

Intrinsically Motivated Autonomy in Human-Robot Interaction: Human Perception of Predictive Information in Robots

Marcus M. Scheunemann¹, Christoph Salge¹, and Kerstin Dautenhahn^{1,2}

¹ University of Hertfordshire, AL10 9AB, United Kingdom

² University of Waterloo, N2L 3G1, Canada

<https://mms.ai/sHRI>

`marcus@mms.ai`

Abstract. In this paper we present a fully autonomous and intrinsically motivated robot usable for HRI experiments. We argue that an intrinsically motivated approach based on the Predictive Information formalism, like the one presented here, could provide us with a pathway towards autonomous robot behaviour generation, that is capable of producing behaviour interesting enough for sustaining the interaction with humans and without the need for a human operator in the loop. We present a possible reactive baseline behaviour for comparison for future research. Participants perceive the baseline and the adaptive, intrinsically motivated behaviour differently. In our exploratory study we see evidence that participants perceive an intrinsically motivated robot as less intelligent than the reactive baseline behaviour. We argue that is mostly due to the high adaptation rate chosen and the design of the environment. However, we also see that the adaptive robot is perceived as more warm, a factor which carries more weight in interpersonal interaction than competence.

Keywords: Robotics · Sustained Interaction · Human-Robot Interaction · Autonomous Human-Robot Interaction · Robot Behaviour · Cognitive Robotics · Robot Control · Autonomous Robots · Intrinsic Motivation · Predictive Information · Information Theory.

1 Introduction

Why use autonomous robots for human-robot interaction (HRI) experimentation [12] instead of teleoperation by human experimenters or scripted behaviour? Scripting reduces the adaptability of the robot to novel situations, limiting the range and flexibility of interaction scenarios. Teleoperation offers more flexibility, but has problems with scalability, introduction of human bias, and experimenters struggling with acting from the robot’s perspective [1]. In contrast, autonomous robots could realise experiments where they freely interact with humans in real environments. Here we want to introduce some exploratory experiments with an intrinsically motivated robot as a pathway towards realising autonomous robots that are interesting to interact with.

This paper is motivated by the authors’ work with a spherical robot [22] (similar to Fig. 1) and children. The children’s interaction patterns were usually very diverse, making a general, pre-scripted robot behaviour for a group of children hard to achieve. Without human interaction or without remotely controlling the robot, the behaviour of the robot was limited, leading to the children losing interest in the robot. While it would be trivial to have a self-controlled robot that exhibits some form of behaviour, the hard question is: What kind of behaviour makes a robot interesting to interact with? We assume that robots are more interesting to interact with if they have perceived agency, allowing the human interaction partner to assign motivations to the robot, support or hinder its goals, or even sympathise with its “joy” when achieving a goal. Once we identify something as an agent, we are likely to direct our attention towards that agent, trying to understand its goals, intentions and behaviour. This interest is, if not synonymous with engagement, a step towards more engagement with the agent. While the simple act of moving [8] or the look of a robot’s “head” [5] can change a human’s perception of the robots, e.g., animacy, the behaviour also needs an observable goal-directedness [11]. We think that behaviour based on actual intrinsic motivation [18] would be a good candidate to create a consistent perception of a robot’s motivation. We think that this enables the possibility of eventually sustaining the interaction with a fully autonomous robot.

In psychology, intrinsic motivation is defined as doing an activity for its inherent satisfaction rather than for some separable consequence or reward [9,20]. The concept has been linked to the idea of autonomy and agency [19], and intrinsic motivations are alternatively defined as those motivations that are an integral, non-instrumental, non-optional part of an agent [17]. The recent interest in computational approaches to intrinsic motivations [18] gives us a rich assortment of formalisms to consider. Most of them have a set of common properties, such as semantic independence, universality and sensitivity in regards to embodiment, and a high degree of robustness. They are usually used to answer questions, such as what is a good general heuristic (i.e. motivation function) if my robot knows nothing about the world or even its own morphology? Ideas such as curiosity or self-maintenance are turned into AI formalism that enhance model learning and were generally used to enhance AI and robot performance [4]. More relevant to our approach, recent work in the HCI domain of games used intrinsically motivated agents to generate more interesting, self-learning agents [16], or to create believable, generic antagonists [13] and companions [14].

2 Predictive Information

The Intrinsic Motivation used to generate the robot’s behaviour in our study is the Predictive Information (PI) formalism, closely following the implementation of Martius et al. [15]. A formal introduction of the measures is omitted due to space constraints. Conceptually, the measure falls into a family of learning rules related to the reduction of the time prediction error in the perception-action loop of the robot. The book *Playful Machines* [10] offers a good introduction.

The book also shows how these approaches can be computed from a robot’s perspective alone, and the zoo of different robots and their behaviours presented within shows how behaviour resulting from the different formalisms is sensitive towards the agent’s specific embodiment.

Predictive Information [2] derives a specific learning rule, that aims to maximise the mutual information between a robot’s past and future sensor states. The relevant literature argues that this produces exploratory behaviour sensitive to the robot’s embodiment. The derivation from information theory also offers an intuitive interpretation of the robot’s adaptation towards being able to reliably predict its own future from its past, while enriching the diversity of its experiences. The approach we use here [15] works by updating the internal neural networks of the robot, one that generates behaviour from sensor input, and the other predicting the futures states. The continuous adaptation, aimed at improving the time-local Predictive Information, moves the robot through a range of behavioural regimes. Importantly, the changes in behaviour are partially triggered by the interaction with the environment, as mediated through the robot’s embodiment. The rate at which those internal neural networks are updated is the one model parameter which we will change between experiments.

To the best of our knowledge, this is the first experiment that uses Predictive Information in the context of HRI, and evaluates how the behaviour based on this intrinsic motivation is perceived by participants.

3 Study Design and Procedure

The challenge in designing this study is that this work is, as far as we know, the first HRI experiment of a robot using Predictive Information. Consequently, we lack an existing baseline for comparison.

3.1 Baseline Behaviour

We considered the following four alternative means of behaviour-generation for serving as a baseline: (1) human remotely controls the robot, (2) random behaviour and (3) pre-adapted reactive behaviour.

Ideally, we want to see how the algorithm compares to a human remotely controlling the robot. However, human controlled behaviour has a high degree of variance, dependent on the particular human controller. Furthermore, it is unclear how much access the human controller should have to environmental information. If the human can directly observe the participants, it could obtain much more information than the robot, giving it an unfair advantage in creating behaviour responsive to the participant. If we limit the human controller to only the robots’ sensors, then the human controller would likely struggle to make sense of this limited input, potentially being unable to control the robot at all.

The problem with using random behaviour as a baseline is that “randomness” actually has a set of parameters that need to be chosen, like how often do values change, or is it the change of value or the value that is being randomised. PI

does not show a real pattern of behaviour switches we could use for the timing of changing speed, heading and/or the overall behaviour. We performed some preliminary trials with random values, but were quickly facing the question of a fair baseline behaviour again. Having the experimenter choose these values leads to basically designing a certain kind of behaviour (chosen from a whole range of behaviours), which makes it problematic as a baseline. We decided to use a pre-adapted reactive behaviour. The pre-adaptation is done with the very same PI implementation and parameters, using the same sensors as the robot will use during the experiments.

3.2 Robot, Environment and Tasks




Fig. 1. Left: The used robot platform BB8 from Sphero. Middle: A 2-D cross-sectional view of the robot. A two-wheel vehicle, kept in position by a heavy weight, moves the sphere when driving. A magnet attached to the vehicle keeps the head on top of the sphere facing in moving direction. Right: The environment the robot explores during the trials from a birds eye perspective. The white area is paper, the black is foam material and the beige coloured area is wood. At the top of the foam material is a hill area and a pit in the lower part. The bottom edge does not have a wall, forcing the participant to interact with the robot.

We want a very simplistic platform with a few degrees of freedom for focusing on the effects induced by Predictive Information. However, for rich behaviour, the robot also has to offer some sensor capabilities. We chose the off-the-shelf spherical robot from the company Sphero, specifically, the BB8 platform as depicted in Figure 1 (left). We chose this version of Sphero because of the advantages the head offered. A magnet keeps the head in driving direction, which gives the user a sense of the robot’s direction. This, and the fact that many people know the robot from movies, provides a better impression of a robot than using a white Sphero. The robot weighs 168 g. It has a 75 MHz ARM Cortex M4 on board powered by two 350 mA h LiPo batteries. The robot has a two-wheel electric vehicle inside the spherical shell, as depicted in the cross-sectional view of the robot in the middle of Figure 1. This is kept in position by a heavy weight, which

is made out of the coil for power inductive charging. A connection of a magnet to the vehicle keeps the attached head in position.

There are two ways of controlling the robot. The vehicle inside the robot has two servos. You can either control the speed and direction of each servo directly or you can use Sphero in “balancing mode”. In this mode the in-built controller tries keeping the robot upright and listens to speed and heading commands. The direct control mode offers a wide range of possible behaviours, e.g. turning on the spot or more wobbly locomotion. We decided however to use the balanced mode, again, for the sake of simplicity.

The BB8 can stream sensor information. It offers readings of a 3-DOF accelerometer, a 3-DOF gyroscope and the current servo position and servo speed. It also offers IMU readings in quaternions or euclidean angles³. The JavaScript API for BB8 has been unsupported since 2016-05-11 and we decided for developing a custom API which is based on C++ and is available from [21]. Figure 1 (right) shows the experimental environment. Two tables form the space where the robot can move around. The area is 180×120 cm in size. It is open to one side where the participant is supposed to stand and interact with the robot. In Figure 1 you can see a participant nudging the robot. The surface of the table differs in friction and height. The black foam area has a hill (top) and a pit (bottom). Additionally, the black area and the white paper area is softer and has higher friction compared to the wooden part.

The participants’ task is to observe the robot and understand whether it has a strategy for exploring the environment. We wanted to encourage the participants to interact with the robot. Therefore, we kept one side open so participants had to interact actively with the robot to prevent it from falling off the table. We hoped this enforcement of interaction would provide the participants with a better understanding of the robot’s capabilities and behavioural richness. The robot itself has no pre-coded task.

3.3 Groups and Conditions

We decided for two different conditions:

REA (reactive): participants interact for approximately 10 min with a reactive robot and are asked about what they have seen.

ADA (adaptive): same as *REA*, but the robot is continuously adapting, based on maximisation of Predictive Information as a motivation to interact with its environment.

The adaptive robot in the *ADA* condition realises behaviour motivated by maximising predictive information, and it continuously updates its internal networks based on that gradient during the experiment. The reactive robot in the *REA* condition starts with the same networks as the adaptive one (based on pre-trial adaptation), which determines how it reacts to sensor input, but it does not further update its internal network during the experiment.

³ We found that the roll angle readings of the IMU are faulty.

We assign participants into two groups: (A) $ADA \rightarrow REA$ and (B) $REA \rightarrow ADA$. The order of REA and ADA is randomly assigned, but balanced over the number of participants. The starting configuration for both conditions (REA and ADA) was generated in two steps. Firstly, we conducted three trials with the robot for 5 min in the previously described environment. At the end of each trial, we saved the robot’s network configuration. In a second step we randomly choose one of these network configurations as the starting configuration.

The PI formalism allows for having different levels of adaptivity to changing environments and new stimuli. The update rate for ADA was determined empirically. We noticed that the robot can get caught in the pit we mentioned earlier. If the robot gets caught in the pit, it would need to adapt to leave and continue exploration. The ADA adaptation rate was set so that the robot would change its behaviour and leave the pit in less than 20s. As we will discuss later, we hypothesised that a high adaptation rate yields a higher perceived intelligence, as the robot would continuously adapt to new stimuli and change the way it would react to certain inputs.

3.4 Robot’s behaviour

The chosen sensors will determine the behaviour to a large extent, as Predictive Information tries to excite sensor input. For example, if you decide to only use the IMU reading of the yaw angle speed, the robot only needs to adapt its heading in order to excite the sensor. It won’t generate any rolling movement. Empirically, we decided for the pitch and roll angles, the x and y component of the accelerometer, as well as the z component of the gyrometer as sensor input.

The aforementioned strategy for using a fixed network for the reactive robot yields a somewhat predictable behaviour for the condition REA . The robot prefers left turns in light of environmental perturbations or human interaction, i.e. if it hits a wall, it will almost always turn left. Its major trajectory is that of circling in different radii. With this in mind, we assumed that the same arguments as discussed in section 3.1 may hold true and people get bored very quickly. However, almost all participants did not recognise the mentioned pattern. Videos for both conditions are available from [21].

The adaptive robot starts with the same network configuration. We chose a very high update rate for its model, as discussed in section 3.1. Its trajectory has a tendency of being straight, if it reaches an obstacle it adapts its heading to be able to continue moving in another direction. However, as soon as a participant interacts with the robot, it is not trivial to understand what the robot will do next to increase sensory stimuli.

3.5 Goals and Hypotheses

This paper aims to answer the research questions: (1) Is an adaptive robot perceived as more competent/intelligent and animal-like than the reactive robot? and (2) do the reactive and the adaptive robot have distinguishable behaviours? We formulate the following null hypotheses for further investigation:

- $H_0(1)$ The reactive and the adaptive robot’s strategy is described the same.
- $H_0(2)$ The median change for the GodSpeed’s factor Animacy is zero.
- $H_0(3)$ The median change for the RoSAS’ factor Competence is zero.
- $H_0(4)$ The median change for the GodSpeed’s factor Perceived Intelligence is zero.

3.6 Measures

We decide to use two standardised questionnaires to compare results with other studies: the GodSpeed scale [3], which been widely used in many experiments, and the Robotic Social Attributes Scale (RoSAS) [6], which is relatively new and has seen little use in HRI so far.

GodSpeed uses a 5-point semantic differential scale and investigates for the factors Anthropomorphism, Animacy, Likeability, Perceived Intelligence and Perceived Safety. The authors of RoSAS do not recommend a specific scale, but recommend having a neutral value, e.g. uneven number of Likert elements. We decided to use a 7-Likert scale. It tests for the factors Warmth, Competence and Discomfort. Although we are mostly interested in the factors Animacy, Perceived Intelligence, Competence and Warmth, we use all the provided items of both scales. This is done to hide the questionnaire intention and to check for other effects in later research. We will show results for the additional factors Anthropomorphism, Likeability, Perceived Safety and Discomfort for completeness purposes, but we will not discuss them in detail in this paper due to space constraints.

We use these scales for the questionnaires after each condition and we ask two open ended questions in addition: (1) “Can you describe the different behaviours of the robot? Did the robot have any particular strategy for exploring?” and (2) “What were the best and/or worst aspects of the robots behaviour?”.

3.7 Methodology

Participants are welcomed to the experimental room, they are handed an information sheet and are asked to sign an informed consent form. Then the environment and the robot is presented and briefly described. Participants are informed that the robot’s aim is to explore the environment. They are asked to observe whether the robot follows a particular strategy to do so, and if they can identify any specific behaviour. They are also asked to prevent the robot from rolling over the open edge. They are shown how to use the hand as a “wall” or nudge the robot to prevent it from falling off the side of the arena that is not enclosed by a wall, or to illicit new behaviour through interaction. Participants then fill in the pre-questionnaire. This gathers information regarding their sex, age and background. Next, the two conditions are presented to the participants on a randomised order. Each lasts approximately 10 min. They fill in two post-questionnaires containing the two scales and the two additional questions discussed earlier in section 3.6. The entire experiment takes 40 to 50 min.

3.8 Sample

We recruited 16 participants (5 female; 11 male) mostly from university staff and students, between the ages of 23 and 60 years ($M = 33.4$). The participants mostly have a background in Computer Science and were asked about how familiar they are with interacting with robots, programming robots and the chosen robot platform. A 5-point Likert scale was chosen with the value 1 for “not familiar” and 5 for “very familiar”. The self-assessed experience for interacting with robots was an average of 4.3. The average familiarity with programming robots was 3.8 and the experience with the chosen robot platform was rated an average of 1.9. The selection of this group was on purpose, as it was assumed that this group have a more realistic view on robot’s capabilities in general. However, all participants were naïve with regards to the purpose of the experiment.

The study is ethically approved by the Health, Science, Engineering & Technology ECDA with protocol number aCOM/PGR/UH/03018. The anonymity and confidentiality of the individual data is guaranteed.

4 Results

We decided to focus on non-parametric tests to analyse the questionnaire data, as these tests are more robust for small sample sizes. For each factor, we investigate whether there is any interaction between the order and the condition. In case of interaction effects (i.e. order and condition) between these factors, investigating the main effects independently would be incomplete or even misleading.

Our study can be expressed as a F1 LD F1 Model with one within factor (condition) and one between factor (order). A non-parametric ANOVA-type test shows that there is no interaction between the condition and the order ($p > .05$).

Participants were not explicitly asked for differences in the seen behaviour. However, we looked into the answers of the open ended questions to see if they spotted and named differences. All participants answered them, however, answers differed in detail and length. Firstly, we checked whether participants described the behaviour and/or exploring strategy differently to the first interaction. 14 participants described the behaviour and or strategy different to the first one. Mostly, it was pointed out that one robot follows the edges more than the other or that one robot tried to leave the arena more often. In addition, we checked whether the participants used terms like “less”, “compared to” or “this one” for directly addressing changes to the first session after the second session. 8 out of 16 participants did so.

One participant mentioned being “unsure, if something was fundamentally different” but described the robots differently. We conclude to reject hypothesis $H_0(1)$ and accept the alternative hypothesis that participants were able to distinguish between the behaviours.

The effect size r is a robust measure for small sample sizes present in this study. The underlying effect is either small ($r = .10$), medium ($r = .30$) or large ($r = .50$) [7]. Table 1 shows the p -values and effect sizes r for all factors computed

Table 1. Wilcoxon Signed Rank Test results between *REA* and *ADA*

factor	95% confidence interval		p	r
	lower bound	upper bound		
Anthropomorphism	-0.3	0.4	0.916	0.037
Animacy	-0.25	0.333	0.69	0.141
Likeability	-0.3	0.4	0.726	0.124
Perceived Intelligence	-0.2	0.8	0.244	0.412
Perceived Safety	-0.667	0.667	0.444	0.271
Warmth	-0.667	0.167	0.366	0.32
Competence	-0.583	0.5	0.798	0.09
Discomfort	-0.833	0.083	0.141	0.52

with the Wilcoxon Signed Rank Test. It can be seen that there is no statistical significance for any factor.

For the factor Animacy there is only a small effect and thus we accept the hypothesis $H_0(2)$, i.e., there is no change for the factor Animacy between both conditions. The factor competence doesn't show any effect between the conditions and we accept the null hypothesis $H_0(3)$ too.

The Godspeed factor Perceived Intelligence, on the other hand, has a medium effect. We thus can reject the hypothesis $H_0(4)$. From the confidence interval we infer that this effect is in favour for the *REA* condition.

Interestingly, there is a medium effect for Warmth ($r > 0.32$) and a large effect for Discomfort ($r > 0.52$). Both factors are from the RoSA scale and both are in favour for *ADA*.

5 Discussion

It is promising that participants could see differences in the behaviour of the reactive and the adaptive behaviour. One concern with choosing a very simplistic platform is that the magnitude of behaviour differences can be very low. However, to our surprise, the perception of the robot's competence and its animal-likeness was not (significantly) increased for the adaptive robot (*ADA*) compared to the reactive one (*REA*). We think this is due to the experimental design. Having the robot adaptive enough to leave the pit quickly was not as exciting for the participants as for us, and the quick adaptation needed to achieve this made the robot too unpredictable. Some participants even preferred the reactive behaviour because of its predictable and stable pattern. Only two participants mentioned in the open ended questions that the reactive robot's trajectory was mostly circling. We assume that a lower update rate for the adaptive robot will make a difference here and we will address this in another study. What also plays in favour for the reactive robot was the fact that it only rarely approaches the edge not enclosed by a wall, i.e., where the robot could fall off the table. The circling pattern makes it seem to be more alert for some participants, which in turn may have influenced their rating. Initially, we thought that people would

feel that the robot tries to approach them rather than trying to fall off the table. However, only one participant mentioned “it may have sought attention”. We designed the interaction at the edge so the participants could experience the adaptation of the robot to interaction. However, the steady, stable behaviour of the reactive robot was preferred by participants, as it seemed easier to keep the robot from rolling over. This is because the reactive robot approaches the edge less, and the robot’s reactions were more predictive. These may all be reasons for the higher score of Perceived Intelligence for the reactive robot. We also have to rethink the initial hypothesis altogether. At the end of the day, the adaptive robot is exploring its sensor space. There is no goal other than exploring, making it more likely that the adaptive robot is indeed perceived as less competent and intelligent.

We assume a different experiment introduction could have also made a difference already. Rather than saying “also, your task is to prevent the robot from falling off the table”, we could have said that “we did not enclose the edge by a wall, so you can better interact with the robot when it is seeking attention, but please take care it is not rolling over the edge”. However, this would have induced a bias towards the robot’s capabilities. We think we rather need to redesign the environment for future experiments.

Although the results are not overwhelmingly convincing for Predictive Information at first glance, they have to be interpreted in the context of their novelty. Additionally, the unexpected medium effect for Warmth is promising. The factor is created from the items Happy, Feeling, Social, Organic, Compassionate and Emotional and it carries more weight in interpersonal interaction than Competence [6].

6 Conclusion and Future Work

In this work we used a minimal robot platform with only proprioceptive sensors updating the PI model. This makes it hard for the model to infer whether perturbations are induced by, e.g., a participant’s hand or an obstacle like a wall. In other words, adaptivity to the environment and participants is limited by design. In addition, 4 participants mentioned in their open ended answers that they did not see the robot having any “memory” of previous obstacles or explored areas. We want to address this with feeding other sensors into the model in future work. One option is an odometry encoder providing the robot with information about its position in the environment. Another option is using a previously developed proximity sensor for mobile robots for sensing interacting humans [22]. Both these sensors will be investigated to see whether they enhance the behaviour and make it more adaptive and interesting for HRI experiments.

One contribution of this work is the presented baseline behaviour for PI. The experiment shows that the reactive robot, based on a pre-trial adaptation of its networks configurations with PI, is a good enough candidate. Future experiments can be conducted with a baseline behaviour, which is pre-adapted for the very

same sensors and the very same environment as the comparing adaptive PI behaviour.

It can be said that the capacity of the robot to leave the pit was not a driving factor for participants when judging the robot's competence positively. Again, the update rate allowing for that skill may have made the adaptive robot appear unnecessarily random. Further studies are needed for finding a good update rate for the robot.

The medium effect for Warmth for the adaptive robot is very promising. In a next step, we want to redesign the experimental setting, as the enforced interaction supposedly has a negative effect. Rather than enforcing the interaction with the robot, we want to implement a more game-like scenario which allows for the participants to interact when they feel the need or the joy to do so, rather than having to interact without them knowing when it is "intended" by the robot. That way we can investigate if the effect for Warmth is indeed caused by the interaction with the intrinsically motivated robot.

Acknowledgements. CS is funded by the EU Horizon 2020 programme under the Marie Skłodowska-Curie grant 705643.

References

1. Adamides, G., Christou, G., Katsanos, C., Xenos, M., Hadzilacos, T.: Usability Guidelines for the Design of Robot Teleoperation: A Taxonomy. *IEEE Transactions on Human-Machine Systems* **45**(2), 256–262 (Apr 2015). <https://doi.org/10.1109/THMS.2014.2371048>
2. Ay, N., Bertschinger, N., Der, R., Güttler, F., Olbrich, E.: Predictive information and explorative behavior of autonomous robots. *The European Physical Journal B* **63**(3), 329–339 (Jun 2008). <https://doi.org/10.1140/epjb/e2008-00175-0>
3. Bartneck, C., Kulić, D., Croft, E., Zoghbi, S.: Measurement Instruments for the Anthropomorphism, Animacy, Likeability, Perceived Intelligence, and Perceived Safety of Robots. *International Journal of Social Robotics* **1**(1), 71–81 (Jan 2009). <https://doi.org/10.1007/s12369-008-0001-3>
4. Barto, A.G.: *Intrinsic Motivation and Reinforcement Learning*, pp. 17–47. Springer Berlin Heidelberg, Berlin, Heidelberg (2013). https://doi.org/10.1007/978-3-642-32375-1_2
5. Blow, M., Dautenhahn, K., Appleby, A., Nehaniv, C.L., Lee, D.C.: Perception of Robot Smiles and Dimensions for Human-Robot Interaction Design. In: *ROMAN 2006 - The 15th IEEE International Symposium on Robot and Human Interactive Communication*. pp. 469–474 (Sep 2006). <https://doi.org/10.1109/ROMAN.2006.314372>
6. Carpinella, C.M., Wyman, A.B., Perez, M.A., Stroessner, S.J.: The Robotic Social Attributes Scale (RoSAS): Development and Validation. In: *Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction*. pp. 254–262. HRI '17, ACM, New York, NY, USA (2017). <https://doi.org/10.1145/2909824.3020208>
7. Cohen, J.: A power primer. *Psychological Bulletin* **112**(1), 155–159 (Jul 1992)

8. Dautenhahn, K.: I could be you: The phenomenological dimension of social understanding. *Cybernetics and Systems* **28**(5), 417–453 (1997). <https://doi.org/10.1080/019697297126074>
9. Deci, E., Ryan, R.M.: *Intrinsic motivation and self-determination in human behavior*. Springer Science & Business Media, New York: Plenum (1985)
10. Der, R., Martius, G.: *The Playful Machine: Theoretical Foundation and Practical Realization of Self-Organizing Robots*, Cognitive Systems Monographs, vol. 15. Springer-Verlag Berlin Heidelberg (Sep 2012). <https://doi.org/10.1007/978-3-642-20253-7>
11. Fukuda, H., Ueda, K.: Interaction with a Moving Object Affects One’s Perception of Its Animacy. *International Journal of Social Robotics* **2**(2), 187–193 (Jun 2010). <https://doi.org/10.1007/s12369-010-0045-z>
12. Goodrich, M.A., Schultz, A.C.: Human–Robot Interaction: A Survey. *Foundations and Trends in Human-Computer Interaction* **1**(3), 203–275 (2008). <https://doi.org/10.1561/11000000005>
13. Guckelsberger, C., Salge, C., Togelius, J.: New And Surprising Ways to Be Mean. In: 2018 IEEE Conference on Computational Intelligence and Games (CIG). pp. 1–8 (2018). <https://doi.org/10.1109/CIG.2018.8490453>
14. Guckelsberger, C., Salge, C., Colton, S.: Intrinsically Motivated General Companion NPCs via Coupled Empowerment Maximisation. In: 2016 IEEE Conference on Computational Intelligence and Games (CIG). pp. 150–157 (Sep 2016). <https://doi.org/10.1109/CIG.2016.7860406>
15. Martius, G., Der, R., Ay, N.: Information Driven Self-Organization of Complex Robotic Behaviors. *PLoS one* **8**(5), 1–14 (May 2013). <https://doi.org/10.1371/journal.pone.0063400>
16. Merrick, K.E., Maher, M.L.: *Motivated Reinforcement Learning. Curious Characters for Multiuser Games*. Springer, 1 edn. (2009)
17. Oudeyer, P.Y., Kaplan, F.: How can we define intrinsic motivation ? In: *Proceedings of the 8th International Conference on Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems*. Lund University Cognitive Studies, Lund:LUCS, Brighton, Brighton, United Kingdom (2008)
18. Oudeyer, P.Y., Kaplan, F.: What is intrinsic motivation? A typology of computational approaches. *Frontiers in Neurorobotics* **1**, 6 (2009). <https://doi.org/10.3389/neuro.12.006.2007>
19. Ryan, R.M.: Agency and organization: Intrinsic motivation, autonomy, and the self in psychological development. *Nebraska Symposium on Motivation. Nebraska Symposium on Motivation* **40**, 1–56 (1993)
20. Ryan, R.M., Deci, E.L.: Intrinsic and Extrinsic Motivations: Classic Definitions and New Directions. *Contemporary Educational Psychology* **25**(1), 54–67 (2000). <https://doi.org/10.1006/ceps.1999.1020>
21. Scheunemann, M.M.: *Intrinsically Motivated Autonomy in Human-Robot Interaction: Human Perception of Predictive Information in Robots*. Online (2019), <https://mms.ai/TAROS2019-supplementary/>, supplementary material
22. Scheunemann, M.M., Dautenhahn, K., Salem, M., Robins, B.: Utilizing Bluetooth Low Energy to recognize proximity, touch and humans. In: 2016 25th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN). pp. 362–367. IEEE (Aug 2016). <https://doi.org/10.1109/ROMAN.2016.7745156>