# Human Perception of Intrinsically Motivated Autonomy in Human-Robot Interaction

Marcus M. Scheunemann<sup>1\*</sup> Christoph Salge<sup>1</sup> Daniel Polani<sup>1</sup> Kerstin Dautenhahn<sup>1,2</sup>

<sup>1</sup>Adaptive Systems Research Group, University of Hertfordshire, Hatfield, United Kingdom <sup>2</sup>Social and Intelligent Robotics Research Laboratory, University of Waterloo, Waterloo, Canada

November 30, 2021

#### Abstract

A challenge in using robots in human-inhabited environments is to design behavior that is engaging, yet robust to the perturbations induced by human interaction. Our idea is to imbue the robot with intrinsic motivation (IM) so that it can handle new situations and appears as a genuine social other to humans and thus be of more interest to a human interaction partner. Human-robot interaction (HRI) experiments mainly focus on scripted or teleoperated robots, that mimic characteristics such as IM to control isolated behavior factors. This article presents a "robotologist" study design that allows comparing autonomously generated behaviors with each other, and, for the first time, evaluates the human perception of IMbased generated behavior in robots. We conducted a within-subjects user study (N = 24)where participants interacted with a fully autonomous Sphero BB8 robot with different behavioral regimes: one realizing an adaptive, intrinsically motivated behavior and the other being reactive, but not adaptive. The robot and its behaviors are intentionally kept minimal to concentrate on the effect induced by IM. A quantitative analysis of post-interaction questionnaires showed a significantly higher perception of the dimension "Warmth" compared to the reactive baseline behavior. Warmth is considered a primary dimension for social attitude formation in human social cognition. A human perceived as warm (friendly, trustworthy) experiences more positive social interactions.

**Keywords** Intrinsic Motivation, User Study, Human-Robot Interaction, Social Cognition, Predictive Information, Embodied Cognition

<sup>\*</sup>Contact corresponding author using marcus@mms.ai or visit https://mms.ai.

# 1 Introduction

This article is part of a larger research program to produce autonomous robots, i.e., robots that are not teleoperated or remotely controlled, yet robust to the unknown perturbations and capable of sustained interaction with humans. In this study, we look at the effects that an intrinsically motivated robot behavior has on the social perceptions of robots by humans, and whether they may engage human participants in an interaction. An example of an intrinsically motivated behavior is a child interacting with a puppy. The child will likely be motivated to play with a puppy, even without an external reward (such as promised money) and even without the existence of an extrinsic reward (such as playing with the puppy as a means to an end, i.e., to train it). Instead, the motivation for the interaction might result purely from wanting to do this activity for its own sake, i.e., the child is intrinsically motivated to play with the puppy. Our long term goal with this research is to better understand how to engage the human, similar to the child in the example. But here we do not focus on modeling the intrinsic motivation of the human interaction partner, but rather focus on using a computational model of intrinsic motivation to generate robot behavior. Our idea is that a robot that is intrinsically motivated (driven by a specific, adaptive IM model) is perceived more like a social other, and thus is more engaging for a human interaction partner.

There are approaches to keep humans engaged in the interaction with a robot. For example, Pinillos et al. (2016) developed an autonomous hotel robot. It attracts the attention of the hotel guests, many of them wanting to know more about the robot itself. They propose that the robot's services (i.e., its competence or usefulness) need to be large in order to keep customers engaged. On the other hand, Kanda et al. (2010) developed a semi-teleoperated mall robot and incrementally added novel behaviors, such as self-disclosure. A field trial indicates that the robot attracted reoccurring visitors, without increasing its services. Engagement is also a concern in the field of social robotics in education (Belpaeme et al., 2018). One existing approach here is to develop robots with a set of hand-designed questions, comments, and statements (Gordon et al., 2015; Ceha et al., 2019). This makes the robots *appear* curious, which elicits curiosity in the humans too, which in turn enhances learning and memory retention (Oudeyer et al., 2016), or is even used synonymously with intrinsic motivation (Schmidhuber, 1991).

The previous studies either constrained the context and focused on a specific task (e.g., Gordon et al., 2015; Pinillos et al., 2016), or were relying on humans teleoperating the robot (e.g., Kanda et al., 2010; Ceha et al., 2019). Teleoperation, or the Wizard-of-Oz model, remains the state of the art for many HRI studies (Clabaugh & Matarić, 2019). This is due to the challenge to define a sufficient set of execution rules (i.e., behaviors) for an HRI task; this holds true even in a laboratory setting. It remains elusive to achieve autonomous, social behavior in an unconstrained environment, i.e., for any given task or goal in the real world (Christensen et al., 2016; Belpaeme et al., 2018). Developing a robot driven by an actual intrinsic motivation formalism, such as the drive to explore its environment and its capabilities, might offer a solution.

If successful, this would provide us with a robust behavior generation mechanism that allows us to "Escape Oz" (Clabaugh & Matarić, 2019), while also producing behavior that appears curious, or similarly engaging to the human interaction partner. This will reduce the reliance on human adaptation or teleoperation, and could provide a promising pathway towards having robots more easily deployed in everyday life.

Our idea is that imbuing a robot with a computational model of intrinsic motivation (IM) makes the robot appear as a genuine *social other* – similar maybe to an animal – and thus be of more interest to a human interaction partner. The concept of intrinsic motivation originates in psychology, initially in close relation to Self-Determination Theory (SDT) (Ryan & Deci, 2000b). SDT posits that humans have an inherent tendency to seek out novelty and challenges, to extend and exercise their capacities to explore and to learn, without having to be coerced by an extrinsic reward. According to SDT, humans have inherent drives for competence, autonomy, and relatedness. Computational models of intrinsic motivation aim to formalize the principles that create those drives to make them operational, i.e., they can be used to create spontaneous exploration and curiosity in an artificial agent (Oudeyer & Kaplan, 2009). Therefore, we hypothesize that they will give an artificial agent a stronger social presence, and thus make them a more interesting interaction partner.

The known models of intrinsic motivation have a range of interesting properties. The idea of universality is of particular interest for this application, in particular the fact that IMs can cope with changes to an agent's environment or its morphology (we discuss this in more detail in subsection 2.1, 2.3, and 3.3). This makes this approach, in principle, suitable to be deployed on any robot and it also allows it to deal with any environment or context. The biggest limitation here is usually computational complexity. The method is also limited by the fact that several approaches at least require agent-centric forward models, similar to sensorimotor contingencies (O'Regan & Noë, 2001), which might not be easily obtainable. Finally, most IMs can be expressed to operate on the immediate perception-action loop of the robot, allowing for tightly coupled or entrained behavior with both the environment or other actors. Both of these properties make IMs an interesting family of approaches to deploy in autonomous human-robot interaction (HRI) robots, as there is a requirement for interactive feedback on a short feedback loop and for the ability to robustly deal with a range of situations. This is particularly relevant, as human social cognition is believed to heavily depend on interaction – and thus any approach that aims to encourage interaction should be robust to the perturbations induced by those social, and possibly physical interactions.

In the remainder of this article, we want to substantiate this main idea with an HRI study involving 24 human participants. As a first step, we evaluate the human perception of robots with different behavior with the help of post-experiment questionnaires. We compare how the introduction of intrinsically motivated behavior affects human perception, and discuss how these factors can lead to the formation of different social attitudes. Our main focus in this article is on the *Warmth* dimension. Warmth and *Competence* are considered the two main dimensions in describing almost all social attitudes in human social cognition, such as friendliness, empathy, admiration, envy, contempt and pity (Fiske et al., 2007; Abele et al., 2016). Warmth is considered the primary dimension for social characterizing peers. This means, when characterizing other people, we firstly judge their intent (Warmth) before judging their capability (Competence) to enact their intent. Warmth is strongly linked to the measure of trust (Fiske et al., 2007; Fiske, 2018). A person who is perceived as warm is also perceived as more trustworthy. For example, Kulms & Kopp (2018) use it as an indicator of people's trust in computers. Importantly, from human social cognition, it is known that human's who are perceived as warm (Fiske et al., 2007). Recent research in HRI has shown that human participants prefer to interact again with the robot if they perceive its behavior as more warm (Oliveira et al., 2019; Scheunemann et al., 2020). Consequently, in order to welcome robots in our everyday life, an understanding is needed for how to enable the perception of Warmth for robots.

We will see that the robot that continues adaptation based on its intrinsic motivation, i.e., the intrinsically motivated robot, generates behavior that participants rated as more warm compared to the baseline behavior. This is a step towards the long-term goal of producing a robot capable of sustained interaction, as it suggests a method to induce a positive social attitude towards the robot in the human. Further studies are, of course, needed to see if this effect for Warmth transfers from human-human interaction to human-robot interaction. We also still need to investigate if higher perceived Warmth for a robot actually leads to more sustained interaction. The interplay between personality and social relationships is still an ongoing – and complex – investigation for human-human interaction (Geukes et al., 2019). Our expectation, which needs to be confirmed in future work, is that a robot which is perceived as warm (friendly, trustworthy) is more likely to receive more positive interactions from people and will facilitate long-term interactions.

#### **Overview**

First, section 2 outlines the background on intrinsic motivation, its computational approaches, and its relation to autonomy, insofar it relates to the present work. Section 3 then introduces time-local predictive information (TiPI), the information theoretic formalism we use to implement intrinsic motivation in our studies. We outline the concrete approximations (and their assumptions) to compute PI. In particular, we highlight how to make this approach suitable for deployment on an actual robot and why it is a good candidate for our research questions.

Section 4 presents the materials and methods of our within-subjects study (N = 24). The study consists of two conditions with the same robot platform: the intrinsically motivated behavior in one condition is generated using predictive information maximization, the behavior in the other condition is a reactive baseline behavior. The focus is on the interplay between the robot and the human participant. We designed a study where the participants interact and observe the robots in order to understand their behavioral differences, but they are not given information about the robot's task and they cannot order the robot to do something. Instead,

the robot and the human participants explore their behavior towards each other. We call this "robotologist" study design (for details see subsection 4.3). The study design is motivated by a preliminary study by Scheunemann et al. (2019) which has been conducted and published prior to this work. We outline the learned lessons from this previous study where needed.

Section 5 presents the results, concentrating on our two main hypotheses: one focusing on the perceived Warmth of the intrinsically motivated robot behavior, and the other on the lack of difference in Perceived Intelligence and Competence between the robot scenarios. We found that our study design makes both robot behaviors appear similarly competent. This is important in order to focus solely on the effect on the Warmth dimension, without interfering by the Competence rating. Most importantly, the study provides evidence that the intrinsically motivated robot displays behavior that is perceived as more warm compared to the reactive baseline behavior. Section 6 discusses the implications of those findings, and how they can be applied to other projects. Section 7 summarizes the study and concludes the article.

# 2 Background

This section provides some background of the previously mentioned concepts relating to this work.

#### 2.1 Intrinsic Motivation

A common definition of intrinsic motivation (IM) in psychology is "doing [...] an activity for its inherent satisfactions rather than for some separable consequence" (Ryan & Deci, 2000a). An intrinsically motivated agent is moved to do something, for the enjoyment of the activity itself, for the "fun or [the] challenge entailed rather than because of external products, pressures, or rewards" (Ryan & Deci, 2000a, pg. 56). Since intrinsic motivations have been considered an instrumental ingredient in the development of humans (Oudeyer et al., 2007), there has also been a great interest in developmental robotics to produce formalized models that can be used to imbue robots with drives for competence and knowledge acquisition. Although concepts like "fun" and "challenge" are presented as crucial for the definition of IM in psychology, the literature lacks consensus on what these concepts are (Oudeyer & Kaplan, 2009). This missing consensus and the resulting vagueness of the definition makes it impossible to transfer it directly onto a robotic system. Oudeyer & Kaplan (2008) characterize intrinsic motivation in the following, broadly accepted way:

An activity or an experienced situation, be it physical or imaginary, is intrinsically motivating for an autonomous entity if its interest depends primarily on the collation or comparison of information from different stimuli [...]. [...] the information that is compared has to be understood in an information theoretic perspective [...], independently of their meaning. As a consequence, measures which pre-suppose the meaning of stimuli, i.e., the meaning of sensorimotor channels (e.g., the fact that a measure is a measure of energy or temperature or color), do not characterize intrinsically motivating activities or situations.

Nowadays, there is a range of formal models that roughly fall under the header of intrinsic motivation, such as the autotelic principle (Steels, 2004), learning progress (Kaplan & Oudeyer, 2004), empowerment (Klyubin et al., 2005), predictive information (Der et al., 2008; Ay et al., 2008), the free energy principle (Friston, 2010), homeokinesis (Der & Martius, 2012) and others. These models have a range of commonalities: they are free of semantics, task-independent, universal and can be computed from an agent's subjective perspective. Most of the work related to IMs focuses on how they create *reasonable* behavior (in some suitable sense) for simulated agents. There has been some work in the domain of computer games that focuses more explicitly on the relationship between intrinsically motivated agents and humans, and how an intrinsic motivation could generate more believable Non-Player Characters (NPCs) (Merrick & Maher, 2009), or produce generic companions (Guckelsberger et al., 2016) or antagonist behavior (Guckelsberger et al., 2018). So far, IMs have been deployed on simulated and physical robots (e.g., Oudeyer et al., 2007; Der & Martius, 2012; Martius et al., 2014), but, as far as we know, there has been no human-robot interaction study yet evaluating the perception of intrinsically motivated robots from the perspective of humans. In this work we use predictive information maximization to implement an autonomous, intrinsically motivated robot. We describe the formalism in more detail in subsection 2.3.

## 2.2 Autonomy

The term *autonomy* is used with multiple meanings (Boden, 2008). When we talk about autonomous robots, we merely mean robots that are not directly controlled by a human operator, autonomy just being a dimension of the experimental design (Huang et al., 2004; Stubbs et al., 2007). In SDT, however, autonomy refers to being in control of one's own life, which can be seen as a close enough analogy for living systems (Paolo, 2004). SDT also assumes that there is a drive to maintain this state of autonomy, which we do not see in general with autonomous robots. We might see autonomy used as the idea that a robot should strive to maintain operational autonomy, i.e., not be in need of external help, but it usually does not refer to a robot striving to not be controlled by a human. Finally, autonomy might also be referring to the concept of self-making or self-law-giving, which is closely related to autopoesis (Maturana & Varela, 1991; Froese & Ziemke, 2009). In robots, this is currently only a theoretical idea (Smithers, 1997). but it is often considered necessary for *true* intrinsic motivation. Any heteronomy during the development or creation of an agent would ultimately make them extrinsic and hence undermine their very nature, i.e., computational models of intrinsic motivations on robots are usually put on those robots by humans, and are thus actually extrinsic. Computational models of intrinsic motivation are an attempt to merely reproduce the behavior or functionality of genuinely intrinsic motivation in organism. This is also the reason that we talk about *perceived* agency

and *perceived* autonomy. One idea behind this is that by using those models for the robots to *pretend* to be intrinsically motivated, humans might indeed perceive the robot as thus. In the following, when we talk about intrinsic motivation on the robot we exclusively refer to the initial, technical meaning, the computational model that aims to mimic intrinsic motivation. The more philosophical underpinnings of autonomy are highly relevant to the larger context of this work and indicate that this approach is useful even if we develop robots with more extensive autonomy, making it a robust approach, even for more self-directed robots in the future. Here, the main purpose of this section was to clarify that there are different levels of autonomy – so it is clear that when we talk about intrinsically motivated autonomy, we do not just speak of a robot that can move by itself, but one that can self-directly change its behavior, based on a goal that is at least aligned with its own agency.

#### 2.3 Predictive Information

This section describes predictive information (PI) maximization, the intrinsic motivation model used for the robot behavior generation in our experiments. PI has been described as early as 1986, termed *effective measure complexity* (Grassberger, 1986) or *excess entropy* (Crutchfield & Young, 1989). Previous work with PI-driven robots in simulation demonstrated its applicability to a large range of different robot morphologies (Der et al., 2008; Martius et al., 2013b; Zahedi et al., 2013; Martius et al., 2014). A range of existing videos from experiments in simulation showcase apparent exploratory, playful, and open-ended behavior of individual robots and robot collectives (see Research Network for Self-Organization of Robot Behavior, 2015). The PI-induced behavior in the videos suggests PI as a promising immediate candidate measure to test our core idea.

Conceptually, when this measure is transformed into a behavior-generating rule, the resulting dynamics essentially fall into a family of learning rules related to the reduction of the time prediction error in the perception-action loop of a robot (see especially the book The Playful Machine, Der & Martius, 2012). The aforementioned book also shows how these approaches can be computed from the robot's perspective alone. Additionally, the variety of different robots and their behaviors presented there shows how different behaviors arise from the same formalism due to the sensitivity towards the agent's specific embodiment.

The predictive information formalism consists of computing a specific learning rule that aims to maximize the mutual information between a robot's past and future sensor states (Ay et al., 2008), i.e., PI quantifies how much information a history of past sensor states contain about future sensor states. More generally, predictive information is defined as the mutual information between the past and the future of a robot's sensor input. A high amount of predictive information requires two things: First, past sensor states should make future sensor states more predictable. This should lead the robot to act so that its actions have predictable consequences. Furthermore, the robot also needs to create a high variety of sensor input. If the robot would always perceive the same sensor input, then there is either insufficient information in the past to predict future sensor states, or an insufficiently varied future for which there is not much to predict. In both cases, an impoverished sensor input reduces the predictive information. Alternatively, if there is strong variation in the sensor input but little structure in the sensor data stream, i.e., the past has little to do with the future, that also leads to low predictive information. Vice versa, a high value for predictive information requires a high entropy in future sensor states, i.e., a richly varied future (a robot motivated to *excite* its sensors to reach a rich variety of different states) which at the same time depends on the observable past (i.e., which the robot can predict well based on the past). The behavioral regime is created by these two counterpoised requirements: predictability and variety. This yields a robot wanting to act so that its future is highly predictable, while exploring and experiencing new sensor states. The PI literature argues that this balancing act produces rich exploratory behavior that is sensitive to the robot's embodiment and argues that predictive information is "the most natural complexity measure for time series" (Bialek et al., 2001; Martius et al., 2013b).

A robot which acts depending on the maximization of PI only compares its sensors channels on an information theoretic level, without the need of pre-defining any meaning to the sensors. The quantity, therefore, falls in the characterization of IM presented by Oudeyer & Kaplan (2008) (see subsection 2.1) and it is a candidate measure to enable intrinsically motivated autonomy in a robot.

Der et al. (2008) and Ay et al. (2008, 2012) presented derivation rules for PI, which allows for computing the model directly for linear systems with stationary dynamics. The next section discusses an extension of their work by Martius et al. (2013b) for the use in nonlinear and nonstationary systems – such as physical robotic systems.

# 3 Time-local predictive information

The predictive information formalism to generate the robot's intrinsically motivated behavior in the studies of this article is closely following the implementation of Martius et al. (2013b). They propose an approximation to compute PI for nonlinear systems with nonstationary dynamics, which allows for behavior development of a self-determined robotic system. They approximate PI with assuming small, Gaussian noise and only consider a time window over the current state of the robot and  $\tau$  steps back in the past, coining it time-local predictive information (TiPI). TiPI allows for going beyond discrete finite-state actions, which still dominates scenarios of information theory-based behavior generation, towards continuous actions. This permits using physical robots in high-dimensional state-action spaces. TiPI enables robot behavior with self-switching dynamics in a simple hysteresis system and spontaneous cooperation of physical coupled systems (Martius et al., 2013b).

TiPI works by updating the two internal neural networks of the robot, one that generates behavior from sensor input and the other that predicts the future states. The continuous adaptation, aimed at improving the TiPI, moves the robot through a range of behavioral regimes. Importantly, the changes in behavior 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 could be adapted for individual preferences (Der & Martius, 2006).

The approach allows to change the robot's morphology without having to redesign the algorithm, but still remaining sensitive to the embodiment of the robot, meaning that the resulting behavior differs, depending on how the robot interacts with the world. The morphology can be changed by changing physical parts or by choosing different sensors as inputs for the robot's neural networks. In both ways, the robot can be guided towards exploring and playing in different ways. For example, by including a sensor for the robot's angular velocity around its main axis, the spherical robot would try to spin clockwise and anticlockwise with changing velocities. If we further include an accelerometer providing measurements of the forward and backward acceleration, the robot would try to explore the relationship between spinning movements and locomotion, yielding a variety of additional motion patterns. If, furthermore, a human is interacting with the robot, this can increase the behavioral diversity, depending on the interaction between the robot and the human.

#### 3.1 Deriving update rules

Martius et al. (2013b) present estimates of the time-local predictive information (TiPI) for general stochastic dynamical systems. For systems with Gaussian noise and with gradient ascent on the TiPI landscape, they derive explicit expressions for exploratory dynamics. This subsection introduces the derivation of the explicit expression. The derivations are kept short providing only the basic concepts of the quantity and introduce the underlying main approximations and assumptions that need to be considered when applying the algorithm to a robot in an HRI scenario. For a detailed treatment, the reader should refer to (Ay et al., 2008; Martius et al., 2013b).

Assume a robot has n sensors and the sensor readings are polled in constant time steps  $(\Delta t = 1)$ . Combine now the result of all sensor values in a vector  $s \in \mathbb{R}^n$ . A series of those sensor readings between points of time a and b (with a < b) can be described as a time-discrete process  $\{S_t\}_{t=a}^b$ , where both boundaries are included. Let the past be defined by the points of time  $a, \ldots, t-1$  and the future by  $t, \ldots, b$ . Bialek et al. (2001) defines the PI for some point in time t for the time series S as the mutual information between the past and the future. Intuitively, the mutual information measures the shared information of two random variables, here  $S_{\text{past}}$  and  $S_{\text{future}}$ , i.e., it measures how much knowledge of the past  $S_{\text{past}}$  reduces the uncertainty of the future  $S_{\text{future}}$ . The predictive information, expressed as mutual information, is thus defined as follows

$$I(S_{\text{future}}; S_{\text{past}}) = \left\langle \ln \frac{p(s_{\text{future}}, s_{\text{past}})}{p(s_{\text{future}})p(s_{\text{past}})} \right\rangle$$
$$= H(S_{\text{future}}) - H(S_{\text{future}}|S_{\text{past}})$$
(1)

with the average taken over the joint probability density distribution  $p(s_{\text{past}}, s_{\text{future}})$ .

The first essential simplification proposed by Martius et al. (2013b) is applying the Markov assumption to Equation 1. If  $\{S_t\}_{t=a}^b$  is a Markov process, all past information relevant to the future is stored in the very last state of the system, i.e.,  $S_{\text{past}} = S_{t-1}$ .

The predictive information in this case reduces to:

$$I(S_t; S_{t-1}) = \sum_{s_{t-1} \in S_{t-1}} \sum_{s_t \in S_t} p(s_t, s_{t-1}) \ln\left(\frac{p(s_t, s_{t-1})}{p(s_t)p(s_{t-1})}\right)$$
$$= H(S_t) - H(S_t|S_{t-1}) .$$
(2)

In general, the Markov assumption only holds true for real-world sensor processes in exceptional cases. Nonetheless, as in the wide use of, e.g., particle or Kalman filters, it is a popular assumption for successfully approximating problems using a Bayesian approach (Thrun et al., 2005). Martius et al. (2013b) use the reduced Equation 2 as the definition of the objective function for deriving the autonomous exploration dynamics.

Equation 2 is a quantity derived for the whole process. However, to create an actual behavior rule that reacts to the current situation, it necessary to compute a local quantity, specific to the current situation. Therefore, instead of computing the probability distribution  $p(s_t)$  over the whole process, we additionally condition the PI on a state  $s_{t-2}$ . The new quantity derived is then

$$I(S_t; S_{t-1}|s_{t-2}) (3)$$

Because of above Markovianity, this is effectively a time-local quantity for PI and therefore it is called *time-local* predictive information (TiPI). To calculate the TiPI, a model of  $S_t$  needs to be learned to predict its time series. Let  $\psi = \mathbb{R}^n \to \mathbb{R}^n$  be a function predicting the time series at t - 2, t - 1, and t via

$$\hat{s}_{t-2} = s_{t-2}$$
 (4)

$$\hat{s}_{t-1} = \psi(s_{t-2}, \theta_{t-2}) \tag{5}$$

$$\hat{s}_t = \psi(\psi(s_{t-2}, \theta_{t-2}), \theta_{t-1})$$
 (6)

preprint

In an example implementation (Martius, 2013),  $\psi$  is realized as a one-layer neural network.  $\theta$  is a set of parameters representing the synaptic weights and biases that are updated each time step in order to increase TiPI. The actual dynamics of the process can be described via

$$s_t = \psi(s_{t-1}, \theta_{t-1}) + \xi_t \tag{7}$$

 $\xi_t$  being the prediction error.

We denote the deviation of the actual dynamics (Equation 7) from the deterministic prediction (Equation 6) as

$$\delta s_{t'} = s_{t'} - \hat{s}_{t'} \tag{8}$$

for any time t' with  $t-2 \le t' \le t$ . Since  $s_{t-2}$  is the initial state for TiPI, there is no deviation at time t-2 and  $\delta s_{t-2} = 0$ , while one step after the initial state  $\delta s_{t-1} = \xi_{t-1}$ . Intuitively,  $\delta s_t$ represents the prediction error(s) accumulated from the start of the prediction (here at t-2) up to time t.

For very small prediction errors the dynamics of  $\delta s$  (Equation 8) can be linearized as an approximation:

$$\delta s_{t'} = L(s_{t'-1})\delta s_{t'-1} + \xi_{t'} + O(||\xi_t||^2) \tag{9}$$

with the Jacobian

$$L_{ij}(s) = \frac{\partial \psi_i(s,\theta)}{\partial s_j}$$

Assuming that the prediction errors  $\xi$  are both small and Gaussian, the TiPI on the deviation process  $\delta S_{t'}$  is the same as on the original process  $S_t$  (see Martius et al., 2013a, sec. A). It is therefore sufficient to concentrate on the error propagation for the computation of the TiPI. This reduces Equation 2 in such a way that only the probability distribution of the deviation  $p(\delta s)$ needs to be known, rather than the probability distribution over the full state p(s).

If we further assume that the prediction error  $\xi$  is white Gaussian, the entropy can be expressed as covariances (Cover & Thomas, 2012). The resulting explicit expression of TiPI on  $\delta S$  becomes:

$$I(\delta S_t; \delta S_{t-1}|s_{t-2}) = \frac{1}{2} \ln |\Sigma_t| - \frac{1}{2} \ln |D_t|$$
(10)

with  $\Sigma = \langle \delta s \, \delta s^T \rangle$  as the covariance matrix of the process  $\delta S$ , and  $D = \langle \xi \xi^T \rangle$  as the covariance matrix of the prediction error. Note that the predictive information becomes meaningful only at t, as the prediction error vanishes at t-2 and at t-1 the two covariance matrices coincide:  $\Sigma_{t-1} = D_{t-1}$ . The covariances are exact for Gaussianity. For the general case they are approximations only. We now give the algorithm used to drive a robot's behavior towards increasing TiPI. Martius et al. (2013b) derive it explicitly for the gradient ascending neural network presented in Equation 6. They argue that the prediction error  $\xi$  is essentially noise and does not depend on the parameter of the controller, and that therefore the term  $\ln |D|$  of Equation 10 can be omitted when computing the gradient. Based on Equation 10, the resulting gradient step executed at each time t is

$$\Delta \theta_t = \epsilon \frac{\partial I}{\partial \theta} = \epsilon \frac{\partial}{\partial \theta} \ln |\Sigma_t| \tag{11}$$

with  $\epsilon$  being the update rate and  $\theta_{t+1} = \theta_t + \Delta \theta_t$ .

Applying Equation 9 to above equations results in explicit gradient step

$$\Delta \theta = \epsilon \left\langle \delta u_t^T \, \frac{\partial L(s_{t-1})}{\partial \theta} \, \delta s_{t-1} \right\rangle \tag{12}$$

where  $\delta s$  and the auxiliary  $\delta u$  are given as

$$\begin{split} \delta s_{t-1} &= s_{t-1} - \psi(s_{t-2}, \theta_{t-2}) \\ \delta s_t &= s_t - \psi(\psi(s_{t-2}, \theta_{t-2}), \theta_{t-1}) \\ \delta u &= \Sigma_t^{-1} \delta s_t \\ \Sigma_t &= \langle \delta s_t \, \delta s_t^T \rangle \end{split}$$

To render  $\Delta \theta$  computable, Equation 12 is further approximated by applying the *self-averaging* property (we explain this in more detail below) of a stochastic gradient

$$\Delta \theta = \epsilon \, \delta u_t^T \, \frac{\partial L(s_{t-1})}{\partial \theta} \, \delta s_{t-1} \tag{13}$$

As per (Der et al., 2008; Martius et al., 2013b), Equation 13 is the equation by which the (approximate) TiPI maximization is ultimately implemented. We remark that increasing  $|\Sigma|$  corresponds to an increase of the norm of  $\delta s$ . In other words, this reflects the amplification of small fluctuations in the motor dynamics, i.e., an increase of the instability of the system dynamics.

#### 3.2 Approximations and assumptions

Along with the above derivation, several approximations and assumptions have been made. When the measure is applied to a real robot in a real-world human-interaction scenario, this requires a careful consideration of the assumptions and approximations, which we do in the following.

#### 3.2.1 Markov assumption

This assumption simplifies the definition of the objective function Equation 2. More importantly, it renders TiPI (Equation 3) computable as it simplifies the conditional probability density distribution. Applying the assumption to robotics-related problems, especially to make Bayesian problems manageable, is common in robotics (Thrun et al., 2005). This approximation therefore can be considered a popular robotics strategy for applying information theory and Bayesian algorithms to the real world.

#### 3.2.2 Conditioning on an initial state from two states back

To compute PI for nonlinear systems with nonstationary dynamics, the proposed solution is to condition the quantity on an initial state from two steps back in time. We stick here to the minimal possible window mainly because computing a larger window online comes to a computational cost challenging to bear on embedded systems.

The sensors used for the input need to be meaningful for the time window. For example, a global position of the robot does not change much within the time window of two steps, so the robot cannot excite the sensor value in the chosen window. It is therefore preferable to choose sensors which display variation within the given time window, such as proprioceptive sensors measuring the acceleration or velocity.

#### 3.2.3 Prediction errors are both: very small and Gaussian

These assumptions are made at various places for deriving the explicit update rules. For example, the assumptions were used to show that TiPI on the process  $\delta S$  (propagation of errors) is equivalent to the one on the original process S (sensor states). This enables the linearization of the error dynamics Equation 9 and eventually, under the same assumptions, the formulation of explicit TiPI expressions (Equation 10). Assuming that the error is very small and Gaussian has implications on choosing the right sensors for the experiments. Therefore, care needs to be taken that the noise of the sensors remains somewhat Gaussian and somewhat small for the duration of the time window. For example, the motor position typically changes in a continuous fashion and therefore the respective sensors are good candidates to fulfill this assumptions.

On the contrary, it would violate the Gaussianity assumption to use a sensor whose values exhibit, e.g., sudden drops, such as proximity sensors based on Bluetooth (Scheunemann & Dautenhahn, 2017). Such sensors measure the signal strength to an external device which is prone to occlusions and can sometimes intermittently fail to provide any reading at all. To mitigate this, it is possible to use filters to smoothen the sensor readings.

#### 3.2.4 Applying the self-averaging property for stochastic gradients

Equation 13 uses the so-called self-averaging property of stochastic gradients, which means that a stochastic gradient over a larger number of steps in a sequence acts as an approximation of averaging over the probability distribution (Van Rensburg et al., 2001). In other words, we can replace the average over multiple independently drawn samples by a one-shot gradient.

Practically, this makes Equation 10 computable, as the density distribution of the gradient is hard to obtain. Martius et al. (2013b) note that using this property is only exactly valid for a small update rate  $\epsilon$  when it is driven to zero eventually. Note that the update rate  $\epsilon$  in our application is quite large to allow for a very fast adaptation process. Martius et al. (2013b) argue that the explicit update rules favor the approach of getting an "intrinsic mechanisms for the self-determined and self-directed exploration", with the exploration being driven only by the sensor values. Thus, the one-shot nature of the gradients favors the explorative nature of the exploration dynamics and increases interesting synergy effects, but is not strictly implementing the average.

#### 3.2.5 Noise is independent of the controller parameters

To derive the explicit update rules (Equation 10), the covariance of the noise  $D = \langle \xi \xi^T \rangle$  is omitted altogether. The propagation in error is only assumed to be pure independent noise in the environment. In other words, the noise is independent of the controller parameter  $\theta$ . Martius et al. (2013b) justify this because of the "parsimonious control" implemented by the formalism.

All these assumptions are of course no longer strictly valid once the robot interacts with the environment, especially humans. Nevertheless, the intended richness of the robot's behavior is not hampered by that. Instead, the formalism gives rise to a varied and manifold repertoire of behaviors, as shown by many studies mentioned in (Ay et al., 2008; Der & Martius, 2012; Martius et al., 2013b, 2014).

#### 3.3 Applicability for an HRI study

Martius et al. (2013b) apply the above maximization of TiPI to simulated robots. As a result, those robots show complex behavior (Martius et al., 2013b). One example is a humanoid robot with 17 degrees of freedom (DoF) controlled by a single high-dimensional controller implementing the PI optimization principle from Equation 13. Importantly, despite using the same rules, the formalism produces different behavioral regimes of the simulated humanoid, depending on the environment it is exposed to. Its universality for different embodiments and nonstationary settings makes it a good candidate for applying it to a robot without concerning oneself too much with the environment or the robot's particular embodiment. However, this formalism on its own does not result in high-level behavior, such as walking or serving a human, which may be directly usable in traditional HRI scenarios. It is not a general control algorithm which may generate extrinsic motivation, such as helping a human to get its battery charged. This raises the question of how to align this kind of behavior generation with behavior a designer of, e.g., an assistance robot would want? Both issues can be addressed by combining IM-based behavior generation with scripted behavior, or behavior based on extrinsically rewarded reinforcement

learning. In the study presented in the next section, we deliberately leave these things out and focus on an empirical study of the effect of intrinsic motivation on its own. In other words, we intend to investigate the effect of intrinsically motivated autonomy in isolation from additional criteria and methods for behavior generation. This evaluation is completely missing from the existing body of work on TiPI or intrinsically motivated autonomy in general. This is the gap this article aims to fill.

# 4 Materials and methods

This section describes the materials and methods. Note that design choices have been motivated by lessons learned from our preliminary study (Scheunemann et al., 2019; Scheunemann, 2021a, Chapter 2). We point out the differences to the present study where relevant.

#### 4.1 Robot



Figure 1: a: The used robot platform BB8 from Sphero. b: A 2-D cross-sectional view of the robot. A two-wheel vehicle (darker shape), kept in position by a heavy weight, moves the sphere when driving. The speed of each servo motor can be set individually, allowing the robot to move straight, to turn and to spin. A magnet attached to the vehicle keeps the head on top of the sphere facing towards the movement direction.

We use the off-the-shelf spherical robot from the company Sphero, specifically, the BB8 platform (Sphero, Inc., 2020), as depicted in Figure 1. BB8 is a character from the *Star Wars* movies (Lucasfilm Ltd., 2015). A magnet keeps the head in the driving direction, which gives the user a sense of the robot's direction. We believe that this helps the human participants to interact with the robot.

A humanoid robot platform may raise expectations of advanced social capabilities in participants (Dautenhahn, 2004; Hayashi et al., 2010). For example, humans would expect the robot to have similar eyesight as themselves, or be able to speak and gesture similarly. This in turn may interfere with the investigation of how IM is being perceived. We wanted to reduce as many factors as possible to maximize the focus on the effects induced by TiPI solely, which is why we decided to use a non-humanoid platform with few degrees of freedom.

The robot's on-board hardware is proprietary which prevented us from flashing it or running our code. However, it is possible to communicate with the robot using Bluetooth Low Energy (BLE). We can, for example, request a stream of sensor information from the robot, or control the robot by either using (i) the robot's balance controller or (ii) directly setting the speed of each servo.

The balancing controller receives speed and heading as input values. The heading is globally initialized to zero degrees when the robot is started. This means, on sending 20° to the controller, the robot always sets its heading towards 20 degrees on top (clockwise turn) of the initial heading. This controller is a closed-loop controller. If the robot gets nudged or turned, it tries to keep the previously set heading constant. This is a closed-loop controller since the resulting servo speed depends on the readings of the robot's inertial measurement unit (IMU).

It is also possible to directly set the speed of the left and right servo. This happens in an open-loop fashion, always setting the speed without any further observations.

As for sensors, the robot offers raw sensor information from a 3-axis accelerometer, a 3-axis gyrometer and the actual motor speed of each servo measured as voltage of the *back electromotive* force (back EMF). The robot can stream data from an IMU represented in quaternions or Euler angles. Additionally, it offers velocity information along a plane in the x and y direction, and also positional data (i.e., odometry) estimated from its starting position. For our studies, we use sensor data from the IMU, the accelerometer, the gyrometer and the speed of the wheels (see subsection 4.6).

We built a custom API to control the robot. The API is based on C++ and can run on embedded hardware. It communicates with the robot using BLE to send commands and to read sensor data (see Scheunemann, 2018).

#### 4.2 Measures

To investigate whether human participants perceive an intrinsically motivated robot as a social other, we employed measures from social cognition. There is a long history in the field of social cognition to understand human impression formation of other peers based on two central dimensions (e.g., Rosenberg et al., 1968; Wojciszke et al., 1998). Wojciszke et al. (1998) found that two central dimensions explain on average 82.25% of the variance of impression formation. A popular model for human impression formation is the stereotype content model (SCM) introduced by Fiske et al. (2007). According to Fiske et al. (2007), people perceived as warm and competent elicit uniformly positive emotions, are in general more favored, and experience more positive interaction by their peers. The opposite is true for people scoring low on these dimensions, meaning they experience more negative interactions. Warmth and Competence, together, almost entirely account for how people perceive and characterize others (Fiske et al., 2007).

Highly simplified, perceived Warmth leads to positive social bias, referred to as active facil-

itation (Cuddy et al., 2007). The Competence dimension mostly moderates this effect. High Warmth and high Competence result in admiration, while high Warmth and low Competence result in pity (Judd et al., 2005). The corresponding effects for low Warmth are envy and contempt. As a result, Warmth can be considered the primary factor for predicting the valence of interpersonal judgments (Fiske et al., 2007; Abele et al., 2016). This means, it primarily predicts whether an impression is positive or negative.

Recent research in HRI has shown that human participants prefer to interact again with the robot when they perceive its behavior as *more* warm compared to another robot behavior (Oliveira et al., 2019; Scheunemann et al., 2020). Our expectation, which would need to be confirmed in future work, is that the behavior of a robot that is perceived as more warm is a good candidate to facilitate long-term interaction.

In our present study we use the Robotic Social Attributes Scale (RoSAS) designed by Carpinella et al. (2017), which tests for the two aforementioned central dimensions *Warmth* and *Competence*. We accompany the scale with the Godpseed scale (Bartneck et al., 2009) because (i) it is a popular scale which may allow to compare the results to other studies and (ii) it encompasses the dimension Perceived Intelligence, which is highly related to the Competence dimension (Carpinella et al., 2017). The latter is important because we designed our study in such a way that participants should perceive one robot as similarly competent as the other robot since the perception of Competence can influence the perception of Warmth (Fiske et al., 2007). However, finding evidence for that, i.e., for a small or no effect on the Competence dimension, can be prone to errors. We therefore investigate whether two related dimensions, Perceived Intelligence and Competence, yield similarly small effects. In case they do support each other, it is unlikely that we missed an effect.

The Godspeed scale uses a 5-point semantic differential scale and investigates the dimensions Anthropomorphism, Animacy, Likeability, Perceived Intelligence and Perceived Safety. The RoSAS tests for the dimensions Warmth, Competence and Discomfort. Carpinella et al. (2017) do not recommend a specific size for the Likert-questions, but recommend including a neutral value, e.g., by having an uneven number of possible responses. Our questionnaire consists of 7-point Likert-type items.

This article focuses on the evaluation and discussion of the dimensions Warmth, Competence and Perceived Intelligence to answer the research questions (see subsection 4.5). However, our questionnaires encompass all the dimensions offered by the RoSAS and the Godspeed scale. The idea is (i) to hide the intent of our questionnaire and (ii) to allow future studies to compare their results with our study. We report the effects of all dimensions. However, we only discuss our target dimensions and any other statistically significant effect which was revealed. This allows us to discuss possible implications of unexpected effects (warranting the formulation of hypotheses and conducting a new study).

Note that we continue to capitalize the dimensions to indicate that we refer to, e.g., the questionnaire dimension Competence, as opposed to true competence. At times, however, we use the adjectives, e.g. competent, where it is clear that we refer to the dimension.

#### 4.3 Robotologist study design

To evaluate the effect of intrinsically motivated behavior generation on human perception of the robot we had to develop a new study design that could fulfill the following criteria:

- 1. Encourage the human to pay attention to the robot.
- 2. Encourage the human participant to interact ideally physically with the robot.
- 3. Do not provide an explicit task assignment for the robot to the human.
- 4. Do not provide a joint or human task assignment that leads to an implicit assumption about the robot's task.

The aim of the study design is to engage the human participant and to direct their attention towards the robot, so that participants can judge the robot behavior based on their interaction with the robot. We also want physical interaction, so we can test that the effect we are looking at is robust in regards to being perturbed by human interaction. This is because it would be counterproductive to identify some robot behavior that could encourage a human to physically interact with the robot and then have this behavior destroyed by the resulting interaction. In general, there are a lot of task-based study designs that can fulfill criteria 1 and 2, but the challenge was to find a design that can fulfill criteria 3 and 4 at the same time. Criteria 3 and 4 are necessary because our research questions focus on how warm the participants' perceive the robot, something that can be biased by the perception of the robot's Competence. Perceived competence in turn tends to be influenced by how well a robot fulfills a stated task – or even just an implicitly assumed task.

For example, in the preliminary study (Scheunemann et al., 2019), we had asked the participants to use their hands to keep the robot from falling off the table, to encourage physical interaction. However, this likely led to an implicit assumption that the robot should stay on the table, and resulted in those robots which were more likely to fall off the table being considered less competent. In this study we wanted to minimize the explicit and implicit task assignment to evaluate the effect of IM behavior generation on Warmth in the least biased way.

In order to achieve the aforementioned requirements, we designed our study so the participants effectively would learn to become *robotologists* in analogy to how anthropologists or naturalists would study animals or humans. The participants' task was to determine if the two presented robot behaviors behave the same or different. We hypothesized that this encourages interaction with the robot because participants usually want to perform well in order to satisfy the perceived needs of the researcher (Orne, 1962). At the same time, however, this task does not influence the participants' expectations of the robot directly and it does not create the expectation that the robot will help or hinder this assessment. We hypothesized that this robotologist study



Figure 2: The picture shows the first author using the interaction tool. He nudged the robot with the white end of the wand. Participants were able to freely chose a position around the table for observing or interacting with the robot.

design encourages the human to interact with the robot, while at the same time reducing their preconceptions.

## 4.4 Environment and Interaction

Figure 2 shows the table that the robot locomoted on. It is circular, with 91 cm in diameter and 72 cm in height. A foam wall of 2.5 cm in height and with 4 cm in width surrounds the border of the table. We decided on these measurements in such a way that the robot cannot fall off the table, even with a very high velocity. Three blankets of a total height of 3 to 4 mm cover the surface (including the walls). This applied some friction and made it easier for the robot to locomote on the otherwise smooth and slippery surface of the wooden table top. The table's distance to the surrounding wall of the room was at least 60 cm, which allowed participants to freely move around the table.

This study design is different from our preliminary study (Scheunemann et al., 2019). There, we positioned the participant at one particular side of the table which had no border and therefore permitted the robot to fall off the table. We hypothesized that this encourages interaction with the robot and that participants would see that the intrinsically motivated robot adapts to their input. Furthermore, we hypothesized that an intrinsically motivated robot, which explores the area and the interaction with the participant would be perceived as having higher agency and competence. This design was not entirely successful. The intrinsically motivated robot could not sense the edge. Therefore, when it approached participants they did not consider it as

approaching them, but judged the robot as rather *suicidal*. In consequence, the participants judged the robot as less competent, despite its ability to adapt (cf. Scheunemann et al., 2019, sec. 5). This is a good example of why it is critical to consider the participants' expectations when choosing a robot and designing a study.

In this study, we decided to keep the robot hardware, but change the design of the environment, so that all borders of the interaction environment are enclosed. If the participant decides to be passive, the robot cannot fall off the table. In addition, the round shape of the table and its position allows participants to reach all borders. We do not assign a specific position to the participants. This further reduces our instructions to the participants, aiming to further reduce our influence on their implicit task assignment for the robot. This design of the interaction environment is key to enable the robotologist study design and supports the requirements outlined above in subsection 4.3.

Figure 2 shows the first author of this article interacting with the robot using the HRI tool referred to as a *wand* that was developed specifically for this study. Participants were asked to use the wand to touch and nudge the robot with the white end. The wand is 50 cm long and weighs 78 g. It consists of a 40 cm long aluminum tube with a diameter of 10 mm. The end is a round, softer sphere. It is made of an off-the-shelf table tennis ball with a diameter of 40 mm.

The style of interaction is another difference to the preliminary study (Scheunemann et al., 2019) where participants used their hands to interact with the robot. The participants in this study were asked to use the wand for interacting. We assumed that this would help to ease the interaction, as some participants in the preliminary study felt uncomfortable with the idea of using their hands for means of interactions. We further hypothesized that the mere existence of a tool would make the participants want to use it and therefore encourages the interaction.

#### 4.5 Research questions

The study concerned itself with two main research questions. First, we wanted to understand if participants perceive an intrinsically motivated robot as a social other and are interested in interacting with such a robot. As discussed in subsection 4.2, we employed the dimension Warmth to measure participants' perception. We are interested in a *change* of the perception because there is not enough research to understand what absolute value of Warmth would be a *good* indicator. In this study, we first want to understand whether IM-driven behavior has a positive effect:

**R1:** Is an intrinsically motivated, adaptive robot behavior perceived as more warm compared to a robot with a reactive baseline behavior?

We hypothesized that we would find evidence that this is the case. We did not expect any other strong effects, but we still report and discuss the statistically significant main effects of all other dimensions. In the preliminary study (Scheunemann et al., 2019), we saw evidence that participants perceive an intrinsically motivated robot as more warm. This, however, needed confirmation because (i) the effect was not statistically significant and because (ii) the study design made participants perceive one of the robots as more competent than the other. This can act as a confounding factor because perceived Competence can influence perceived Warmth. In this study, we aimed to produce a fairer set-up to allow an evaluation of Warmth that is not biased by perceived differences of the robot's competence. The above subsections presented what we call the participants' task to become a *robotologist* and the environment and the interaction we employed in order to achieve that goal.

Drawing a conclusion from *not* finding an effect for Competence is prone to the risk that we simply might have missed an effect. To be certain, we therefore used the dimension Perceived Intelligence, a dimension related to Competence (see subsection 4.2), and expected both of them to show no effect. Our second question is therefore:

**R2:** Does the robotologist study design help to make the two robot behaviors appear to the human participants as similarly competent and intelligent?

We hypothesized that the answer to the question is "yes". This means that for the dimensions Competence and Perceived Intelligence we would expect to see no evidence for an effect.

#### 4.6 Conditions

This experiment consists of two conditions of behavior generation with the following characteristics:

- ADA: The robot is adapting continuously, based on maximization of TiPI and directly applies the resulting IMs as servo speed.
- *REA*: The robot uses its balanced mode for locomotion, the network controlling the robot has been pre-adapted using PI and it remains constant.

The reactive robot in the *REA* condition uses the same binary and starts with the same networks as the robots in the preliminary study (cf. Scheunemann et al., 2019, sec. 3.1). The weights are received based on pre-trial adaptation. This determines how it reacts to sensor input, but it does not further update its internal network during the experiment. There are two reasons for taking the *REA* robot from our preliminary study (Scheunemann et al., 2019). Firstly, the behavior is a good baseline behavior. The robot was interesting to the participants and the behavior showed enough variety for them to not see any patterns (Scheunemann et al., 2019). Secondly, keeping the baseline constant, but changing other variables, allows for a better comparison to the previous findings and the previous intrinsically motivated, adaptive robot.

The intrinsically motivated robot in the ADA condition realizes behavior motivated by TiPI maximization, and it continuously updates its internal networks based on that gradient during

the experiment. In contrast to the preliminary study, the robot changed the speed of its two servos directly, instead of using the balancing controller (cf. Scheunemann et al., 2019, sec. 3.3). This way the robot's behavior is only influenced by its IMs, unconstrained by additional software such as the closed-loop balancing controller. In particular, this balancing controller might have added meaning to the robots output (i.e., staying upright), which may not yield an intrinsically motivated robot in accordance to the definition presented in subsection 2.1. Therefore, this change allowed to further focus the analysis on the perception of intrinsically motivated autonomy.

The robot sensor input is again the linear acceleration for the forward/backward and left/right axis from the accelerometer, and the angular velocity around the upright axis received by the gyrometer. Instead of using the absolute position of the robot received via its pitch and roll angles from the IMU as in the preliminary study (Scheunemann et al., 2019), we now input the speed of the two servos. This allows us to directly couple the output of the controller changing the servo speed and the actual measured servo speed.

We wanted both robots to behave similarly at the beginning of the two conditions, to avoid the formation of very different first impressions. For example, one robot starting off smooth and slow, and the other accelerating very fast and bumping into the wall may form a first impression in the participant which influences their overall responses. Therefore, we tweaked the starting weights of the network of the intrinsically motivated robot by hand. As there is direct coupling between the servo speed readings and the controller output, i.e., the set speed for the servos, the weights were set in such a way that a reading on the left servo would amplify the output for the left servo, and vice versa. This way we could create a slow-pace forward movement for the first few seconds, which looks similar to the reactive baseline robot. An example video of the two conditions of one session accompanies the article (see Scheunemann, 2021b). As seen in the video, the resulting two robot behaviors look overall very similar and are hard to distinguish visually. We believe this is a strength of the experimental set-up, as it allows us to rule out other, incidental reasons for the observed change in human perception. This makes it more likely that it is the adaptation to stimuli received by the robot, driven by TiPI maximization, that is responsible for the differences in human perception.

#### 4.7 Procedure

Participants are welcomed to the experimental room and were then handed an information sheet. They were encouraged to discuss concerns related to their participation. If they were happy to proceed with the study, they were asked to sign an informed consent form. Then the environment and the robot are presented and briefly described. It was then emphasized that they could leave the study whenever they feel uncomfortable, stressed or bored. Participants then complete a pre-questionnaire. This gathers information regarding their gender, age and background.

We then formulated the task for the participants, namely that they should find out whether the two presented robots are any different. For understanding differences in the robots' behaviors, they can use the HRI tool: the wand. They are allowed to nudge the robot or block it. Both of these actions are demonstrated to the participants. However, no other information is provided.

Next, the two conditions are presented to the participants in a randomized but counterbalanced order, each lasting approximately 5 min. They complete a post-questionnaire containing the two scales after each condition. The entire experiment takes 50 to 60 min per participant.

## 4.8 Participants

We recruited 24 participants (10 female; 14 male) mostly from university staff and students, between the ages of 18 and 64 years (M = 31.7, SD = 12.6). The participants were undergraduates or post-graduates from the university, but all naïve towards the objectives of the experiment. Eight participants had a background in HRI, whereas 9 participants never participated in any prior HRI study. All were asked how familiar they are with interacting with robots, programming robots and the chosen robot platform. 5-point Likert-questions were chosen with the value 1 for "not familiar" to 5 for "very familiar". The self-assessed experience for interacting with robots showed an average of 3.5 (Mode = 5). The average familiarity with programming robots 3.2 (Mode = 5) and the experience with the chosen robot platform was rated an average 2.1 (Mode = 1). The familiarity with the movie series Star Wars was rated an average of 3.2 (Mode = 4).

These data were collected primarily to understand what type of participants were attracted to the study. We did not form any hypotheses, but we expected that many participants may be aware of the Star Wars movies or the characters and their impression of the robot's behavior may vary depending on their expectations. However, in this study we are interested whether participants perceive one robot behavior as more warm than the other to understand their perception of intrinsically motivated autonomy (cf. subsection 4.5). This is why we decided for a within-subject design and presented both conditions to each participant. This way, we can investigate the participants' change in their responses, rather than comparing responses of participants with different expectations of the robot's behavior and capabilities. This allows us to concentrate on the changes of their responses without considering their expectations of the robot's behavior.

The study was conducted on the premises of the University of Hertfordshire and was ethically approved by the Health, Science, Engineering & Technology ECDA with protocol number aCOM/PGR/UH/03018(3). The anonymity and confidentiality of participants' data are guaranteed.

#### 4.9 Data analysis

Our research questions asked whether the responses to a questionnaire dimension of a participant differ between conditions. In particular, we want to know whether participants perceived one robot behavior as more warm than the other (RQ1). Since we are interested in the change of participants' perception, we designed the study so that the robot behavior generation is a within-subjects variable, i.e., all its conditions are presented to each participant. This has the benefit that we can answer the research questions with the use of statistical tests for pairwise comparisons of the participants' responses.

Since our study consists of exactly two conditions, we decided to employ two-sample location test of matched pairs: the Wilcoxon signed-rank test. In contrast to the paired *t*-Test, it is a non-parametric test which is both: (i) robust for small sample sizes and (ii) usable for the assessed Likert-scale data. We decided to use the two-sided version of the test to investigate for effects in both directions. This way, effects contrary to our hypotheses are revealed too. We use wilcox.test(*ADA*, *REA*, paired=TRUE) to compute the test, a method that is part of R's built-in stats package.

#### 4.10 Data preparation

In order to use the Wilcoxon signed-rank test we first had to analyze the data from participants' responses in two steps: (i) we analyzed the scale reliability in order to amend scale items if needed and (ii) we tested whether the condition responses are independent of their presented order.

#### 4.10.1 Scale reliability

To prepare the data for analysis, we tested the score reliability of the scales of both standardized questionnaires using Cronbach's  $\alpha$ . Table 1 presents all test results for the used dimensions. We found that the item *quiescent-surprised* of the dimension Perceived Safety of the Godspeed scale is negatively loaded on the dimension. Even if reversed, the reliability is poor with  $\alpha = .54$ . We decided to remove Perceived Safety altogether. All other dimensions reveal a good score reliability ranging from .75 to .85, or acceptable reliability for the dimension Anthropomorphism:  $\alpha = .67$ . This is evidence that the scale dimensions are reliable and can be used for further investigation.

Table 1	1:	Internal	consistency	reliability	measured	with	Cronbach'	s	α.

	dimension	items	$\alpha$
SAS	Warmth	6	.80
	Competence	6	.85
$\mathbf{R}_{\mathbf{O}}$	Discomfort	6	.79
	Anthropomorphism	5	.67
eec	Animacy	6	.74
$\frac{1}{2}$	Likeability	5	.82
005	Perceived Intelligence	5	.84
0	Perceived Safety	3	.37

#### 4.10.2 Interaction effects

We then analyzed whether the randomized and counterbalanced assignment of the order of conditions to the participants was successful in that the data does not show interaction effects between the condition responses and the order of the condition.

	dimension	F	df1	p
SAS	Warmth	0.001	1	.976
	Competence	1.473	1	.225
Ч	Discomfort	1.787	1	.181
Godspeed	Anthropomorphism	0.164	1	.685
	Animacy	0.665	1	.415
	Likeability	1.455	1	.228
	Perceived Intelligence	< 0.001	1	.984

Table 2: ANOVA-type test results for interaction effects.

An analysis of variances (ANOVA) is commonly used for investigating interaction effects. Following our above argumentation for employing non-parametric tests, we use a non-parametric ANOVA-type test (Noguchi et al., 2012). For computing the ANOVA-type test we used the R package nparLD. The study can be expressed as F1-LD-F1 Model with the one within-subjects variable behavior generation (levels: *REA*, *ADA*) and the one between-subjects variable order (two levels). The nparLD package offers the function f1.ld.f1() for computing such models. Table 2 shows the results. For a 5% significance level there is no statistical significance and there is not enough evidence for an interaction effect for any of the dimensions.

The randomized but counterbalanced ordering of conditions resulted in no evidence for interaction effects. Therefore, we can safely investigate the main effects independently of their order. This means we can use the Wilcoxon signed-rank test and compare the responses to both conditions independently of whether the participants were exposed to, e.g., ADA, at the beginning of the experiment or at the end.

## 5 Results

Our two research questions asked whether the participants' responses to the robot behavior differ between conditions. We therefore employed a paired difference test: the Wilcoxon signed-rank test. It is a non-parametric test which is known to be robust for small sample sizes and can be used for the assessed Likert-scale data.

Table 3 shows the results of the two-sided Wilcoxon signed-rank test for all dimensions of both scales comparing the condition REA and ADA. For each dimension, we report the test statistic V, the p value, a point estimate and its corresponding confidence interval. The point estimate (short: estimate) is the median of the difference between REA and ADA. It provides

			95% confide				
	dimension	V	lower bound	upper bound	estimate	p	r
RoSAS	Warmth	27.5	-1.333	-0.333	-0.833	*.007	.555
	Competence	138.5	-0.833	0.583	0	.988	.003
	Discomfort	54.0	-1.250	0.417	-0.250	.287	.217
Godspeed	Anthropomorphism	26.0	-1.200	-0.500	-0.900	$^{*}.002$	.635
	Animacy	30.5	-1.250	-0.417	-0.833	*.002	.636
	Likeability	38.0	-0.700	-0.100	-0.400	$^*.038$	.424
	Perceived Intelligence	126.0	-0.600	0.500	-0.100	.715	.075

Table 3: Results of the two-sided Wilcoxon signed-rank test for comparing the difference between REA and ADA.

a magnitude and a direction for how much the participants prefer one condition. For example, if the estimate for REA - ADA equals -0.833, this means that on average the participants responded to Warmth with 0.833 units higher in the ADA condition compared to the REA condition. In other words, participants perceive the intrinsically motivated robot as more warm than the baseline behavior. The units here are the responses to the Likert-type scale ranging from 1 to 7 (RoSAS) or the differential scale ranging from 1 to 5 (Godspeed). Along with the estimate, we further report the standardized effect size r. It allows investigating the size of a potential effect independently of the sample size. We follow Yatani (2016) who suggested to use Cohen's interpretation: small effect for r = .1, a medium effect for r = .3 and a large effect for r = .5.

The results show that there are statistically significant effects for the dimensions Warmth (p = .007), Anthropomorphism (p = .002), Animacy (p = .002), and Likeability (p = .038). The magnitude of the effects (estimates) for these effects are negative, which means that participants respond higher on each dimension if the robot is intrinsically motivated. We only hypothesized one effect, namely that most participants perceive the intrinsically motivated robot (ADA) as more warm than the robot in the reactive baseline condition (REA). This directly answers our first research question (cf. subsection 4.5, RQ1), namely that an intrinsically motivated robot (as the one in the ADA condition) is perceived as more warm. We will discuss the other unexpected findings in section 6.

Figure 3 visualizes the magnitude of the effect. The magnitude of the effect increases with an increasing distance of the estimate to zero. The figure also visualizes the certainty that the point estimate is indeed the true effect. The smaller the error bars, i.e., the confidence interval, the more certain we can be about the point estimate. Figure 3 confirms that there is a large effect for Warmth in favor of the ADA condition.

An even larger effect can be observed for the two dimensions Anthropomorphism and Animacy. For both dimensions, their perception differs and is statistically significant. The estimate and r > .5 again indicates that there is a large effect in favor of the *ADA* condition.



Figure 3: The figure visualizes the results of the two-sided Wilcoxon signed-rank test reported in Table 3. The median of the difference (magnitude of the effect) is plotted as point estimates and the 95% confidence interval as error bars. The figure shows that the intrinsically motivated robot (ADA) is perceived more warm compared to the baseline robot (REA). The figure confirms that there is no magnitude of an effect for either of the two dimensions Perceived Intelligence and Competence.

The p value for the dimension Likeability is statistically significant (p = .038) and the effect size is r = .424 (medium). The magnitude of the effect is the smallest among all statistically significant effects. Interestingly, it shows the narrowest confidence interval, indicating that the magnitude of the effect, while small, is most certain.

As hypothesized, there is no evidence that the robot is perceived differently for either of the two dimensions Perceived Intelligence and Competence. From Table 3 we see that the standardized effect size r for Competence (r < .003) and Perceived Intelligence (r < .075) indicate that there is no effect. Moreover, Figure 3 shows that the magnitude of an effect is almost zero for both of the dimensions. The confidence interval is almost equally distributed around zero, indicating that there is no certainty for an effect in any direction. Both dimensions Competence and Perceived Intelligence, which measure a similar concept, show a very similar result for testing the perceived difference, which provides further support that there is no or only a very small effect. This helps to answer our second research question (cf. subsection 4.5, RQ2), namely whether our robotologist study design helps to make the two robot behaviors appear similarly competent.

# 6 Discussion

The study results provide evidence that the intrinsically motivated robot behavior is perceived as more warm than the behavior of the reactive baseline robot. The dimension Warmth is one of the universal dimensions for humans judging social attributes of other humans (e.g., Judd et al., 2005; Fiske et al., 2007; Abele et al., 2016). Notably, a high scoring for Warmth is considered positive, i.e., desirable, and leads to more positive interactions with peers. There is first evidence that this also applies to robots perceived as more warm too. Recent research on robot behavior in HRI has shown that human participants prefer to continue interacting with the robot they perceive as most warm (Oliveira et al., 2019; Scheunemann et al., 2020). Our results show that the intrinsically motivated robot displays behavior that caused participants to rate the robot higher in Warmth compared to the behavior of the reactive baseline robot. This is an indicator that intrinsically motivated autonomy in robots may prove to be relevant for human-robot interaction. Further support is provided by the statistically significant effect for the dimension Likeability, which indicates that participants do like the intrinsically motivated robot more.

The current robotologist study design helped us to focus on the Warmth dimension. Participants did not perceive the robot differently for either of the two dimensions Perceived Intelligence and Competence in any of the conditions. This reflects that the participants did not assign a similar (or any) implicit goal to the robot. Although Competence and Warmth are mainly considered unique dimensions, some interference and correlations have been pointed out between them (Fiske et al., 2007; Abele et al., 2016). The lack of an effect for both Perceived Intelligence and Competence is therefore an important feature of our study design, which allows for an isolated observation of the influences of IM on the perception of Warmth.

We also – unexpectedly – observed that participants perceived the intrinsically motivated robot as more animated (and as more anthropomorphic). There is evidence that humans perceive a robot higher in Animacy when the robot moves more "naturally" (Castro-González et al., 2016). In fact, any object is considered animated if it changes speed and direction without visible influences (Tremoulet & Feldman, 2000). Another influence of the perception of Animacy is the reactivity of the robot (Fukuda & Ueda, 2010). We designed our baseline behavior to provide both: similar movement variety, and reaction to sensor input, to allow for a fair comparison and focus on the effects of the intrinsic motivation. In the current study, the control mechanism for the IM robot was changed (cf. subsection 4.6). Other than the reactive baseline behavior, where the robot could only move forward and was kept mostly upright due to the balancing controller, the IM robot had a different behavioral regime. It could go backward and forward, and because the servo speed was set directly and individually, it could produce different behavioral regimes like, e.g., a wobbling locomotion. Therefore there are three possible explanations for the baseline behavior being perceived as less animated: (i) the motion patterns, (ii) the reactivity, or (iii) the intrinsic motivation. With the current data, we cannot answer this question sufficiently. Although the baseline behavior was shown to be feasible in the preliminary study (Scheunemann et al., 2019), we will therefore investigate possible changes in a follow-up study, where the baseline will be designed to be perceived as similarly animated as the intrinsically motivated robot here.

It should be noted that the robot with the baseline behavior is not perceived as inanimate in an absolute sense. Instead, participants simply perceived the intrinsically motivated robot as more animated and the baseline as less animated. Although this is an indication for the baseline behavior to have less natural movements (as discussed), there is no evidence in the literature that the rating for Warmth has been significantly influenced. In the preliminary study (Scheunemann et al., 2019), for example, participants perceived the baseline behavior as more animated (small effect), but they perceived the intrinsically motivated robot as more warm (medium effect). Given the results of both studies, we argue that there is evidence that the different participant responses for Warmth between the two behavior conditions were mainly caused by the robot's intrinsic motivation.

What remains unclear is how much the knowledge from human social cognition transfers to human-robot interaction. For example, much like in human social cognition, will a robot perceived as warm also experience more positive social interaction? Despite recent advances (e.g., Oliveira et al., 2019; Scheunemann et al., 2020), future work has to understand whether the concepts from human social cognition transfer to *physical* interaction with a robot. If that is the case, our study shows that a robot which has intrinsic motivation can help to increase the interest of a human to interact with it, and that an intrinsically motivated robot is likely to receive more positive social interactions.

Our robotologist study design can play a vital role in answering the above question, and it can support future research to deepen the understanding of the effect of intrinsically motivated robot behavior on the human interaction partner. The study design can be easily amended by (i) using different baseline behaviors, (ii) investigating more complex robot behaviors enabled by IM, or (iii) using alternative engagement measures. We already discussed how we plan to compare the intrinsically motivated robot with a baseline behavior that is more animated. Other baseline behaviors could also include behaviors that are similarly adaptive, but not intrinsically motivated. Once the isolated effect of IM on human perception and its causal mechanisms have become more clear, the design can be amended using more complex interactions. This could include robot behavior that is not only exploratory but also reuses learned behavior. We did not investigate such behavior in this study. On the one hand, it is currently an open question of how to align IM and learned behavior (more on that later), on the other hand, this study serves as first evidence of the specific effect of IM on human participants. Depending on future advances for (i) and (ii), the study design also allows using alternative engagement measures for further investigation. The current study was chosen to use questionnaires using established measures from social cognition as a first step to understand the effect of IM. However, with increasing behavioral complexity the study would likely need to be amended with additional measures for the participant's perception. A popular measure for engagement is the time a participant chooses to be involved in an interaction with another agent (e.g., Bickmore et al., 2010). A future study could employ the time of interaction as an engagement measure and compare the length of the human-robot interaction between conditions. With growing interaction complexity, this may need further amendments of other engagement measures that investigate distinct interaction parts, such as communication or gaze (e.g., Rich et al., 2010). Note that we purposefully did not investigate more complex scenarios in this study to focus on the effect IM adaptions produce in an even minimal system. This simplicity in our study design, in addition, was chosen to limit

the possibility of alternative explanations for the observed effect other than the specific effect between the two conditions.

Coming back to the future plans to investigate more complex human-robot interaction scenarios. We already know from recent research that more complex, anthropomorphic robot platforms that *mimic* being intrinsically motivated (e.g., curious), are more engaging for humans (e.g., Gordon et al., 2015; Ceha et al., 2019). To close the gap between the robots mimicking IM, i.e., with behavior implemented and designed by humans, and intrinsically motivated robots per se, several tasks need to be addressed first. Two main concerns here are scalability and alignment. We know from related work that TiPI, and other intrinsic motivations, can be scaled up to more complex behaviors. For example, Martius et al. (2013b) presented TiPI-controlled simulated humanoids that generate different patterns of behavior in different environments. This might not be seen as a complex scenario in the field of HRI with its focus on reliable task completion. This raises the question of how to align this kind of behavior generation with behavior a designer would want, such as a robot assisting a human? Both issues can be addressed by combining IM-based behavior generation with scripted behavior, or behavior based on extrinsically rewarded reinforcement learning (see Singh et al., 2004; Jaques et al., 2019, for broadly related initial research). This can provide a complex scaffolding of behaviors, which then gets further enhanced towards more perceived Warmth, by using IMs, such as TiPI, for behavior modification and local behavior generation. In this article, we deliberately left these things out, to focus on an empirical study of the effect of intrinsic motivation on its own. In other words, we investigate the effect of IM in isolation from other factors.

One challenge that this approach poses from the perspective of traditional HRI research, is the question of how to identify the isolated factors, such as salient behavioral patterns, that cause the robot's behavior to be perceived as more warm. One motivation for this could be, for instance, that the two resulting robot behaviors are hard to distinguish visually. This contrasts with more established approaches, where a specific behavior is the tested condition, thus providing a more proximal explanation of how certain human perceptions were elicited (see subsection 4.6). We note that there is research in HRI that does exactly that: it scripts robots to display behavior associated with various intrinsic motivations, and measures their effect on humans (e.g., Gordon et al., 2015; Ceha et al., 2019). By design, such scripts constitute isolated factors. Our study does not do this, instead, it investigates the perception of a robot where an intrinsic motivation model actually generates and adapts the behavior. The robot is not just going through the motions of behaving as if it is curious, it actually seeks out new and predictable stimuli for its sensors. Nevertheless, we do still establish two clearly distinct conditions, defined by our controller and the mathematical properties they operate under, on which we performed an interventional, double-blind test with significantly different results. It would be, of course, of great interest to us to better understand the resulting, salient behavior differences that provide a proximal, mechanical explanation for these different perceptions. This would help to close the gap between IM generation and more scripted or task-focused ap-

proaches. However, we emphasize that using classic HRI techniques for this study to understand isolated behavior patterns is a task far from trivial. This is because of the core property of the information-theoretic measure for IM: universality. Robots that are truly intrinsically motivated can cope with a broad range of changes to an agent's environment or its morphology (see subsection 2.1, 2.3, and 3.3). We argued that this is key to enable robots to interact with humans despite complex environments and unknown human characteristics. However, it is this precise property that makes it inherently hard to investigate for isolated behavioral factors. Firstly, our exclusive focus on the robot's intrinsic motivation comes with the limitation that we cannot, inherently, impose specific behaviors or motivations as isolated factors. Doing so would make the robot's behavior externally motivated, which is something we did not set out to investigate in the present study. Secondly, information decomposition of the TiPI maximization that is used to generate IM behavior could ultimately help to identify and isolate factors of the behavior generation. However, the information-theoretic work on the question of information decomposition, which has been under intense investigation in the last decade, makes it clear that the question on how to systematically disentangle different contributions to a truly complex behavior is far from straightforward to answer at this stage (e.g., Ay et al., 2020; Rosas et al., 2020). As we can neither control the robot behavior externally nor decompose the information that yields the robot's behavior, we would need to research methods that extract isolated factors. We know that observing robot behavior and being part of an interaction with the same robot behavior can result in different perceptions (e.g., Fukuda & Ueda, 2010). Therefore, the employed methods should be more than just extracting isolated factors from the robot behavior, but rather be a careful treatment of the influence of the robot's behavior on the human and vice versa. This means a detailed understanding of the interaction loop between both the robot and the human is needed. Suitable measures to address this satisfactorily do not yet exist. In fact, this is why most HRI research uses scripts or teleoperators that mimic certain aspects (such as IM). These scripts or the actions of a teleoperator do already constitute isolated factors. We argue that our approach is an important addition to this technique as it allows the investigation of fully autonomously generated behavior that cannot currently be decomposed; something that, to the best of our knowledge, has not yet been researched in HRI.

Despite those long-term challenges, our findings offer some direct applications to more current robots. The here presented TiPI formalism can be used to implement a generic motion for situations when a robot does not exhibit a specific behavior, i.e., no human is interacting with the robot. We provided evidence that this may attract more humans due to their perception of the robot behavior being friendly (Warmth). This would reduce the number of times researchers need to hand-tweak natural or affective behavior. To make a robot more engaging and to elicit curiosity in the human interaction partners, we discussed that some researchers proposed that novel behaviors, or a larger variety, are important. These designed behaviors (e.g., uttering questions or statements) are often randomly chosen in autonomous robots (e.g., Gordon et al., 2015). We propose that a more naturalistic selection could be applied by using an intrinsic motivation measure. Using TiPI directly is not the best candidate, as designed behaviors or scripts cannot be represented by a continuous variable. However, TiPI could be used as a reward signal for a selection algorithm based on, e.g., reinforcement learning. Alternatively, researchers could decide for another formalism implementing IM, such as empowerment (Klyubin et al., 2005). We argue that the study design presented here can help to prototype an affective behavior, or affective behavior selection, for a variety of IM formalisms.

# 7 Conclusion

We started this research with the question if intrinsically motivated autonomous robots can be beneficial for designing engaging human-robot interaction (HRI). We conducted a within-subjects study (N = 24) where participants interacted with a fully autonomous Sphero BB8 robot with two conditions covering different behavioral regimes: one realizing an adaptive, intrinsically motivated behavior and the other being reactive, but not adaptive. We used time-local predictive information (TiPI) maximization as one candidate measure to produce IM-based behavior, and produced, to our knowledge, the first study quantitatively investigating human perception of intrinsically motivated robots. Of particular interest is the high similarity between both conditions in Perceived Intelligence (r = .075, p = .715) and Competence (r = .003, p = .988), which gives support to our non-task oriented interaction design. This was particularly important as Competence ratings can influence the perception of Warmth, which is the dimension we focused on in our study.

Our main result is that an intrinsically motivated robot generates behavior that was perceived more warm compared to a baseline robot that was not intrinsically motivated (r = .555, p = .007). The baseline behavior includes both: similar movement and reaction to sensor inputs – meaning that it is highly likely that the difference in perception arises from the robot's adaptation to the physical interaction. The dimension Warmth is, as mentioned previously, an important factor for attitude formation in human-human social cognition. However, it is not immediately clear if this higher perceived Warmth leads to a positive attitude or preferences in human-*robot* interaction. If future work would demonstrate this, then we believe the formalism presented here could be utilized to create a preference or positive attitude towards a robot in a large range of HRI scenarios.

The open questions going forward are now: Can we confirm the results by using a baseline behavior with more similar motion regimes to further strengthen the focus on the intrinsic motivation (IM) of the agent and the interaction? Does the universal applicability of the formalism also translate into a universal, or at least widespread, evocation of Warmth across different robot morphologies? Does this effect persist over time? And does a positive social attitude lead to more engagement? All these questions are empirically testable, and given the positive results here are possible directions for future research.

## Acknowledgements

We would like to thank Rebecca Miko for her help with recruiting participants and with her help in proof-reading.

## Declaration of conflicting interests

The authors declare that there is no conflict of interest.

## Funding

MS and DP acknowledge support by the socSMCs FET Proactive project [grant number H2020-641321]; and KD acknowledges funding from the Canada 150 Research Chairs Program.

# References

- Abele, A. E., Hauke, N., Peters, K., Louvet, E., Szymkow, A., & Duan, Y. (2016). Facets of the Fundamental Content Dimensions: Agency with Competence and Assertiveness—Communion with Warmth and Morality. *Frontiers in Psychology*, 7. doi: 10.3389/fpsyg.2016.01810
- Ay, N., Bernigau, H., Der, R., & Prokopenko, M. (2012). Information-driven self-organization: the dynamical system approach to autonomous robot behavior. *Theory in Biosciences*, 131(3), 161–179. doi: 10.1007/s12064-011-0137-9
- Ay, N., Bertschinger, N., Der, R., Güttler, F., & Olbrich, E. (2008, 01). Predictive information and explorative behavior of autonomous robots. *The European Physical Journal B*, 63(3), 329–339. doi: 10.1140/epjb/e2008-00175-0
- Ay, N., Polani, D., & Virgo, N. (2020). Information Decomposition based on Cooperative Game Theory. *Kybernetika*, 56(5), 979–1014. doi: 10.14736/kyb-2020-5-0979
- Bartneck, C., Kulić, D., Croft, E., & Zoghbi, S. (2009). Measurement Instruments for the Anthropomorphism, Animacy, Likeability, Perceived Intelligence, and Perceived Safety of Robots. *International Journal of Social Robotics*, 1(1), 71–81. doi: 10.1007/s12369-008-0001-3
- Belpaeme, T., Kennedy, J., Ramachandran, A., Scassellati, B., & Tanaka, F. (2018). Social robots for education: A review. *Science Robotics*, 3(21). doi: 10.1126/scirobotics.aat5954
- Bialek, W., Nemenman, I., & Tishby, N. (2001). Predictability, Complexity, and Learning. Neural Computation, 13(11), 2409–2463. doi: 10.1162/089976601753195969
- Bickmore, T., Schulman, D., & Yin, L. (2010). Maintaining Engagement in Long-term Interventions with Relational Agents. Applied Artificial Intelligence, 24(6), 648–666. doi: 10.1080/08839514.2010.492259

- Boden, M. A. (2008). Autonomy: What is it? *Biosystems*, 91(2), 305–308. doi: 10.1016/j.biosystems.2007.07.003
- Carpinella, C. M., Wyman, A. B., Perez, M. A., & Stroessner, S. J. (2017). 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). New York, NY, USA: ACM. doi: 10.1145/2909824.3020208
- Castro-González, Á., Admoni, H., & Scassellati, B. (2016). Effects of form and motion on judgments of social robots' animacy, likability, trustworthiness and unpleasantness. *International Journal of Human-Computer Studies*, 90, 27–38. doi: 10.1016/j.ijhcs.2016.02.004
- Ceha, J., Chhibber, N., Goh, J., McDonald, C., Oudeyer, P.-Y., Kulić, D., & Law, E. (2019). Expression of Curiosity in Social Robots. In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems - CHI'19. ACM Press. doi: 10.1145/3290605.3300636
- Christensen, H. I., Okamura, A., Mataric, M., Kumar, V., Hager, G., & Choset, H. (2016). Next Generation Robotics (Tech. Rep.). Retrieved from https://arxiv.org/abs/1606.09205v1
- Clabaugh, C., & Matarić, M. (2019). Escaping Oz: Autonomy in Socially Assistive Robotics. Annual Review of Control, Robotics, and Autonomous Systems, 2(1), 33-61. doi: 10.1146/annurev-control-060117-104911
- Cover, T. M., & Thomas, J. A. (2012). Elements of Information Theory. John Wiley & Sons.
- Crutchfield, J. P., & Young, K. (1989). Inferring statistical complexity. *Physical Review Letters*, 63(2), 105–108. doi: 10.1103/physrevlett.63.105
- Cuddy, A. J. C., Fiske, S. T., & Glick, P. (2007). The BIAS map: Behaviors from intergroup affect and stereotypes. *Journal of Personality and Social Psychology*, 92(4), 631–648. doi: 10.1037/0022-3514.92.4.631
- Dautenhahn, K. (2004). Robots we like to live with?! a developmental perspective on a personalized, life-long robot companion. In RO-MAN 2004. 13th IEEE International Workshop on Robot and Human Interactive Communication (IEEE Catalog No.04TH8759). IEEE. doi: 10.1109/roman.2004.1374720
- Der, R., Güttler, F., & Ay, N. (2008). Predictive information and emergent cooperativity in a chain of mobile robots. In S. Bullock, J. Noble, R. Watson, & M. A. Bedau (Eds.), The 11th International Conference on the Synthesis and Simulation of Living Systems (Artificial Life XI) (pp. 166–172). MIT Press. Retrieved from https://mitpress-request.mit.edu/ sites/default/files/titles/alife/0262287196chap22.pdf
- Der, R., & Martius, G. (2006). From Motor Babbling to Purposive Actions: Emerging Selfexploration in a Dynamical Systems Approach to Early Robot Development. In S. Nolfi et al.

(Eds.), From Animals to Animats 9 (Vol. 4095, pp. 406–421). Springer Berlin Heidelberg. doi: 10.1007/11840541<sup>34</sup>

- Der, R., & Martius, G. (2012). The Playful Machine (Vol. 15). Springer-Verlag Berlin Heidelberg. doi: 10.1007/978-3-642-20253-7
- Fiske, S. T. (2018). Stereotype Content: Warmth and Competence Endure. Current Directions in Psychological Science, 27(2), 67–73. doi: 10.1177/0963721417738825
- Fiske, S. T., Cuddy, A. J. C., & Glick, P. (2007). Universal dimensions of social cognition: warmth and competence. *Trends in Cognitive Sciences*, 11(2), 77–83. doi: 10.1016/j.tics.2006.11.005
- Friston, K. (2010). The free-energy principle: a unified brain theory? Nature Reviews Neuroscience, 11, 127–138. doi: 10.1038/nrn2787
- Froese, T., & Ziemke, T. (2009). Enactive artificial intelligence: Investigating the systemic organization of life and mind. Artificial Intelligence, 173(3-4), 466–500. doi: 10.1016/j.artint.2008.12.001
- Fukuda, H., & Ueda, K. (2010). Interaction with a Moving Object Affects One's Perception of Its Animacy. International Journal of Social Robotics, 2(2), 187–193. doi: 10.1007/s12369-010-0045-z
- Geukes, K., Breil, S. M., Hutteman, R., Nestler, S., Küfner, A. C. P., & Back, M. D. (2019). Explaining the longitudinal interplay of personality and social relationships in the laboratory and in the field: The PILS and the CONNECT study. *PLOS ONE*, 14(1). doi: 10.1371/journal.pone.0210424
- Gordon, G., Breazeal, C., & Engel, S. (2015). Can Children Catch Curiosity from a Social Robot? In Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction - HRI'15. ACM Press. doi: 10.1145/2696454.2696469
- Grassberger, P. (1986). Toward a quantitative theory of self-generated complexity. International Journal of Theoretical Physics, 25(9), 907–938. doi: 10.1007/bf00668821
- Guckelsberger, C., Salge, C., & Colton, S. (2016). Intrinsically motivated general companion NPCs via Coupled Empowerment Maximisation. In 2016 IEEE Conference on Computational Intelligence and Games (CIG) (pp. 150–157). IEEE. doi: 10.1109/cig.2016.7860406
- Guckelsberger, C., Salge, C., & Togelius, J. (2018). New And Surprising Ways to Be Mean. In 2018 IEEE Conference on Computational Intelligence and Games (CIG) (pp. 1–8). IEEE. doi: 10.1109/cig.2018.8490453
- Hayashi, K., Shiomi, M., Kanda, T., & Hagita, N. (2010). Who is appropriate? A robot, human and mascot perform three troublesome tasks. In 19th International Symposium in Robot and Human Interactive Communication. IEEE. doi: 10.1109/roman.2010.5598661

- Huang, H.-M., Messina, E., Wade, R., English, R., Novak, B., & Albus, J. (2004). Autonomy Measures for Robots. In ASME 2004 International Mechanical Engineering Congress and Exposition (pp. 1241–1247). doi: 10.1115/IMECE2004-61812
- Jaques, N., Lazaridou, A., Hughes, E., Gulcehre, C., Ortega, P., Strouse, D., ... De Freitas, N. (2019). Social influence as intrinsic motivation for multi-agent deep reinforcement learning. In K. Chaudhuri & R. Salakhutdinov (Eds.), *Proceedings of the 36th international conference on machine learning* (Vol. 97, pp. 3040–3049). PMLR. Retrieved from https://proceedings .mlr.press/v97/jaques19a.html
- Judd, C. M., James-Hawkins, L., Yzerbyt, V., & Kashima, Y. (2005). Fundamental dimensions of social judgment: Understanding the relations between judgments of competence and warmth. *Journal of Personality and Social Psychology*, 89(6), 899–913. doi: 10.1037/0022-3514.89.6.899
- Kanda, T., Shiomi, M., Miyashita, Z., Ishiguro, H., & Hagita, N. (2010). A Communication Robot in a Shopping Mall. *IEEE Transactions on Robotics*, 26(5), 897–913. doi: 10.1109/tro.2010.2062550
- Kaplan, F., & Oudeyer, P.-Y. (2004). Maximizing Learning Progress: An Internal Reward System for Development. In F. Iida, R. Pfeifer, L. Steels, & Y. Kuniyoshi (Eds.), *Embodied Artificial Intelligence* (Vol. 3139, pp. 259–270). Springer Berlin Heidelberg. doi: 10.1007/978-3-540-27833-7<sup>19</sup>
- Klyubin, A. S., Polani, D., & Nehaniv, C. L. (2005). Empowerment: A Universal Agent-Centric Measure of Control. In 2005 IEEE Congress on Evolutionary Computation (Vol. 1, pp. 128– 135). IEEE. doi: 10.1109/cec.2005.1554676
- Kulms, P., & Kopp, S. (2018). A Social Cognition Perspective on Human–Computer Trust: The Effect of Perceived Warmth and Competence on Trust in Decision-Making With Computers. *Frontiers in Digital Humanities*, 5. doi: 10.3389/fdigh.2018.00014
- Lucasfilm Ltd. (2015). *Star Wars.* (https://www.starwars.com/films/, accessed November 2020)
- Martius, G. (2013). Implementation of predictive information maximization. (https://github .com/georgmartius/lpzrobots/blob/d2e6bbd164d902cdaa57eef154ed353ee0027236/ selforg/controller/pimax.cpp, accessed December 2020)
- Martius, G., Der, R., & Ay, N. (2013a). Appendix with derivations and technical detail. doi: https://doi.org/10.1371/journal.pone.0063400.s001
- Martius, G., Der, R., & Ay, N. (2013b). Information Driven Self-Organization of Complex Robotic Behaviors. PLOS One, 8(5), 1–14. doi: 10.1371/journal.pone.0063400

- Martius, G., Jahn, L., Hauser, H., & Hafner, V. V. (2014). Self-exploration of the Stumpy Robot with Predictive Information Maximization. In *From Animals to Animats 13* (pp. 32– 42). Springer International Publishing. doi: 10.1007/978-3-319-08864-8'4
- Maturana, H. R., & Varela, F. J. (1991). Autopoiesis and Cognition: The Realization of the Living (Vol. 42; R. S. Cohen & M. W. Wartofsky, Eds.). Springer Science & Business Media. doi: 10.1007/978-94-009-8947-4
- Merrick, K. E., & Maher, M. L. (2009). Motivated Reinforcement Learning (1st ed.). Springer Berlin Heidelberg. doi: 10.1007/978-3-540-89187-1
- Noguchi, K., Gel, Y. R., Brunner, E., & Konietschke, F. (2012). nparLD: AnRSoftware package for the nonparametric analysis of longitudinal data in factorial experiments. *Journal of Statistical Software*, 50(12). doi: 10.18637/jss.v050.i12
- Oliveira, R., Arriaga, P., Correia, F., & Paiva, A. (2019). The stereotype content model applied to human-robot interactions in groups. In 2019 14th ACM/IEEE international conference on human-robot interaction (HRI). IEEE. doi: 10.1109/hri.2019.8673171
- O'Regan, J. K., & Noë, A. (2001). A sensorimotor account of vision and visual consciousness. Behavioral and brain sciences, 24(5), 939–973. doi: 10.1017/S0140525X01000115
- Orne, M. T. (1962). On the social psychology of the psychological experiment: With particular reference to demand characteristics and their implications. *American Psychologist*, 17(11), 776–783. doi: 10.1037/h0043424
- Oudeyer, P.-Y., Gottlieb, J., & Lopes, M. (2016). Intrinsic motivation, curiosity, and learning. In *Motivation - theory, neurobiology and applications* (pp. 257–284). Elsevier. doi: 10.1016/bs.pbr.2016.05.005
- Oudeyer, P.-Y., & Kaplan, F. (2008). How can we define intrinsic motivation ? In Proceedings of the 8th international conference on epigenetic robotics: Modeling cognitive development in robotic systems. Brighton, United Kingdom. Retrieved from https://hal.inria.fr/ inria-00420175
- Oudeyer, P.-Y., & Kaplan, F. (2009). What is Intrinsic Motivation? A Typology of Computational Approaches. Frontiers in Neurorobotics, 1, 6. doi: 10.3389/neuro.12.006.2007
- Oudeyer, P.-Y., Kaplan, F., & Hafner, V. V. (2007). Intrinsic Motivation Systems for Autonomous Mental Development. *IEEE Transactions on Evolutionary Computation*, 11(2), 265–286. doi: 10.1109/tevc.2006.890271
- Paolo, E. A. D. (2004). Unbinding Biological Autonomy: Francisco Varela's Contributions to Artificial Life (Vol. 10) (No. 3). MIT Press - Journals. doi: 10.1162/1064546041255566

- Pinillos, R., Marcos, S., Feliz, R., Zalama, E., & Gómez-García-Bermejo, J. (2016). Long-term assessment of a service robot in a hotel environment. *Robotics and Autonomous Systems*, 79, 40–57. doi: 10.1016/j.robot.2016.01.014
- Research Network for Self-Organization of Robot Behavior. (2015). Research Network for Self-Organization of Robot Behavior: Videos. (https://robot.informatik.uni-leipzig.de/ videos, accessed February 2020)
- Rich, C., Ponsler, B., Holroyd, A., & Sidner, C. L. (2010). Recognizing engagement in humanrobot interaction. In 2010 5th ACM/IEEE International Conference on Human-Robot Interaction (HRI) (pp. 375–382). doi: 10.1109/HRI.2010.5453163
- Rosas, F. E., Mediano, P. A. M., Jensen, H. J., Seth, A. K., Barrett, A. B., Carhart-Harris, R. L., & Bor, D. (2020). Reconciling emergences: An information-theoretic approach to identify causal emergence in multivariate data. *PLOS Computational Biology*, 16(12), e1008289. doi: 10.1371/journal.pcbi.1008289
- Rosenberg, S., Nelson, C., & Vivekananthan, P. S. (1968). A multidimensional approach to the structure of personality impressions. *Journal of Personality and Social Psychology*, 9(4), 283–294. doi: 10.1037/h0026086
- Ryan, R. M., & Deci, E. L. (2000a). Intrinsic and Extrinsic Motivations: Classic Definitions and New Directions. Contemporary Educational Psychology, 25(1), 54–67. doi: 10.1006/ceps.1999.1020
- Ryan, R. M., & Deci, E. L. (2000b). Self-determination theory and the facilitation of intrinsic motivation, social development, and well-being. *American Psychologist*, 55(1), 68–78. doi: 10.1037/0003-066X.55.1.68
- Scheunemann, M. M. (2018). Code repository: spherocpp. Retrieved 2020-12-12, from https://
  gitlab.com/scheunemann/spheropp (https://gitlab.com/scheunemann/spheropp, accessed December 2020)
- Scheunemann, M. M. (2021a). Autonomous and Intrinsically Motivated Robots for Sustained Human-Robot Interaction (Doctoral dissertation). doi: 10.18745/TH.23936
- Scheunemann, M. M. (2021b). Supplementary Material for Human Perception of Intrinsically Motivated Autonomy in Human-Robot Interaction. (https://mms.ai/adb2021, accessed November 2021)
- Scheunemann, M. M., Cuijpers, R. H., & Salge, C. (2020). Warmth and Competence to Predict Human Preference of Robot Behavior in Physical Human-Robot Interaction. In Proceedings of the 29th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN) (pp. 1340–1347). IEEE. doi: 10.1109/RO-MAN47096.2020.9223478

- Scheunemann, M. M., & Dautenhahn, K. (2017). Bluetooth low energy for autonomous humanrobot interaction. In Proceedings of the companion of the 2017 ACM/IEEE international conference on human-robot interaction. ACM. doi: 10.1145/3029798.3036663
- Scheunemann, M. M., Salge, C., & Dautenhahn, K. (2019). Intrinsically Motivated Autonomy in Human-Robot Interaction: Human Perception of Predictive Information in Robots. In K. Althoefer, J. Konstantinova, & K. Zhang (Eds.), *Towards autonomous robotic systems* (pp. 325–337). Cham: Springer International Publishing. doi: 10.1007/978-3-030-23807-0'27
- Schmidhuber, J. (1991). Curious model-building control systems. In [proceedings] 1991 ieee international joint conference on neural networks (pp. 1458–1463). IEEE. doi: 10.1109/ijcnn.1991.170605
- Singh, S., Barto, A. G., & Chentanez, N. (2004). Intrinsically motivated reinforcement learning. In Proceedings of the 17th international conference on neural information processing systems (p. 1281–1288). Cambridge, MA, USA: MIT Press.
- Smithers, T. (1997). Autonomy in Robots and Other Agents. Brain and Cognition, 34(1), 88–106. doi: https://doi.org/10.1006/brcg.1997.0908
- Sphero, Inc. (2020). Licensed products: BB-8. (https://support.sphero.com/category/ kxwbdyqeyq-bb-8, accessed December 2020)
- Steels, L. (2004). The Autotelic Principle. In *Embodied Artificial Intelligence* (Vol. 3139, pp. 231–242). Springer Berlin Heidelberg. doi: 10.1007/978-3-540-27833-7<sup>17</sup>
- Stubbs, K., Hinds, P. J., & Wettergreen, D. (2007). Autonomy and Common Ground in Human-Robot Interaction: A Field Study. *IEEE Intelligent Systems*, 22(2), 42-50. doi: 10.1109/MIS.2007.21
- Thrun, S., Burgard, W., & Fox, D. (2005). Probabilistic Robotics. MIT Press.
- Tremoulet, P. D., & Feldman, J. (2000). Perception of Animacy from the Motion of a Single Object. Perception, 29(8), 943–951. doi: 10.1068/p3101
- Van Rensburg, E. J. J., Rechnitzer, A., Causo, M. S., & Whittington, S. G. (2001). Selfaveraging sequences in the statistical mechanics of random copolymers. *Journal of Physics A: Mathematical and General*, 34(33), 6381. Retrieved from http://stacks.iop.org/ja/34/ 6381
- Wojciszke, B., Bazinska, R., & Jaworski, M. (1998). On the Dominance of Moral Categories in Impression Formation. *Personality and Social Psychology Bulletin*, 24(12), 1251–1263. doi: 10.1177/01461672982412001
- Yatani, K. (2016). Effect Sizes and Power Analysis in HCI. In J. Robertson & M. Kaptein (Eds.), Modern Statistical Methods for HCI (pp. 87–110). Springer. doi: 10.1007/978-3-319-26633-6'5

Zahedi, K., Martius, G., & Ay, N. (2013). Linear combination of one-step predictive information with an external reward in an episodic policy gradient setting: a critical analysis. *Frontiers in Psychology*, 4. doi: 10.3389/fpsyg.2013.00801