p \leq .001, \eta_p^2 = .37$. Regarding the relevance of spatial characteristics of the environment, subjects rated the physical world as more impacting on the agents behavior on the context based condition $M = 4; SE = .19$, than on the random policy condition, $M = 2.5; SE = .19$; $F(1, 90) = 37.7, p \leq .001, \eta_p^2 = .3$. The ratings attributed to agency show the opposite pattern as subjects rated the cubes as being significantly lower in agency when the policy was context-based $M = 2.98; SE = .16$, than on the random policy condition, $M = 3.89; SE = .12$; $F(1, 90) = 23.25, p \leq .001, \eta_p^2 = .21$.

The *open-ended question* was the last task, and, as such, any conclusion needs to be drawn carefully since all participants had previously responded to SMIAS, which most likely affected their answers to this question. We asked participants to explain, using their own words, the main character's behavior on both videos. We collected 91 answers to our open-ended question, and 13 of those were considered invalid. These answers did not paint a clear picture of the participant's understanding of the videos and were mostly single-word answers. For example, the answers solely included short statements such as "*robotic*", "*precise and intelligent*", and "*uncertain*".

Some participants made a clear distinction between the cubes, providing an answer for each video, whereas others did not. For example, one participant wrote, "*Both cubes proactively changed color when entering a new room (...)*", and another participant wrote, "*One cube assigned a specific color to each room. The other cube randomly chose a color for each room*". Based on whether people

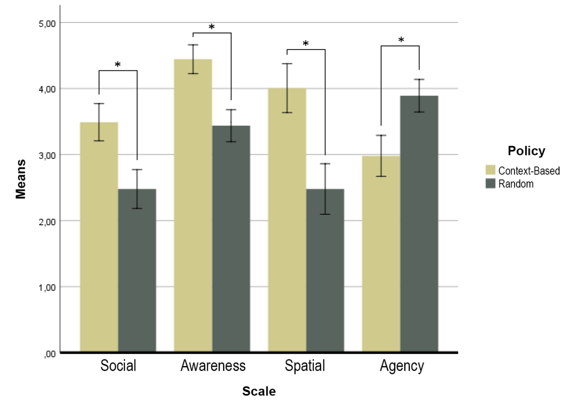

Figure 4: Scale Factors (Social, Awareness, Spatial, Agency) grouped by Policy

distinguished the cubes in both videos, we divided the 78 answers in two groups: on G_D 31 (40%) made clear distinctions, and on $G_{\sim D}$ the remaining 47 (60%) grouped the videos.

On $G_{\sim D}$, 32 (68%) are descriptions of the behavior that do not provide any intentions or motivations; they simply describe movement. For example:"*they scooted around a weird indirect path and would randomly pop up spin and change color. (...)*". Additionally, 12 answers (26%) of $G_{\sim D}$, mentioned the cubes decision power or awareness, "*(..) the cube could change its color regardless of the room but preferred to change the color based on entering a new room. So the room played a factor but the cube still had control*". Another 2 answers of $G_{\sim D}$ mentioned the room controlling the behavior of the cubes "*(...) Each room represents a different color so no matter what each room the cube goes it will change colors either way*". The remaining answer did not fit any of the previous categories and was therefore classified as "others".

In G_D, it is possible to distinguish between the policies (random vs. context-based). For the random policy, the most frequent answer related to choice and free will (19). For example, one participant wrote "*(...) the second one is making its own choices*". Nine participants made simple descriptions of the policy behind the random policy video, for example, "*The first cube seemed to change colors randomly(...)*". Three answers did not fit into any of these categories. Regarding the context-based policy, we noted that the more frequent answers were descriptions of the policy (20), e.g., "*The first cube seemed to change color in each room, and it was constant: always a certain color in a certain room"*. The second most frequent answer for the CB video (10) mentioned the room, the cube's ability to choose, and their relation (e.g., ""*The second cube changed based on the room it was in. It did seem to be adapting to it's environment*". Finally one answer did not seem to fit in any of previous categories.

4.5 Discussion

Since the behavioral changes were triggered when the agent entered a new context, participants identified the context-based behavior as more spatially motivated than the random behavior. However, based on our results, the presence of a physical cue (a symbol in each context) did not influence the observer's perception of the behaviors' motivations. As such, Hypothesis 2 was not validated. Still, the lack of evidence to support the aforementioned hypothesis may suggest that a stronger representation that the agent identifies a new room, and thus adapting its behavior, is required. As Gibson suggested, one's perception is guided by the opportunities the environment affords [17] and, by requesting an external interpretation of the agent's motivator without displaying any hints about the perception of such symbol, might not be sufficient for an external observer to relate the perception of a specific context with an adequate behavior. Moreover, as Semin and Garrido's findings highlight [37], ascribing motivations solely based on environmental physical features might not be sufficient since one's sensor capabilities also play a role in behavior adaption. As such, the existence of visual tokens in each context might not be enough to identify an interdependence between such context and the behavioral change. Instead, it may be necessary to explicitly display the agent's sensory outcome that accounts for the recognition of such environmental cues may be necessary.

The results also indicate that participants ascribe higher values for social motivations to context-based behavioral changes when compared to the random behavioral changes. As such, the evidence supports Hypothesis 1. As proposed by Smith and Semin, socially situated cognition happens in a specific context to guide an adaptive action [39]. Aligned with their claim, our findings suggest that external observers ascribed higher values for social motivations to the agents that exhibit a context-based behavior than when random behavioral changes are displayed. Considering that no priming or social expectations about the adequate behavior for each context was carried to the experimental setting, our results indicate that just by displaying a consistent policy for behavioral changes, external observers seem interpret such actions as more socially motivated. The most frequent answer to the open-ended question for CB policy videos was a description of the policy; this does not give us information about intentions but assures us that people could identify the cube with the same color for the same room. The second most frequent answer mentioned the choice based on the room. The cube either assigned a color for each room based on its perceptions or knew which color to take for each room.

The remaining two factors of our scale, Awareness, and Agency, indicate opposing patterns from each other. Regarding the first, observers interpreted that the agent is more aware when it exhibits a context-based behavior. Despite mean values for the perceived awareness being above average, this result may reveal that participants might not recognize agents following a random policy as capable of identifying distinct contexts, which aligns with the previously mentioned issue of not explicitly displaying the outcome of the sensory process. However, regarding the agent's agency, the ones that exhibited a random behavior received higher values for agency than the context-based ones. Perhaps, this significant difference indicates a perceived lack of freedom to determine its actions when the agent changes its behavior based on the context. In the open-ended question, participants ascribed free will and choice more frequently to the random policy video describing the cube as making its own choices regardless of the rules. The random policy video was the only group where the most frequent answer was not a mere description of the policy. Additionally, by randomizing the agent's behavior, participants might have interpreted the agent's actions as being influenced by intrinsic motivations (e.g., goals and drives) besides external ones (e.g., spatial and social). Participants may have interpreted the agent as being less susceptible to social conformity. Also, as other researchers investigated in robots [14], the unpredictable behavior of the cubes helps to anthropomorphize the agent. Still, considering the previous remarks, what explains the attribution to the cube's awareness and agency is yet unclear. Thus, more work to identify the factors contributing to the perceived awareness and agency of virtual characters is necessary.

One important limitation of this work is the fact that we use a scale that contains social motivations rather than letting participants interpret the videos freely like Heider & Simmel [19]. Although we are in fact providing cues to the type of attribution we expect, the results show significant differences on the values attributed to both policies with higher values for social motivation when the policy is context based.

5 CONCLUSION

In this work, we explored how external observers interpret a virtual agent's behavior, namely the motivations behind its behavioral changes. While endowed with mechanisms to adequate its action to a location, the reasons to perform such changes might not be apparent for a bystander. Using a minimalist scenario, we conducted an experimental study to understand what motivations are identified by the participants. To measure such motivations, we proposed and employed a new scale SMIAS to measure intelligent agents' social motivation. Our results suggest that, although no other agents are present and there is no previous knowledge about the adequate behavior for each context, participants ascribe higher social motivations to agents that change their behavior based on the context defined by the physical location. Our findings reveal that people are sensitive to changes in behavior to adapt the surroundings, and they attribute to it higher social motivations. As such, while developing social agents, architectures and computational models should include adaptive mechanisms that enable such artificial agents to situate their behavior based on the physical location. Moreover, previous research suggests that external observers expect more consistency in an individual's behavior compared to that of a group [41]. Towards creating more socially believable virtual agents, demands a better understanding of the individual and group social reasons. As such, studying and comparing if such attributions also applies to groups of virtual agents is a future research direction worth of exploring.

Additionally, we contributed with a versatile framework to study external observer's interpretation of intelligent agents' behavior in virtual worlds. This creates an opportunity to use virtual worlds and intelligent virtual agents to study human-agent interaction and, to an extent, human-human behavior. We aim at further exploring other valences of social context, more specifically when other agents are present. Also, since our hypothesis on the symbols was not supported by our results, it is worth enhancing the agents' perception of the symbols to see if it affects the attribution of social motivations as well as introduce self-contextual symbols, such as emojis or cultural icons. Furthermore, the SMIAS scale was designed for this experiment and more work is necessary to validate it. Due to the absence of a self-reported measure for ascribing social motivations to agents, we consider a robust validation of SMIAS as an important contribution to this research community.

To the extent of our knowledge, there were no empirical studies that rely on artificial agents in virtual worlds to study the effect of the context in the perception of social intelligence. Moreover, other researchers also identify the scarcity of empirical work to verify some assumptions and test predictions about social influence [15]. With this work, we expect that by studying the interpretations of social agents' actions, researchers will be able to understand the motivation and importance of elements that compose the social context, thus creating models that recognize and understand human social cognition, and are also able to reproduce such capabilities.

ACKNOWLEDGMENTS

This work was supported by national funds through Fundação para a Ciência e a Tecnologia (FCT) with reference UIDB/50021/2020. Diogo Rato acknowledges his FCT grant (SFRH BD/131024/2017)

REFERENCES

[1] Ashraf Abdul, Jo Vermeulen, Danding Wang, Brian Y Lim, and Mohan Kankanhalli. 2018. Trends and trajectories for explainable, accountable and intelligible systems: An hci research agenda. In *Proceedings of the 2018 CHI conference on human factors in computing systems*. ACM, 582.

[2] Unai Alegre, Juan Carlos Augusto, and Tony Clark. 2016. Engineering context-aware systems and applications: A survey. *Journal of Systems and Software* 117 (2016), 55–83.

[3] Christoph Bartneck, Dana Kulić, Elizabeth Croft, and Susana Zoghbi. 2009. Measurement instruments for the anthropomorphism, animacy, likeability, perceived intelligence, and perceived safety of robots. *International journal of social robotics* 1, 1 (2009), 71–81.

[4] Colleen M Carpinella, Alisa B Wyman, Michael A Perez, and Steven J Stroessner. 2017. The robotic social attributes scale (RoSAS) development and validation. In *Proceedings of the 2017 ACM/IEEE International Conference on human-robot interaction*. 254–262.

[5] Cristiano Castelfranchi. 1994. Guarantees for autonomy in cognitive agent architecture. In *International Workshop on Agent Theories, Architectures, and Languages*. Springer, 56–70.

[6] Thierry Chaminade and Gordon Cheng. 2009. Social cognitive neuroscience and humanoid robotics. *Journal of Physiology-Paris* 103, 3-5 (2009), 286–295.

[7] Thierry Chaminade, Jessica Hodgins, and Mitsuo Kawato. 2007. Anthropomorphism influences perception of computer-animated characters' actions. *Social cognitive and affective neuroscience* 2, 3 (2007), 206–216.

[8] Paul Cisek and John F Kalaska. 2010. Neural mechanisms for interacting with a world full of action choices. *Annual review of neuroscience* 33 (2010), 269–298.

[9] William J Clancey. 1991. Situated cognition: Stepping out of representational flatland. *AI Communications The European Journal on Artificial Intelligence* 4, 2/3 (1991), 109–112.

[10] Joshua R De Leeuw. 2015. jsPsych: A JavaScript library for creating behavioral experiments in a Web browser. *Behavior research methods* 47, 1 (2015), 1–12.

[11] Frank Dignum. 2018. Interactions as Social Practices: towards a formalization. *CoRR* abs/1809.08751 (2018). arXiv:1809.08751 http://arxiv.org/abs/1809.08751

[12] Frank Dignum, Rui Prada, and Gert Jan Hofstede. 2014. From autistic to social agents. In *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems*. International Foundation for Autonomous Agents and Multiagent Systems, 1161–1164.

[13] Virginia Dignum and Frank Dignum. 2014. Contextualized planning using social practices. In *International Workshop on Coordination, Organizations, Institutions, and Norms in Agent Systems*. Springer, 36–52.

[14] Friederike Eyssel, Dieta Kuchenbrandt, and Simon Bobinger. 2011. Effects of anticipated human-robot interaction and predictability of robot behavior on perceptions of anthropomorphism. In *Proceedings of the 6th international conference on Human-robot interaction*. ACM, 61–68.

[15] Andreas Flache, Michael MÃ¤s, Thomas Feliciani, Edmund Chattoe-Brown, Guillaume Deffuant, Sylvie Huet, and Jan Lorenz. 2017. Models of Social Influence: Towards the Next Frontiers. *Journal of Artificial Societies and Social Simulation* 20, 4 (2017), 1–2.

[16] Sigmund Freud. 1913. On beginning the treatment, Standard.

[17] James J. Gibson. 1979. The ecological approach to visual perception. Boston, MA, US.

[18] Li Gong. 2008. How social is social responses to computers? The function of the degree of anthropomorphism in computer representations. *Computers in Human Behavior* 24, 4 (2008), 1494–1509.

[19] Fritz Heider and Marianne Simmel. 1944. An experimental study of apparent behavior. *The American journal of psychology* 57, 2 (1944), 243–259.

[20] Jacob B Hirsh, Raymond A Mar, and Jordan B Peterson. 2012. Psychological entropy: A framework for understanding uncertainty-related anxiety. *Psychological review* 119, 2 (2012), 304.

[21] Graeme D Hutcheson and Nick Sofroniou. 1999. *The multivariate social scientist: Introductory statistics using generalized linear models*. Sage.

[22] Katherine Isbister and Patrick Doyle. 2002. Design and evaluation of embodied conversational agents: A proposed taxonomy. In *The first international joint conference on autonomous agents & multi-agent systems*. Citeseer.

[23] Yoshihisa Kashima, Allison McKintyre, and Paul Clifford. 1998. The category of the mind: Folk psychology of belief, desire, and intention. *Asian Journal of Social Psychology* 1, 3 (1998), 289–313.

[24] Laurence Kaufmann and Fabrice Clément. 2007. How culture comes to mind: from social affordances to cultural analogies. *Intellectica* 46, 2-3 (2007), 221–250.

[25] Sören Krach, Frank Hegel, Britta Wrede, Gerhard Sagerer, Ferdinand Binkofski, and Tilo Kircher. 2008. Can machines think? Interaction and perspective taking with robots investigated via fMRI. *PloS one* 3, 7 (2008), e2597.

[26] Christian U Krägeloh, Jaishankar Bharatharaj, Senthil Kumar Sasthan Kutty, Praveen Regunathan Nirmala, and Loulin Huang. 2019. Questionnaires to measure acceptability of social robots: a critical review. *Robotics* 8, 4 (2019), 88.

[27] Lewin Kurt. 1936. Principles of topological psychology. *New York-London* (1936).

[28] J. M. Linacre. 2020. Winsteps® Rasch measurement computer program. Winsteps. com

[29] Bertram F Malle. 2011. Attribution theories: How people make sense of behavior. *Theories in social psychology* 23 (2011), 72–95.

[30] João Marôco. 2018. *Análise Estatística com o SPSS Statistics.: 7ª edição.* ReportNumber, Lda.

[31] Rijk Mercuur, Frank Dignum, and Yoshihisa Kashima. 2017. Changing habits using contextualized decision making. In *Advances in social simulation 2015*. Springer, 267–272.

[32] Tim Miller. 2019. Explanation in artificial intelligence: Insights from the social sciences. *Artificial Intelligence* 267 (2019), 1–38.

[33] Sara Mitri, Steffen Wischmann, Dario Floreano, and Laurent Keller. 2013. Using robots to understand social behaviour. *Biological Reviews* 88, 1 (2013), 31–39.

[34] Clifford Nass, Katherine Isbister, Eun-Ju Lee, et al. 2000. Truth is beauty: Researching embodied conversational agents. *Embodied conversational agents* (2000), 374–402.

[35] Matthew J O'Laughlin and Bertram F Malle. 2002. How people explain actions performed by groups and individuals. *Journal of Personality and Social Psychology* 82, 1 (2002), 33.

[36] Diogo Rato, Samuel Mascarenhas, and Rui Prada. 2019. Towards Social Identity in Socio-Cognitive Agents. *Springer Socio-Cognitive Systems* (2019). preprint: https://arxiv.org/abs/2001.07142.

[37] Gün R Semin and Margarida V Garrido. 2012. A systemic approach to impression formation: From verbal to multimodal processes. *Social thinking and interpersonal behavior* (2012), 81–96.

[38] Maarten Sierhuis, William J Clancey, and Ron JJ Van Hoof. 2007. Brahms: a multi-agent modelling environment for simulating work processes and practices. *International Journal of Simulation and Process Modelling* 3, 3 (2007), 134–152.

[39] Eliot R Smith and Gün Semin. 2004. Socially situated cognition: Cognition in its social context. *Advances in experimental social psychology* 36 (2004), 57–121.

[40] S. Sturgeon, A. Palmer, J. Blankenburg, and D. Feil-Seifer. 2019. Perception of Social Intelligence in Robots Performing False-Belief Tasks. *International Symposium on Robot and Human Interactive Communication (RO-MAN)* (Oct 2019). http://par.nsf.gov/biblio/10112946

[41] Joshua Susskind, Kristin Maurer, Vinita Thakkar, David L Hamilton, and Jeffrey W Sherman. 1999. Perceiving individuals and groups: Expectancies, dispositional inferences, and causal attributions. *Journal of Personality and Social Psychology* 76, 2 (1999), 181.

[42] Adam Tapal, Ela Oren, Reuven Dar, and Baruch Eitam. 2017. The Sense of Agency Scale: A Measure of Consciously Perceived Control over One's Mind, Body, and the Immediate Environment. *Frontiers in Psychology* 8 (2017), 1552. https://doi.org/10.3389/fpsyg.2017.01552

[43] Gary R VandenBos. 2007. *APA dictionary of psychology*. American Psychological Association.

[44] Alessandro Vinciarelli, Maja Pantic, and Hervé Bourlard. 2009. Social signal processing: Survey of an emerging domain. *Image and vision computing* 27, 12 (2009), 1743–1759.

[45] Agnieszka Wykowska, Thierry Chaminade, and Gordon Cheng. 2016. Embodied artificial agents for understanding human social cognition. *Philosophical Transactions of the Royal Society B: Biological Sciences* 371, 1693 (2016), 20150375.

[46] Andreas Zimmermann, Andreas Lorenz, and Reinhard Oppermann. 2007. An operational definition of context. In *International and Interdisciplinary Conference on Modeling and Using Context*. Springer, 558–571.

A SOCIAL MOTIVATION OF INTELLIGENT AGENTS SCALE

The Social Motivation of Intelligent Agents Scale has 18 items divided into the following 4 factors: Social Motivations, Spatial Motivations, Awareness and Agency. The items identified with *(R)* are reversely scored.

(1) Social Motivations
  (a) No social rules are influencing the cube's color *(R)*
  (b) The colors have no social meaning *(R)*
  (c) Each color has a social meaning
  (d) Each room has a social meaning
  (e) When changing colors, the cube is following social rules
(2) Spatial Motivations
  (a) Each room changes the cube's color
  (b) Entering the room changes the cube's color
(3) Awareness
  (a) The cube changes colors as a reaction to its perceptions
  (b) The cube decides to follow a rule about when to change colors
  (c) The cube knows what color to choose for each room
  (d) The cube perceives the rooms as being different from each other
  (e) The cube realizes it is in a different room
(4) Agency
  (a) Changing colors is not a choice but an imposition *(R)*
  (b) The cube's color always changes when it changes rooms *(R)*
  (c) The cube decides when to change colors
  (d) The cube has full control of its color
  (e) The cube intelligently changed color
  (f) The cube is not forced to change colors

The Social Motivation of Intelligent Agents Scale supports other agents, actions and contexts, besides the ones addressed on our scenario: *cube*, *color changes* and *room*. As such, the scale can be applied to other use cases with the different entities, with their own behaviors and respective settings.