Autonomous and Intrinsically Motivated Robots for Sustained Human-Robot Interaction

Marcus M. Scheunemann

A thesis submitted in partial fulfilment of the requirements of the University of Hertfordshire for the degree of

Doctor of Philosophy

June 1, 2020

Abstract

A challenge in using fully autonomous robots in human-robot interaction (HRI) is to design behavior that is engaging enough to encourage voluntary, long-term interaction, yet robust to the perturbations induced by human interaction. It has been repeatedly argued that intrinsic motivations (IMs) are crucial for human development, so it seems reasonable that this mechanism could produce an adaptive and developing robot, which is interesting to interact with. This thesis evaluates whether an intrinsically motivated robot can lead to sustained HRI.

Recent research showed that robots which 'appeared' intrinsically motivated raised interest in the human interaction partner. The displayed IMs resulted from 'unpredictably' asking a question or from a self-disclosing statement. They were designed with the help of pre-defined scripts or teleoperation. An issue here is that this practice renders the behavior less robust toward unexpected input or requires a trained human in the loop.

Instead, this thesis proposes a computational model of IM to realize fully autonomous and adaptive behavior generation in a robot. Previous work showed that predictive information maximization leads to playful, exploratory behavior in simulated robots that is robust to changes in the robot's morphology and environment. This thesis demonstrates how to deploy the formalism on a physical robot that interacts with humans.

The thesis conducted three within-subjects studies, where participants interacted with a fully autonomous Sphero BB8 robot with two behavioral regimes: one realizing an adaptive, intrinsically motivated behavior and the other being reactive, but not adaptive. The first study contributes to the idea of the overall proposed study design: the interaction needs to be designed in such a way, that participants are not given any idea of the robot's task. The second study implements this idea, letting participants focus on answering the question of whether the robots are any different. It further contributes ideas for a more 'challenging' baseline behavior motivating the third and final study. Here, a systematic baseline is generated and shows that participants perceive it as almost indistinguishable and similarly animated compared to the intrinsically motivated robot. Despite the emphasis on the design of similarly perceived baseline behaviors, quantitative analyses of post-interaction questionnaires after each study showed a significantly higher perception of the dimension 'Warmth' for the intrinsically motivated robot compared to the baseline behavior. Warmth is considered a primary dimension for social attitude formation in social cognition. A human perceived as warm (i.e. friendly and trustworthy) experiences more positive social interactions.

The Robotic Social Attribute Scale (RoSAS) implements the scale dimension Warmth for

the HRI domain, which has been validated with a series of still images. Going beyond static images, this thesis provides support for the use and applicability of this scale dimension for the purpose of comparing behaviors. It shows that participants prefer to continue interacting with the robot they perceive highest in Warmth.

This research opens new research avenues, in particular with respect to different physical robots and longitudinal studies, which are ought to be performed to corroborate the results presented here. However, this thesis shows the general methods presented here, which do not require a human operator in the loop, can be used to imbue robots with behavior leading to positive perception by their human interaction partners, which can yield sustained HRI.

Contents

Chapter 1.

Introduction

1.1. Motivation

This thesis investigates whether intrinsically motivated, fully autonomous robots are a compelling tool to pave the way toward sustained human-robot interaction (HRI).

I started my work with the aim to create a robot mediator for the therapy given to children with autism. My goal was to find dyads of children, who share the interest to interact with a robot, to bring them together and to help develop friendships between the children: a rare occurrence among children with autism. I quickly faced the challenge of implementing *interesting* behavior, which would enable and sustain interactions between the children and the robot. This became a tedious task, keeping the robot predictable enough to have the children with autism feel comfortable interacting with it, but also giving the robot enough variability to keep the interaction interesting, to avoid boredom among the children. In other words, I tried to sustain the interaction. I quickly found myself in a situation where I tweaked sets of parameters per each of the 20 children I worked with. Also, more generally, and independent of the particular child, I started to have different behavior implementations for the robot, to keep the children interested. It became clear that this task required a specialist for tweaking the right parameters for each of the children, similarly to the requirement of a specialist to operate robots for other therapeutic uses, or general therapy for children with autism. My goal, however, has been to bring a robot to children's homes, to families who cannot afford to have their children visit a special care unit, or to families living so remote that a care unit is simply out of scope.

In HRI, scientists have focused on the interaction with robots in constrained environments. For example, robots which offer guidance and help in a hotel lobby (Pinillos et al. 2016) or robots which play games to educate children (e.g. Wainer et al. 2014). This way robots could follow a set of given rules, which they could handle with pre-scripted behavior. What is still far out of reach is an understanding of *how* a robot can *develop* its behavior. How will they learn to cope with humans in an environment shaped by humans? This thesis contributes to close that gap.

To achieve this, I took a step back, and looked into the research of what is known about behavior generation. Broadly speaking, how can we have a robot grow with us, so that it learns the needs of the humans, rather than getting those needs induced by some code enabling a set of rules. When thinking about this, the idea of *autonomy* comes to mind. Autonomy is a complex term without a central definition across research fields (Boden 2008; Ziemke 2008). Autonomy can merely mean robots that are not directly controlled by a human operator, with autonomy used as a dimension of the experimental design (Huang et al. 2004; Stubbs et al. 2007). Autonomy might also be referring to the concept of self-making or self-law-giving, which is closely related to autopoiesis (Maturana and Varela 1991; Froese and Ziemke 2009). In Self-Determination Theory (SDT) 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 is usually not referred to in this way.

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. Closely related to SDT, and originating from psychology, the concept of intrinsic motivation (IM) describes these drives (Ryan and Deci 2000).

IM has been extensively studied by psychologists since the 1940s, yielding a large body of work. Both a multitude of approaches for describing the phenomenon and a variety of definitions for IM exist. Hull (1943) introduced the theory of drives, which are construed as stressful perturbations (e.g. hunger) an organism tries to reduce. This idea was challenged by White (1959), with the argument that activities based on IMs have different, in particular non-homeostatic, dynamics. Ryan and Deci (2000), to present one of the most popular examples, define intrinsic motivation as the "doing [...] of an activity for its inherent satisfactions rather than for some separable consequence". An 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" (ibid., pg. 56).

Concepts like *fun* and *challenge* are presented as crucial for the definition of IM in psychology and yet the literature lacks consensus on what these concepts are (Oudeyer and Kaplan 2009). This missing consensus and the resulting vagueness of the definition makes it impossible to transfer it directly onto a robotic system. Research in HRI circumvents this issue by mimicking IM designed by humans. For example, 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 recurring visitors, without increasing its services. One existing approach in social robots in education 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 as well, which in turn enhances learning and memory retention (Oudeyer, Gottlieb, et al. 2016). Curiosity is part of the broader concept of IM (ibid.), or is even used synonymously for intrinsic motivation (Schmidhuber 1991). While the studies above are promising to use IM in robots in order to sustain HRI, they raised the same issues posed at the beginning of this introduction: they either needed a constrained context, 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 and Matarić 2019). This is caused by 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). However, developing a robot driven by an actual IM implementation might offer a solution to both problems. If successful, this would provide us with a robust behavior generation mechanism that allows us to "Escape Oz" (Clabaugh and Matarić 2019), while also producing behavior that appears curious, or similarly engaging to the human interaction partner. This would reduce the reliance on human adaptation or teleoperation, and could provide a promising pathway toward having robots more easily deployed in our everyday lives.

However, the question remains, how to build a computational model based on relatively vague psychological definitions? Already early on, Schmidhuber (1991) looked into the computation of artificial curiosity, where an agent chooses actions to maximize its learning progress. Later, Oudeyer and Kaplan (2008) characterize IM 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.

Ultimately, any heteronomy during the development or creation of an agent would make them extrinsic and hence undermine their very nature, i.e. computational models of IMs on robots are usually put on those robots by humans, and are thus actually extrinsic. Therefore, computational models of IM are another attempt to merely reproduce the behavior or functionality of genuinely IM in organisms, with the added benefit that no human operator is present in the loop. When I talk about IM on a robot, I exclusively refer to the above definition of IM from a computational perspective.

Nowadays, there is a range of formal models that roughly fall under the header of intrinsic motivation, such as the free energy principle (Friston 2010), homeokinesis (Der and Martius 2012), empowerment (Klyubin et al. 2005), learning progress (Kaplan and Oudeyer 2004), the autotelic principle (Steels 2004). 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 IM could generate more believable non-player characters (Merrick and Maher 2009) or produce generic companions (Guckelsberger et al. 2016). So far, IMs have been deployed on simulated and physical robots (e.g. Oudeyer, Kaplan, and Hafner 2007; Der and Martius 2012; Martius, Jahn, et al. 2014), but, as far as I know, there has not yet been an HRI study that evaluated the perception of intrinsically motivated robots from the perspective of humans.

This thesis uses predictive information (PI) maximization as a realization of a computational model for IM, an "internal, task independent motivation for the development of [...] behavior" (Der, Güttler, et al. 2008). The PI formalism fits the above characterization: it is an information-theoretic measure in the sensor space, which quantifies the information contained in a preceding series of sensory input for a consecutive sequence. An agent maximizing its predictive information will derive a rich and variable sensor future while, at the same time, keeping it as predictable as possible from past sensory input. Martius, Der, et al. (2013b) developed an approximation for computing PI for non-linear and non-stationary dynamics, which allows applying the measure to a real robotic system, inhabiting a real world: time-local predictive information (TiPI). The core idea is to only maximize PI based on states that are a few steps back in time, rather than the full history. For animals and humans, it is essential to modify their behavior intrinsically motivated to enable further development. An algorithm that is capable of enabling this in a robot might be perceived as more *familiar* by a human participant, increasing the human's interest in the robot as an interaction partner, and thus bringing the search of sustainable HRI a step closer. One contribution of this thesis is the evidence that the TiPI maximization can be applied to a physical robot in an HRI scenario.

Applying the above formalism to a robot, and letting it generate its behavior by itself, may raise skepticism for some people. One particular concern is likely to be the predictability of its behavior. Section 2.3 outlines that the above formalism, by design, is not predictable and results in exploratory behavior. The authors argue that it allows for self-organization of complex behavior (ibid.). In other words, such a robot is able to independently explore its environment by exploring its sensor space. This, in turn, as it has been argued above, is seen as a key feature in autonomy for both natural and artificial systems (Boden 2008).

In some HRI scenarios this may yield discomfort: humans probably expect a robot to be predictable when it comes to *critical* tasks. Otherwise, depending on the robot's strength, there is a risk that the robot will cause damage or harm. The literature in behavioral science is non-conclusive when it comes to the perception of unpredictable behavior in interpersonal interaction. However, there is evidence that a degree of unpredictability might be a cause for the human's interest. Grillon et al. (2004) showed that unpredictable and aversive stimuli lead to a more sustained level of anxiety when compared to stimuli that were predictable and aversive. In contrast, agent behavior research by Bickmore et al. (2010) indicates that a degree of unpredictability in the behavior of an animated human might be essential to sustain engagement of the human participants in longitudinal interactions. In other words, if a longitudinal interaction is too predictive, the human participant might simply get bored. Fukuda and Ueda (2010) conducted an HRI experiment investigating the difference between a fully reactive robot, and a robot which is *mostly* reactive: that is a robot which responded as predicted in most scenarios, but a level of noise also let it react unpredictably. They found that the participants who *observed* the HRI scenario preferred the fully reactive robot. However, participants who *interacted* with the robot seemed to value a level of unpredictability. Interestingly, maximizing TiPI creates behavior somewhere between chaos and predictability. These properties give more weight to the hypothesis that an intrinsically motivated robot makes a human participant want to interact further with the robot, and therefore sustains the interaction.

This leads us to the critical question: how can the human's intention to continue engaging in an interaction – the sustainability of HRI – be measured? Unarguably, the best way would be to deploy robot companions into households and measure how much people engage in interacting with the robots. This, however, yields a chicken and egg problem. We know such robots do not exist as of now, so there still needs to be development in order to pass this test. An alternative approach, as the one employed in this thesis, is to indirectly measure sustainability with the help of other dimensions. In social cognition, one popular measure to understand human's social attitude formation and their resulting behaviors are the two central (or universal) dimensions *Warmth* and *Competence*. Fiske et al. (2007) argue that, in order for humans to be perceived positively and experience more interaction by their peers, humans need to be perceived high in Competence and Warmth. The dimension Warmth is primary to understand whether this perception is *positive*, while Competence measures its intensity (e.g. Cuddy et al. 2007; Abele, Hauke, et al. 2016). Carpinella et al. (2017) designed the Robotic Social Attribute Scale (RoSAS) transferring the dimensions of Warmth and Competence to the domain of HRI. This questionnaire provides an ideal handle for investigating how the behavior generated by TiPI maximization is perceived. Intuitively, being able to generate behavior capable of sustaining the interaction should suffice to score high in Warmth. How high in Warmth is high enough? This is a matter of investigation. However, one main contribution of this thesis is that it provides evidence that there is a link between participants' responses to Warmth and their preferred behavior. Specifically, the thesis shows that human participants prefer to continue interacting with the robot behavior they perceive highest in Warmth.

In summary, there are two main research questions this thesis addresses. Firstly, how to measure sustainability in HRI? Secondly, can an intrinsically motivated robot sustain HRI? In order to answer these research questions, a few intermediate research objectives emerged. One has been already outlined: is there a computational model of IM that can be applied to a robot in an HRI scenario. This thesis shows that TiPI is a good candidate. Another question concerns the design of a suitable study to understand the impact of intrinsically motivated autonomy in robots. This thesis provides a design that addresses several issues related to these questions, such as a suitable robot platform, a suitable baseline behavior and a way to encourage unbiased interaction.

The thesis consists of three interaction studies that approach these intermediate objectives incrementally and contributes to answering the main research questions outlined in section 1.2. Section 1.3 provides an overview of this thesis, while section 1.4 summarizes its contributions in regard to the research questions and objectives.

1.2. Research questions

RQ1 Can dimensions of social cognition be employed to measure human participants' preferences of robot behavior in order to understand what may sustain the interaction between humans and robots?[1]

RQ2 Can an autonomously, intrinsically motivated robot, sustain the interaction with humans?

In order to address the main research questions above, two intermediate research objectives emerged:

O1 To identify a suitable computational model of IM that can be applied to a robot in an HRI scenario.

O2 To develop a good study design for investigating the human perception of intrinsically motivated robots.

1.3. Overview

Chapter 2. This chapter presents the background and technical developments that are needed for this thesis. It first focuses on *what* exactly constitutes as intrinsic motivation (IM), and *how* it could be applied to the robot. The chosen robot platform was the BB8 version of a Sphero, which is described in section 2.5. The limited, non-humanoid robot was chosen in order to limit participants' expectations of the robot's capabilities. A computational

[1]Note that the focus on social cognition dimensions was a result of the first study of chapter 3.

model of IM needs to fulfill two main criteria in order to be robustly applicable to this robot platform in a real-world human-robot interaction (HRI) scenario: (i) it needs to cover an infinite number of states, i.e., it needs to be able to work on a large range of continuous sensor input and (ii) it needs to be computable.

The chapter proposes a computational model for IM which maximizes the time-local predictive information (TiPI) introduced by Martius, Der, et al. (2013b). TiPI uses past observations to quantify how many of the future states are predictable, therefore to maximize this quantity the robot has to generate a rich sensory input. The chapter derives the formulas to compute TiPI, along with a discussion of the needed approximations and their implications when applying them to robots. Two approximations have a direct impact on designing the studies: (i) prediction errors need to be very small and Gaussian and (ii) the noise must be independent of controller parameters. It is explained that approximation (i) determined the choice of sensory input, which made proprioceptive sensors, such as speed or acceleration sensors, a sensible choice. However, section 2.6 shows that for the chosen robot platform the measured servo speed is not solely dependent on the control for that servo. This means that when using the measured servo speed as an input, the error of that proprioceptive sensor depends on the controller parameters, which violated approximation (ii). The chapter presents a simple motion model based on linear equations to break this dependency and shows that the use of this model enabled a rich behavior generation.

Another technical development presented in this chapter is the sensor system which was used in the final study, presented in chapter 5. It used the received signal strength of Bluetooth Low Energy (BLE) to derive proximity information between the robot and the human. It can also enable a robot to distinguish between humans in its vicinity and enable it to recognize touch gestures, but these functionalities were not employed in this thesis. The sensor system is a contribution of this thesis.

The main contribution of the chapter, however, is that it addresses the first research objective O1: TiPI maximization is a possible computational model for IM, which can be applied to a real, minimal robot in an HRI scenario.

Chapter 3. This chapter presents the first of three interaction studies and discusses its results. It was a within-subjects study ($N = 16$) which compared an intrinsically motivated robot to a reactive baseline behavior. The main contribution of this chapter is the systematic development of this baseline behavior. Before conducting a study, the intrinsically motivated robot was placed in the study environment and it adapted its parameters. The adaptation was then stopped and the robot used constant, but adapted, parameters to serve as a baseline during the study.

The robots in the study were both reactive to the same input sensors and both used the built-in balancing controller to locomote. The controller keeps the robots upright by using the heading and speed information as input. This also created a wide variety of behaviors in the baseline.

The chapter also presents the first approach to design a suitable study. The main idea was to create a situation that *enforced* HRI, while allowing participants to see the robots' capability to adapt. Therefore the participants were tasked with preventing the robot from falling, while a robot moved on a table with an open edge. Additional complexity was added to the environment by adjusting the altitude and friction on different areas of the table. This was included to further emphasize the strength of the adaptive robot.

The chapter then discusses the results of the study, which highlighted areas that needed to be redesigned. Evidently, the task created the assumption that the robot's goal was to remain on the table, and participants considered any deviation from this as *faulty* or suicidal behavior. As a result, and contrary to the hypothesis, the intrinsically motivated robot was considered the least competent. However, the results also showed that a medium effect of Warmth was perceived in the intrinsically motivated robot. The chapter goes on to explain how the dimensions Warmth and Competence are central to social cognition, and that humans who are perceived as high in Warmth experience more positive social interactions from their peers (e.g. Abele and Wojciszke 2007). The chapter then guides the focus of this research toward the effects of Warmth.

Chapter 4. This chapter presents the second study ($N = 24$) and the specific design changes which followed the previous study. These changes were (i) concentrate on the perception of Warmth and (ii) enable behavior generation solely based on the robot's IMs.

Firstly, the environment was changed in order to fully focus on the perception of Warmth. The table was now circular, had no friction or altitude variations and was fully enclosed. Without the open edge, a second change became necessary to motivate interaction. The participants were now presented with a *game-like* task, with instructions "to find out whether the two presented robots are different". The game was thought to prevent the participants from implicitly assuming a robot's goal, which does not match the robot's behavior. This game design also encouraged interaction by exploiting the participant's interest to perform *well* (e.g. Orne 1962). The study introduces the wand-shaped tool. The idea is that a tool that was provided for interacting with the robot causes the participants to feel the urge to use it, which in turn further encourages interaction.

Secondly, the intrinsically motivated robot used the motion model (cf. 2.6) to directly change the speed of its two servos, instead of using the balancing controller. This way the robot's behavior is only influenced by its IMs, unconstrained from additional software. This allowed to further focus the analysis on the perception of intrinsically motivated autonomy.

The results in this chapter indicated that the changes to the study design were successful, as the participants perceived both robots as similarly competent and intelligent. Additionally, the results on the Warmth dimension were not influenced by the perception of the robot's Competence (cf. Fiske et al. 2007). And, most importantly, the chapter presents results that show the intrinsically motivated robot was perceived statistically significantly warmer than the baseline behavior.

However, results also showed that both behaviors were perceived very differently and the intrinsically motivated robot was even perceived as more animated. It is known that humans perceive robots and even objects as animated if the cause of their movement changes are not obvious to the observer (e.g. Castro-González et al. 2016). This raised the concern that the motion regimes of the baseline robot were too different (maybe even too predictable) in comparison to those of the intrinsically motivated robot.

Chapter 5. This chapter presents the final study ($N = 36$) of the thesis, which focused on confirming the effect of Warmth. Therefore, the chapter describes the changes that were made while addressing the concerns found in the previous study: was the intrinsically motivated robot perceived statistically significantly more warm because (i) it was more animated and had a different behavioral regime than the baseline, or because (ii) participants could see that the robot was responding more directly to their perturbations, or because of (iii) the exploratory, playful behavior generated by the robot's IMs?

The first change was a new baseline behavior, which included the implementation of *fake* adaptivity. This meant that the robot now received updates that were recorded during a previous run of an intrinsically motivated robot and then *replayed* for the baseline. The chapter explains that the updates were not random, but also not truly adaptive. Additionally, the final study now had two conditions with intrinsically motivated robots: one with a proximity sensor and one without. The sensor enabled the robot to perceive human proximity by using BLE signal strength between the wand-shaped tool and the robot (cf. 2.7). This allowed the robot to respond directly to the participants' input.

The results presented in this chapter show that the new baseline behavior, based on a *parameter replaying controller*, was perceived as much more similar, including similarly animated and competent, when compared to the intrinsically motivated robots. This made it a good candidate for comparison. Furthermore, despite these similarities, the two intrinsically motivated robots were both perceived as more warm. This underlined evidence from the previous studies: an intrinsically motivated robot in an embodied social cognition scenario elicits a feeling of Warmth. Additionally, it was shown that the proximity amplified the feeling of Warmth. Evidently, this means that an intrinsically motivated robot that can adapt toward the proximity of the human interaction partner will elicit a stronger feeling of Warmth. The chapter interprets these results as strong evidence that this effect was mainly routed in the robot's IMs, and not because of (i) or (ii).

The chapter presents another important contribution to the thesis. It describes details of the investigation into whether the knowledge of social cognition transfers to HRI: do participants prefer to interact with the robot which they perceived highest in Warmth? The results showed that robots perceived highest in Warmth are likewise the robots that human participants preferred to interact with the most. The chapter explains the importance of these results, giving two underlying reasons. Firstly, it gave more weight to the previous study results and provided evidence that all intrinsically motivated robots were preferred over the baseline robot. Secondly, it showed that there was a link between human attitude formation toward peers and robots. If further evidence can confirm the results, this provides future research with a good measure of human preference, by using tools established in social cognition.

Chapter 6. This chapter concludes the thesis and addresses this thesis' evidence that a fully autonomous, intrinsically motivated robot elicits a feeling of Warmth in the human interaction partner, which in turn shows that it can potentially sustain the interaction. The chapter provides ideas for possible future studies, which can further confirm these results. One central idea is to measure the actual interaction time with robots which are already placed in human-inhabited environments, such as the therapeutic robot Paro or vacuum cleaning robots.

1.4. Contributions of the thesis

This thesis, to the best of my knowledge, is the first to investigate the impact that intrinsically motivated autonomy has on human perception. It used a computational model, based on information theory approaches, to implement intrinsic motivation onto a robot and then analyzed the effect on human perception. Furthermore, the thesis investigated whether human attitude formation studied in social cognition is transferable to human attitude formation toward robots. Therefore, this thesis contributes knowledge to the four disciplines: robotics, HRI, information theory, and social cognition.

1.4.1. Information theory

This work contributes to information theory by presenting an information-theoretically based computational model of IM, which is shown to have had an impact on human perception. Intrinsically motivated autonomy was implemented using predictive information (PI) maximization, arguably "the most natural complexity measure for time series" (e.g. Bialek et al. 2001). The behavior generated from PI maximization has been previously judged by observations, which claimed that it results in "playful and exploratory" behavior (Martius, Der, et al. 2013b). This and all other computational models usually followed theoretically sound approaches, but how they performed in a real-world setting had not yet been analyzed. This thesis therefore contributes a quantitative argument to the list. Without a doubt, humans are our best judges of how humans perceive robotic behavior. This thesis shows that human participants perceived an intrinsically motivated robot as more warm, which is the primary concept to understand *positive* human attribute perception. This means that PI maximization, as a candidate for a computational model of IM, created a feeling of Warmth toward an artificial agent.

On the more practical side, issues with the original TiPI implementation by Martius, Der, et al. (ibid.) were discovered during the process of actually implementing the computational model. In particular, this concerned the parameter computation for more than two steps back. These observations were addressed to the authors and were also addressed publicly (see Scheunemann 2018d). In addition, there are contributions that enable compiling the `lpzrobots simulator` on current architectures, a simulator created by the authors and used for their simulation experiments (see Scheunemann 2018c).

1.4.2. Robotics

The thesis contributes to the field of robotics by providing guidance for *how* to implement PI maximization onto a robot. It outlines the approximations of the formalism that was used and the implications for applying them to a robot. This contributes to further research in this area, especially for research that uses the same approach as this thesis. The guidance can also be used as an orientation for related computational models of IM. In general, similar approaches could follow up on the same approximations to make the computation of the underlying information theory concepts possible.

On the more technical side of robotics, this work contributes code enhancements to run the Sphero robot. In fact, the off-the-shelf robot was chosen to expedite the start of research. The next chapter explains the hardware decisions in more detail. It turned out, however, that the code to run the robot had a variety of issues. The provided official `JavaScript` framework, for example, did not parse the protocol correctly, which yielded a stuck robot behavior that forced a re-start of the controller. In addition, sensor values, such as the roll angle of the robot, were faulty, along with the order of how the servo speed was set. These insights were forwarded to the company or made into public pull requests to the (back then) official framework. The four contributions which are most closely related to this thesis are the following. The first allows the retrieval of quaternion readings (see Scheunemann 2017d) and another fixed the documentation (see Scheunemann 2017b). Arguably, the most important are the two more technical contributions: one which allows making a connection to multiple robots simultaneously (see Scheunemann 2017c) and the other enables parsing the robot's protocol correctly to prevent the robot from crashing (see Scheunemann 2017a).

When the above contributions were proposed to the official framework, it occurred already that the framework development would not continue for long. This was one reason why the knowledge was also transferred into an own framework development based on `C++`. The other reason was that the `C++` language allows writing code that is closer to the hardware, which enabled controlling the robot from embedded, computationally limited systems. The `C++` library is publicly available (see Scheunemann 2017e; Scheunemann 2018b). The first attempts of this thesis involved working with autistic children in a nursery. In this context, the library was successfully used on a *Raspberry Pi 2* during the initial play sessions with children.

The public availability of the framework creates a simpler way of reproducing the experiments. More importantly, it allows other researchers to start straight away without the developmental issues present prior to this thesis. The robot platform eventually became discontinued altogether in 2018. This is a very common issue to robotic related contributions: the developed code is bound to a specific robot platform and the contribution is therefore only short-term. However, a technical contribution with a longer *date of expiry* is the derived motion model discussed in this thesis (cf. section 2.6). More specifically, the way the motion model was derived. The ideas can be extended to a variety of other robot platforms.

For the work on this thesis, a proximity sensor based on BLE was developed. It was successfully applied to the wand-shaped tool in chapter 5, allowing the robot to sense the proximity of the human interaction partner. To the best of my knowledge, BLE has not been used in the context of robotics for a proximity sensor. Related research also shows that it can be used to prototype a touch sensor and that it can help to distinguish between people (e.g. Scheunemann, Dautenhahn, Salem, et al. 2016b). The technology is relatively cheap and easily applicable, which can contribute to faster robot development and a faster design process for robot-related experiments. Instead of developing a whole vision pipeline in order to recognize humans, a researcher can start investigating by using the cheap BLE technology first. The technology can also be applied to service robots already present in human inhabitant environments. For example, one method to increase the functionality of service robots is to distinguish humans and recognize reoccurring visitors. Instead of relying on images to accomplish this, which consumes modeling time and computational power, the robot can rely on the phone signals or a visitor's badge instead. The code for the sensor system is publicly available (see Scheunemann 2018e; Scheunemann 2018a).

1.4.3. Human-robot interaction

This thesis also contributes knowledge to the field of HRI. It has a substantially different approach compared to many other HRI research and focused on the generation of robust HRI. In some HRI studies, human interference could cause a potential risk to the reliability of the system. Certain interactions would need to be explained to the participants prior to the experiment and too much deviation could cause system failures or unwanted results. Instead, the thesis will present a robot which handled interactions *intrinsically*. This allowed for a study design which only needed to provide minimal information to the participants, without the need to control how they interacted with the robot. It also allowed for conditions which could be presented totally independent from the experimenter.

A key contribution to the study design is the systematic way to generate a baseline behavior. The design of a good baseline behavior is very critical and it needs to fulfill two tasks: it needs to be reproducible by other researchers and it needs to be systematically sound. For example, if I had compared an intrinsically motivated robot to a straight driving robot, the results would have had less quality. This thesis extensively discusses the thought process involved and it is believed that the same process can be applied on different and more complex robot platforms, like the one proposed in section 6.4.

The most central contribution to knowledge is that the TiPI formalism, which enables a fully autonomous robot behavior generation, had an impact on human perception of robots. The difference to existing work on IM is that the robot did not mimic IMs, but was in fact truly intrinsically motivated, driven by its *interest* to explore the world through a hysteresis of predictability and change. It is hoped that one of the contributions of this work is an increased motivation of the HRI community in computationally modeled IM. However, the biggest contribution is perhaps that the step toward this research has been taken. I believe that for robots to be situated in our societies and everyday life, they need to be *interested* in mastering new situations. There is an ample amount of evidence that this *robustness* cannot be hard-coded into a robot. Instead, the robot needs to be adaptive toward changes in its environment. So far, this research resulted in two publications, and others are currently being written: (Scheunemann, Salge, and Dautenhahn 2019; Scheunemann, Salge, Polani, et al. 2021).

Last but not least, the thesis presents a first step to understand whether the knowledge of social cognition transfers to physical HRI. Human-human interactions and relationships have been extensively studied over decades and it remains an active research area. If a connection between human-human *interactions* and HRI can be further fostered, this would allow HRI research to use a large set of existing methods and tools to evaluate robot behavior.

The thesis shows that, for an interaction scenario with robots of the same morphology, there is a potential link between our perception of Warmth in humans and our perception of Warmth in robots (RQ1). In both cases we like to sustain the interaction with the agent we perceive highest in Warmth. This evidence gives more weight to the contributions mentioned before: intrinsically motivated autonomy in robots is perceived as warm by human interaction partners, which provides evidence that IM is key to sustain the interaction in HRI (RQ2). A list of scientific publications and disseminations can be found in Appendix A.

Chapter 2.

Background and Developments

This chapter presents the background and developments needed for the work on this thesis. Section 2.1 presents how intrinsic motivation (IM) is understood in psychology and from a computational perspective. Section 2.2 describes the term autonomy. In this thesis, the robot reproduces the behavior of genuine IM in organisms. In other words, the robot *autonomously* generates behavior. Section 2.3 explains the concept behind predictive information (PI) maximization, the formalism used in this thesis for realizing intrinsically motivated autonomy on a robot. The section shows the derivation of update rules for a time-local heuristic of PI. The focus of this section lies in extracting the needed approximations, along with a detailed explanation of what is needed when applying it to a real robot. Section 2.5 then presents the off-the-shelf robot platform Sphero in its BB8 version, along with its sensors and actuators. The application of PI maximization on the robot is briefly discussed. Section 2.6 presents and discusses a motion model to let PI maximization control the robot's servos directly. Section 2.7 then presents a wand-shaped human-robot interaction tool. The tool was used by participants to interact with the robot, since the preliminary study revealed that some participants felt uncomfortable using their hands directly. Furthermore, changes were made to the robot so it can sense the proximity of the tool using Bluetooth Low Energy. Section 2.4 presents measures from social cognition to describe the perception of others. The section links these findings to human-robot interaction (HRI) and shows that the same measures can help to understand human perception of robots.

2.1. Intrinsic motivation

Psychologists studied IM since the 1940s, yielding a large body of work. Both a multitude of approaches for describing the phenomenon and a variety of definitions for IM exist. Hull (1943) introduced the theory of drives to explain IM, which he believed are construed as stressful perturbations (e.g., hunger) an organism tries to reduce. This idea was challenged by White (1959), with the argument that activities based on IMs have different, in particular none-homeostatic, dynamics. Ryan and Deci (2000) defined[1] it as the "doing [...] of an activity for its inherent satisfactions rather than for some separable consequence". An agent

[1]The definition by Ryan and Deci (2000) presents one of the most popular definitions for IM in psychology.

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 and Deci 2000, pg. 56).

This, for example, can be observed in a child interacting with a puppy. The child will likely be motivated to do so, 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 such as training the puppy). Instead, the motivation for the interaction might results purely from wanting to do this activity for its own sake, i.e., the child is intrinsically motivated to play with the puppy.

Concepts like *fun* and *challenge* are presented as crucial for the definition of IM in psychology (ibid.), and yet the literature lacks consensus on what these concepts mean. This triggered the investigation of IM from a bottom-up or developmental perspective. Oudeyer and Kaplan (2008) characterize IM 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.

Generally speaking, a computational model of IM can use any source of information, as long as there is no hard-coded meaning for the sensor input. The sensor input can correspond directly to a physical variable, such as the measured light intensity or the speed of a servo wheel. It can also correspond to a more high-level quantity, such as the number of people in an image. However, the model should derive actions for the robot primarily on the basis of comparing these sensor values, without having any pre-judgment of the corresponding values of the sensors. For example, a robot that uses the information of its energy level can be intrinsically motivated. However, an implementation that makes the robot drive to a power station because the power level is critically low does not constitute intrinsically motivated behavior. I present an example of a computational model of IM later in section 2.3, but first I look at their conceptual relation to HRI research.

In HRI, computational models of IM have not yet been studied. However, there are two HRI studies which use robots in education which elicit curiosity. Curiosity-driven behavior and behavior driven by IM are closely related. In the theory of learning progress, for example, curiosity is a state of experiencing intermediate novelty and complexity, which a person is lead to by their IMs (Oudeyer, Gottlieb, et al. 2016). The theory about the learning process considers curiosity essential for learning, motivating research for education robots in HRI.

One idea behind this is by making robots *appear* curious, that this elicits curiosity in the humans too, which in turn enhances learning and memory retention (ibid.).

Gordon et al. (2015) studied the effect of a robot appearing curious in an educational HRI scenario. A child played a tablet app together with a robot that was portrayed as a "younger peer", to investigate whether the child's learning aim was increased by a curious robot. The robot's behavior was based on a pre-defined protocol, which also included some randomness regarding its questions and statements. While the effects on the children's learning gain were not conclusive, they showed that the robot could elicit curious behavior in children. In particular, they argued that the *curious* robot behavior effected the children's curiosity and not the "engagement or affects of the children toward the robot". Ceha et al. (2019) investigated this in the context of a game played with a curious, social robot and focused on the resulting verbal behavior of the participant. They pre-defined a set of curious questions and revealing statements. They tested their hypothesis in an educational game that taught participants about different kinds of rocks. In a game session, a researcher teleoperated the robot to trigger factual and revealing statements, along with questions showing curiosity. They showed that a robot which displays curious behavior "produce[s] both emotional and behavioral curiosity contagion effects in participants". An interview a week after the experiment showed that participants who played the game with a curious robot had more questions about rocks and wanted to learn more about them. Similar to (Gordon et al. 2015), they could not show an effect on the participants' learning outcome. However, both studies show that a robot could elicit curiosity in the participants.

What remains unstudied in HRI is intrinsically motivated autonomy, i.e., a robot which is intrinsically motivated and seeks these states of curiosity by itself. Will this have an effect on our perception? This is one of the main research questions outlined in this thesis (RQ2).

2.2. Autonomy

The term *autonomy* is overloaded (Boden 2008) and used with ambiguous meanings. For example, when some researchers talk about autonomous robots, they 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 Self-Determination Theory (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 autopoiesis (Maturana and Varela 1991; Froese and Ziemke 2009).

In robots, this is currently a theoretical idea only, but it is often considered necessary for *true* IM. 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 IM on robots are usually put on those robots by humans, and are thus actually extrinsic. Computational models of IM are an attempt to merely reproduce the behavior or functionality of genuine IM in organisms.

From this point onward, when I talk about IM on a robot I will exclusively refer to the initial, technical meaning, the computational model that aims to mimic IM. While the more philosophical underpinnings of autonomy are highly relevant to the larger context of this work and will make this approach useful even if we develop robots with more extensive autonomy, I will set them aside for the present work.

2.3. Predictive information

This section describes predictive information (PI) maximization, the computational model of intrinsic motivation (IM) used for the robot behavior generation in the thesis' experiments. The initial motivation to use this computational model was that it is, by design, computable for continuous sensor input and that it has been applied to physical robotic systems prior to this work.

PI has been described as early as 1986, termed *effective measure complexity* (Grassberger 1986) or *excess entropy* (Crutchfield and Young 1989). Previous work with PI-driven robots in simulation demonstrated its applicability to a large range of different robot morphologies (Der, Güttler, et al. 2008; Martius, Der, et al. 2013b; Zahedi et al. 2013; Martius, Jahn, et al. 2014). A range of existing videos (Martius, Der, et al. 2013c) from experiments in simulation showcase apparent exploratory, playful and open-ended behavior of individual robots and robot collectives. The PI-induced behavior in the videos suggests PI maximization 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 falls 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 and 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 in computing a specific learning rule that aims to maximize the mutual information between a robot's past and future sensor states (Ay, Bertschinger, et al. 2008), i.e., PI quantifies how much information a history of past sensor states contain about future sensor states. More generally, PI is defined as the mutual information between the past and the future of a robot's sensor input. A high amount of PI 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 PI. Alternatively, if there is strong variation in sensor inputs but little structure in the sensory data stream, i.e., the past has little to do with the future, that would also lead to low PI. Vice versa, a high value for PI 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 PI is "the most natural complexity measure for time series" (Bialek et al. 2001; Martius, Der, et al. 2013b).

Der, Güttler, et al. (2008), Ay, Bertschinger, et al. (2008), and Ay, Bernigau, et al. (2012) presented derivation rules for PI, which allows for computing the model directly for linear systems with stationary dynamics. The next subsections present an extension of their work by Martius, Der, et al. (2013b) for the use in nonlinear and non-stationary systems – such as physical robotic systems. The main idea is that instead of computing the full system dynamics, only the system's time-local dynamics are considered to compute the PI values. This quantity is called time-local predictive information (TiPI) and is the one used in this work. Subsection 2.3.1 provides an overview of TiPI and introduces the measure. This is followed by subsection 2.3.2 presenting the derivation of the explicit update rules used for my studies. The derivations are kept short to provide the basic concepts of the quantity and introduce the underlying main approximations and assumptions. Subsection 2.3.3 discusses these approximations and assumptions with respect to applying TiPI in a human-robot interaction (HRI) scenario. Subsection 2.3.4 summarizes this section.

2.3.1. Overview

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, Der, et al. (ibid.). They propose an approximation to compute PI for nonlinear systems with non-stationary 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 τ steps back in the past: time-local predictive information (TiPI). TiPI allows for going beyond discrete finite-state actions, which still dominates scenarios of information-theory-based behavior generation, toward continuous actions. This permits the use in 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, Der, et al. 2013b).

It 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 and Martius 2006).

The approach allows changing the robot's morphology without having to redesign the algorithm, but will still remain 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 toward 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 will 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 will 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.

2.3.2. Deriving update rules

Martius, Der, 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. I do not aim to provide a full mathematical background of the method. For a detailed treatment, the reader should refer to (Ay, Bertschinger, et al. 2008; Martius, Der, 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_{past} and S_{future}, i.e., it measures how much knowledge of the past S_{past} reduces the uncertainty of the future S_{future}. The predictive information, expressed as mutual information, is thus defined as follows:

$$
\begin{aligned}
I(S_{\text{future}}; S_{\text{past}}) &= \left\langle \ln \frac{p(s_{\text{future}}, s_{\text{past}})}{p(s_{\text{future}})p(s_{\text{past}})} \right\rangle \\
&= H(S_{\text{future}}) - H(S_{\text{future}}|S_{\text{past}})
\end{aligned} \tag{2.1}
$$

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

The first essential simplification proposed by Martius, Der, et al. (2013b) is applying the Markov assumption to Equation 2.1. 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 PI in this case reduces to:

$$
\begin{aligned}
I(S_t; S_{t-1}) &= \sum_{s_{t-1} \in S_{t-1}} \sum_{s_t \in S_t} p(s_t, s_{t-1}) \ln\left(\frac{p(s_t, s_{t-1})}{p(s_t)p(s_{t-1})} \right) \\
&= H(S_t) - H(S_t|S_{t-1}) \ .
\end{aligned} \tag{2.2}
$$

In general, the Markov assumption will only hold 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, Der, et al. (2013b) use the reduced Equation 2.2 as the definition of the objective function for deriving the autonomous exploration dynamics.

Above Equation 2.2 is a quantity derived for the whole process. However, to create an actual behavior rule that reacts to 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

$$
I(S_t; S_{t-1}|s_{t-2}) \tag{2.3}
$$

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

$$
\hat{s}_{t-2} = s_{t-2} \tag{2.4}
$$

$$
\hat{s}_{t-1} = \psi(s_{t-2}, \theta_{t-2}) \tag{2.5}
$$

$$
\hat{s}_t = \psi(\psi(s_{t-2}, \theta_{t-2}), \theta_{t-1}) \tag{2.6}
$$

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

$$s_t = \psi(s_{t-1}, \theta_{t-1}) + \xi_t \tag{2.7}$$

ξ_t being the prediction error.

I denote the deviation of the actual dynamics (Equation 2.7) from the deterministic prediction (Equation 2.6) as

$$\delta s_{t'} = s_{t'} - \hat{s}_{t'} \tag{2.8}$$

for any time t' with $t-2 \leq t' \leq t$. 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, δs_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 δs (Equation 2.8) can be linearized as an approximation:

$$\delta s_{t'} = L(s_{t'-1})\delta s_{t'-1} + \xi_{t'} + O(||\xi_t||^2) \tag{2.9}$$

with the Jacobian

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

Assuming that the prediction errors ξ 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, Der, et al. 2013a, sec. A). It is therefore sufficient to concentrate on the error propagation for the computation of the TiPI. This reduces Equation 2.2 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 ξ is white Gaussian, the entropy can be expressed as covariances (Cover and Thomas 2012). The resulting explicit expression of TiPI on δS becomes:

$$I(\delta S_t; \delta S_{t-1} | s_{t-2}) = \frac{1}{2} \ln |\Sigma_t| - \frac{1}{2} \ln |D_t| \tag{2.10}$$

with $\Sigma = \langle \delta s \, \delta s^T \rangle$ as the covariance matrix of the process δS, and $D = \langle \xi \xi^T \rangle$ as the covariance matrix of the prediction error. Note that the PI 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.

I now give the algorithm used to drive a robot's behavior toward increasing TiPI. Martius, Der, et al. (2013b) derive it explicitly for the gradient ascending neural network presented in Equation 2.6. They argue that the prediction error ξ is essentially noise and does not depend on the parameter of the controller, and that therefore the term $\ln|D|$ of Equation 2.10 can be omitted when computing the gradient. Based on Equation 2.10, the resulting gradient step executed at each time t is

$$\Delta\theta_t = \epsilon \frac{\partial I}{\partial \theta} = \epsilon \frac{\partial}{\partial \theta} \ln|\Sigma_t| \tag{2.11}$$

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

Applying Equation 2.9 to above equations results in the explicit gradient step

$$\Delta\theta = \epsilon \left\langle \delta u_t^T \frac{\partial L(s_{t-1})}{\partial \theta} \delta s_{t-1} \right\rangle \tag{2.12}$$

where δs and the auxiliary δu are given as

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

To render $\Delta\theta$ computable the Equation 2.12 is further approximated by applying the *self-averaging property* (this is explained in more detail below) of a stochastic gradient

$$\Delta\theta = \epsilon \, \delta u_t^T \frac{\partial L(s_{t-1})}{\partial \theta} \delta s_{t-1} \tag{2.13}$$

As per (Der, Güttler, et al. 2008; Martius, Der, et al. 2013b), Equation 2.13 is the equation by which the (approximate) TiPI maximization is ultimately implemented. I remark that increasing $|\Sigma|$ corresponds to an increase of the norm of δs. 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.

2.3.3. Considerations for applying to real robots

Martius, Der, et al. (2013b) apply the above maximization of TiPI to simulated robots. As a result, those robots show complex behavior. 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 2.13. Importantly, despite using the same rules, the formalism produces different behavioral regimes of the simulated humanoid, depending on the environment it is exposed to. Along 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 I do in the following.

Markov assumption

This assumption simplifies the definition of the objective function Equation 2.2. More importantly, it renders TiPI (Equation 2.3) 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.

Conditioning on an initial from two states back

To compute PI for nonlinear systems with non-stationary dynamics, the proposed solution is to condition the quantity on an initial state being two steps back in time. I 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 that display variation within the given time window, such as proprioceptive sensors measuring the acceleration or velocity.

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 δS (propagation of errors) is equivalent to the one on the original process S (sensor states). This enables the linearization of the error dynamics Equation 2.9 and eventually, under the same assumptions, the formulation of explicit TiPI expressions (Equation 2.10). The assumption 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 these 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, Salem, et al. 2016b). Such sensors measure the signal strength to an external device that 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.

Applying the self-averaging property for stochastic gradients

Equation 2.13 uses the so called self-averaging property of stochastic gradients, that 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 with a one-shot gradient.

Practically, this makes Equation 2.10 computable, as the density distribution of the gradient is hard to obtain. Martius, Der, et al. (2013b) note that using this property is only exactly valid for a small update rate ϵ when it is driven to zero eventually. Note that the update rate ϵ in my application is quite large to allow for a very fast adaptation process. Martius, Der, et al. (ibid.) 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.

Noise is independent of the controller parameters

To derive the explicit update rules (Equation 2.10), 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 θ. Martius, Der, et al. (ibid.) 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, Bertschinger, et al. 2008; Der and Martius 2012; Martius, Der, et al. 2013b; Martius, Jahn, et al. 2014).

2.3.4. Summary

The TiPI method generates the aforementioned variety of different behavioral patterns for a robot. This makes TiPI-maximization a promising candidate for use in HRI settings. Its universality for different embodiments and non-stationary 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. Completely missing from the existing body of work on TiPI, however, is the actual evaluation of the behavior when it is induced by the interaction with humans. This is the gap this thesis aims to fill. One hypothesis of this thesis is that an intrinsically motivated robot, enabled by TiPI-maximization, is perceived as more social by human participants, and that this will eventually help to sustain human-robot interaction (RQ2). As this has not been investigated before, the design of the studies needs to be carefully considered. With a particular focus on understanding the perception of intrinsically motivated autonomy, it was crucial to isolate the perception from the robot's capability to fulfill a specific task. This is a main contribution of this work and the three studies in chapter 3, 4 and 5 guides through the process of finding a suitable study design to address O2.

2.4. Social cognition

A large body of literature in social psychology connects our social perception, often in forms of stereotypes, to behavior formation (e.g., Judd et al. 2005; Fiske et al. 2007; Cuddy et al. 2007; Abele, Hauke, et al. 2016). When looking into the literature of social cognition to understand the social perception of peers (person perception), a popular approach is to model the complex human stimuli with fundamental dimensions. This section briefly introduces two models for human attitude formation. In both, two dimensions suffice to explain more than 80 % of our perception of others. The last section shows examples of applications in HRI research, and motivates the measure of the dimension *Warmth* in this thesis.

2.4.1. Person perception

Wojciszke, Bazinska, et al. (1998) discussed that moral categories (i.e., stereotypes) are dominant in impression formation of others. Generally speaking, the consideration of stereotypical behavior is mostly connotated negatively in the general public. In the scientific community, however, stereotypes became a powerful tool to understand and predict behavior formation for individuals and groups (e.g., Czopp et al. 2015; Bodenhausen et al. 2012).

One reasoning is that mapping humans' complex stimuli structure quickly into simple categories is an essential survival skill for humans. When two humans meet, their first judgment relates to whether the other human is going to harm them or not. In other words, their *behavioral intent* is judged, whether they are friends or foes (Fiske 2018). The next crucial question is whether they can *enact* that intent or not. The stereotype content model (SCM) labels these two dimensions of judgments as *Warmth* (trustworthiness, sociability) and *Competence* (capability, agency) (e.g., Fiske et al. 2007; Fiske 2018). One important finding when researching the SCM was "that people perceived as warm and competent perceive uniformly positive emotions and behavior, whereas those perceived as lacking warmth and competence elicit uniform negativity" (Fiske et al. 2007). Cuddy et al. (2007) described further how intergroup affects and emotions measured by these stereotypes yield certain behavior[2]. They

[2]The framework for behaviors from intergroup affect and stereotypes (BIAS).

explained that Warmth describes the valence of behavior and Competence describes the intensity.

Another actively researched framework is the dual perspective model (DPM), which concerns the two central dimensions of *Communion* (benevolence, trustworthiness, morality) and *Agency* (competence, assertiveness, decisiveness), which relates to goal-achievement and task functioning. The DPM has been researched for self-concept (Abele, Hauke, et al. 2016) and interpersonal attitudes (Wojciszke, Abele, et al. 2009). There, liking-disliking (Communion) reflects personal preferences, and respect-disrespect (Agency) reflects deference. It has long been acknowledged and later been confirmed that these two dimensions can account for more than 80% of the variance in individual impressions (Wojciszke, Bazinska, et al. 1998; Abele and Wojciszke 2014; Fiske 2018). Both SCM and DPM consider two logically independent cues to describe the complex interplay and attitude formation in humans. The importance and weight of these cues differ between the different perceptions, such as perception of others, perception of groups and self-perception. For example, in self-perception Agency and Competence content receives more weight than Communion or Warmth. Both models, however, show that in the perception of others either Warmth or Communion receives more weight than Competence and Agency (Abele and Wojciszke 2014; Fiske et al. 2007; Carpinella et al. 2017). They further show that humans perceived high in Warmth and Competence perceive uniformly more positive, social behavior of others (e.g., Fiske et al. 2007).

2.4.2. Robot perception

The knowledge of person perception has been used in the context of HRI. For example, stereotypes have been studied in HRI such as facial gender cues (Eyssel and Hegel 2012) or gender and the relationship of personality and gender perception (Tay et al. 2014). Carpinella et al. (2017) designed the Robotic Social Attribute Scale (RoSAS) to transfer the dimensions of Warmth and Competence to the domain of HRI. The scale has been since picked up and further analyzed (see Stroessner 2020). The validation of the scale has been conducted based on still-images of human-like and robotic-like pictures, with gender-specific characteristics. They have been further confirmed on a larger set of images (Mieczkowski et al. 2019).

Oliveira et al. (2019) investigated how the display of high and low Warmth and Competence of robots in a card game affects the emotions of human interaction partners. These emotions were then analyzed depending on the role of the human. The 4-player card game was played competitively between two collaborating human-robot teams. The researchers manipulated the robot's Competence with its skills to master the card game, and they manipulated Warmth with utterances of the robot. Interestingly, they found that the formed stereotypes of the robots remained prevalent in the memory of the participants after they were approached a week after conducting the study. The participants' questionnaire responses suggested that participants preferred to play again with the robot they perceived as the most warm. Each participant played with two robots. If both of them were manipulated to be warm, only then did the participants prefer to play again with the robot they judged highest in Competence. This is evidence that the observation from social cognition, namely that Warmth carries more weight in our judgment of others, may transfer to robot perception.

Contrary to the above findings, Stroessner (2020) provided evidence that Warmth is a "poor predictor for contact desirability". In addition, Mieczkowski et al. (2019) further evaluated the SCM (and its extension BIAS) with the use of still-images. They found evidence that attitude formation toward robots mostly follows the rules from social cognition. They further argued, however, that people judge a robot's Warmth and Competence solely on their physical characteristics, and less on the robot behavior.

The thesis' research question RQ1 is to understand whether dimensions of social cognition offers a good measure for the perception of robot behavior to understand participants' preference. The working hypothesis derived from the first study of chapter 3 is that Warmth, a measure for how social a robot is perceived by human participants, can further serve as a good indicator for the participants' preference to continue interacting with the robot. The final study of chapter 5 addressed this hypothesis and provides evidence that the perception of Warmth on a behavioral level can help to predict participants' intent for future interaction. In the final study, participants did not perceive any of the robots differently on the Competence dimensions. However, the results suggested that participants preferred to interact again with the robot they perceived highest in Warmth. This was subsequently published (see Scheunemann, Cuijpers, et al. 2020). The impact of this evidence is twofold. Firstly, it links the findings of social cognition to robot perception. It shows that robots that are perceived as high in Warmth are more favored by human participants, as is the case with humans that are perceived as high in Warmth. Secondly, this allows focusing research on a single dimension. This is important to the experiments of this thesis because it studies behavior isolated from a specific robot's task fulfillment.

2.5. Robot

This section describes the robot platform used for the interaction studies in this thesis. Figure 2.1a shows the robot: the BB8 version of *Sphero* (Sphero, Inc. 2020a). It was developed by the eponymous company Sphero and resembles the BB8 character from the *Star Wars* movies (Lucasfilm Ltd. 2015).

Subsection 2.5.1 describes the reasons for choosing this specific robot platform, followed by an overview of the technical details of the robot (2.5.2). The last sections describe the software to control the robot (2.5.3) and its motion control (2.5.4).

2.5.1. Choice of robot platform

The robot platform had to meet certain requirements to make it applicable in a human-robot interaction (HRI) scenario and to be able to run an implementation of time-local predictive information (TiPI) maximization. This resulted in the use of a small, spherical and non-humanoid robot platform.

Note that the choice was driven by purely technical standpoints, which are addressed in detail in this subsection. The fact that the robot resembled a movie character does of course raise expectations of its behavior. This is addressed in later study designs with the use of a within-subjects design and the focus on perceptual change rather than on absolute perception. More details will be discussed in chapter 3, 4, and 5.

Purchasable Off-the-shelf

The anticipated time for testing the algorithms and developing the studies could not be accurately predicted, mainly due to the lack of previous work in using an implementation of intrinsic motivation (IM) in an HRI scenario. A rough estimation for implementing and thoroughly testing the study was predicted to be over a year. This of course bears the potential risk that over time the robot's hardware would wear out considerably. The idea behind using an off-the-shelf robot is that prototyping and designing can be done with one robot, while the studies are conducted with another robot of the same batch, using a robot with sufficiently similar characteristics. Another idea was that an off-the-shelf robot may offer a tested software stack to start the study quickly.

Physical robustness

The main reason for choosing a spherical robot was its physical robustness. This was important for three reasons: firstly, the investigation started with very young children with autism (2 to 4 years of age). Secondly, it was not clear *how* the interaction would turn out. Lastly, it was unclear in which environment the HRI investigation would take place. Would it be in casual and everyday places, or rather in a laboratory setting?

Therefore, the robot needed to be quite robust to tailor for a variety of interaction scenarios. The Sphero company showed advertisement videos where the robot was used outside and even in water. Initial tests with the robot showed that it could fall off a table onto a hard surface, and could cope with being kicked without breaking.

Spherical shape

The stable and reliable locomotion of a robot is still a challenging task. For example, humanoid robots cannot yet master stable walking on different terrains or grounds. Additionally, given the novelty of the research in this thesis, it was not expected that a full behavior generation can be realized which can master locomoting with robot morphologies which provide many degrees of freedom. Therefore wheeled robots were considered, as stable locomotion is almost guaranteed. However, they are constrained to even ground in order to locomote. Following the example above, it was unclear what the interaction and study design would eventually look like. In the end a spherical robot was chosen: like with wheeled robots, its locomotion is always stable, but it is also able to locomote in uneven/changing terrain.

Minimality of the platform

The choice above describes the decision for a robot which has as few degrees of freedom (DOF) as possible, but still enables stable and reliable locomotion. This is also one of the arguments for the requirement of *minimality*. Fewer DOF also have more practical benefits. Firstly, it makes purchasing a batch of robots feasible, since fewer DOF means fewer servos, which results in lower costs. Secondly, if it is possible to apply IMs to robots in HRI, it will surely render computable on a robot with minimal hardware. Furthermore, there are disadvantages of using a more complex platform with more DOF. A humanoid platform "might raise false expectations regarding the cognitive and social abilities that the robot cannot fulfill" (Dautenhahn 2004). A humanoid robot could decrease a human participant's interest in the robot when its behavior does not meet their expectations. On the other hand, human participants could get overexcited about a robot that accidentally conducts *gesture-like* motions. For example, they could anticipate a waving gesture, which could give human participants a sense of a *will* to communicate. All this can blur the participants' perception of the robot.

These above reasons lead to the choice of a non-humanoid, even minimal platform, in order to fully concentrate on the impacts the robot's IMs have on the human participant. This way, if the minimal robot elicits any positive human perception, it would give more weight to its behavior generated by its IMs.

Visibility of moving direction

The minimalism mentioned above can also have negative side effects. For example, a fully spherical robot could make it difficult for a human interaction partner to understand its movements. For example, if the robot would simply spin on the spot, it may be hard to perceive. However, the BB8 version of the Sphero robot has a head attached. The head always points in the driving direction, which provided the naïve participant with a sense of direction.

Applicability to use TiPI

A behavior generation controller that uses TiPI maximization needs sensor input that fulfills the requirements outlined in subsection 2.3.2. One core assumption of the derived TiPI formulas is that the prediction errors are both very small and Gaussian. This rules out sensor input which has sudden drops, which is an issue for sensors that have to deal with occlusions, such as camera sensors. Subsection 2.3.3 argues that speed or acceleration sensors are good candidates since they fulfill these requirements. The robot platform offers a range of such sensor readings, such as sensors for linear acceleration, for angular velocities and for its servo speeds.

2.5.2. Technical details

The Sphero robot is a small, spherical robot that has a 74 mm diameter and weighs 168 g. Figure 2.1a depicts what the robot looks like. It has a 75 MHz ARM Cortex M4 on board, powered by two 350 mA h LiPo batteries. The batteries can be charged with power inductive charging. By design, the inner parts of the robot cannot be accessed.

The robot contains a two-wheeled electric vehicle inside its spherical shell, as depicted in the cross-sectional view of the robot in the center of Figure 2.1b. The vehicle and the head

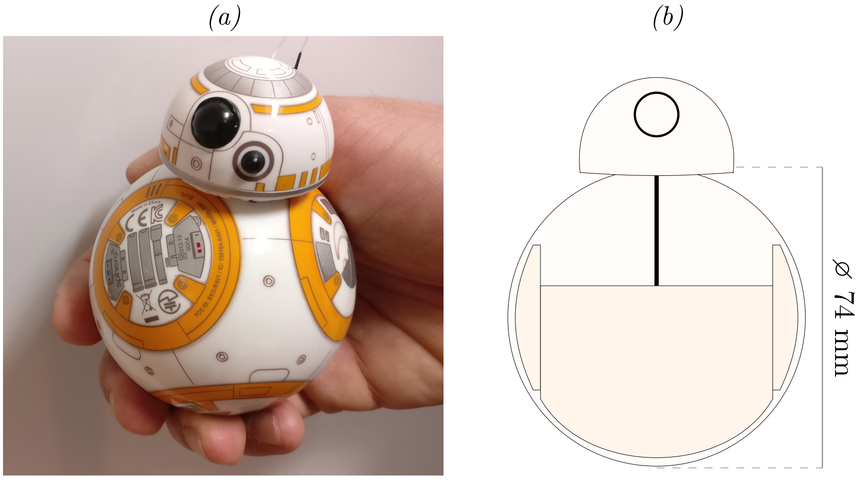

Figure 2.1.: (a) The robot platform BB8 from Sphero. (b) A 2-D cross-sectional view of the robot. A two-wheeled vehicle (darker shape), kept in position by a heavy weight, moves the sphere when driving. The speed of each servo motor can be set individually, allowing the robot to move straight, to turn and to spin. A magnet attached to the vehicle keeps the head on top of the sphere, facing the direction of movement.

both have magnets that attach the head to the outside of the sphere. This also allows the head to be kept in the driving direction, providing any observer with a sense of the robot's direction. There is a coil for power inductive charging on the bottom of the shell, which also acts as a weight to keep the vehicle in place. The two servo motors provide the robot with two DOF, allowing the robot to locomote backward, forward and to spin on the spot.

The robot can stream sensor information. It offers raw sensor readings of a 3-DOF accelerometer (linear acceleration), a 3-DOF gyroscope (angular velocity) and the servo speed (in back EMF). These sensors are fused on-board, delivering further sensor readings such as inertial measurement unit (IMU) data (in quaternions or Euclidean angles), the robot global position on a 2-D map (with the reference frame as the robot's starting position) and the robot speed on a 2-D plane, which is delivered in components for the forward/backward and sideward speed.

2.5.3. Software

The robot came with a mobile phone application that could be used to remotely control the robot either directly or by drawing a path that the robot then follows. It further allows setting the robot in a mode where it could automatically *explore* the environment. Additionally, there was an official API written in the script language `JavaScript`. It used Bluetooth Low Energy (BLE) to send control commands to the robot, and receive sensor readings from the robot. The package in `Node.js` enabled writing custom programs to control the robot.

The API was a good starting point for exploring the robot. However, there were a few reasons which eventually triggered the development of a custom framework. First of all, the development for this thesis started in 2017 and at the time the last commit to the repository dated back to 2016-05-11. Furthermore, necessary pull requests to make the robot operable were not taken care of anymore. I contacted the original API developer to learn that there there is no further support planned. I had submitted four pull requests (PR) to the repository, ranging from corrections to the documentation (see Scheunemann 2017b) to more crucial PRs which made the robot operable for my needs. This includes code to enable connecting and controlling more than one robot (see Scheunemann 2017c), enable direct access to quaternion readings (see Scheunemann 2017d), and, most importantly, fixing issues with parsing packages sent by Sphero (see Scheunemann 2017a). This fix was needed as the robot could get stuck entirely and would make the robot inoperable, especially when the communicated packages were quite large. This happens, in particular, when a lot of sensor readings are streamed, something which was absolutely crucial to enough data for TiPI maximization, but also to log the actual sensor state of the robot for later analysis. In December 2019, the repository was eventually archived and the documentations to controlling the robot using JavaScript disappeared from the company server.

I decided to implement my own API using the programming language `C++` (see Scheunemann 2017e; Scheunemann 2018b). The language leaves implementation details to the developer, which allows programming with little hardware demands. This enables running developed code on embedded systems. This way the study could be carried out independently of a laboratory setting. The first attempts of this thesis involved working with autistic children in a nursery. In this context, the library was successfully used on a *Raspberry Pi 2* during the initial play sessions with children.

It needs to be noted that the robot had issues with the on-board firmware as well. For example, when it sent commands to “set the raw motors”, as mentioned in the API documentation, the speed values for the left and right servos were swapped. More crucially, the sensor readings provided by the built-in IMU were faulty. For example, the quaternions did not provide the true orientation of the robot. This is probably best explained in Euler angles. If the robot was rolled more than 90° (positively or negatively), the absolute angle value decreased again. This had no direct implication for the studies of this thesis, as the robot does not reach such a roll angle by itself. However, all these findings counteracted the idea of an off-the-shelf robot, as it slowed down development and made excessive testing of the robot's sensors necessary.

2.5.4. Robot control

The vehicle inside the robot has two servos and there are two ways of controlling them, either by directly setting the speed for each of the servos or by using a built-in balanced mode: a closed-loop controller that keeps upright while processing speed and heading commands. The two options, i.e., *direct control* and *balanced control* are explained in this section.

Balanced control

BB8 provides a built-in *balancing mode* controller. This controller expects heading and speed information, which it then translates to commands to the robot's servos. The speed here is referring to the forward speed of the robot and the heading is an absolute position of the robot's head. When the robot is started, the current heading is 0°. If the robot receives a command of 20° for its heading, it is set to 20° with respect to the starting position (i.e., global frame). According to the Sphero documentation, a parameter can be set to control the speed for changing the robot's heading direction. However, this parameter did not have an effect on the moving speed of the head. Furthermore, it is documented that the balance controller uses a fusion of accelerometers, gyroscopes and the wheel encoders (Sphero, Inc. 2020b). Other than that, there is no further information on the closed-source controller. What can be observed is that the robot tries primarily to remain upright (i.e., the robot's head stays on top) while trying to reach the requested heading and speed. The robot has several possible motion patterns while always being upright: it can locomote forward, it can turn while moving in the forward direction with various radii, or it can spin on the spot. In the first study in chapter 3, both conditions are based on the balanced mode.

Direct control

It is possible to control the speed and direction of both servos directly and individually. This enables the possibility of having an open-loop control, which means that the servos control is independent of any sensor information. Importantly, the set servo speed and direction are not manipulated further by the robot firmware, allowing the implementation of TiPI maximization to fully control them. In contrast to the balanced control, the direct control allows for a larger variety of motion patterns. It can cover the patterns above, i.e., moving straight forward, turning and even turning on the spot (spinning), but it can also create more wobbly locomotion and moving and turning backward. In the first study (the preliminary study), the direct control was not used. This has much to do with the challenge of applying the algorithm directly, which is discussed in the next section (2.6).

2.6. Motion control

The interaction studies used two different motion controls: the built-in and closed source balanced motion control and the direct motion control. Both have been described earlier in subsection 2.5.4.

The first study used the balancing controller for the robots in all conditions. The main motivation was that it was relatively straight forward to implement intrinsic motivations (IMs) in the robot, which allowed gaining insights into the usage of intrinsically motivated robots and the resulted adaptive and exploratory behavior.

However, there are two main reasons to avoid using the balancing controller. Firstly, the implementation is closed-source, which leaves it open to whether the control eventually fulfills the time-local predictive information (TiPI) requirements outlined in subsection 2.3.3. For example, is there a control delay that spans much further than 2 timesteps in the past? This would mean that the current measured sensor errors are not realized by the TiPI maximization output from the previous step. Secondly, a controller plugged between the TiPI maximization output signals and the robot output may restrict the robot's self-motivation. For example, if the robot is intrinsically motivated to excite its sensor input associated with its accelerometer, it may gradually change the associated output for the speed command. However, the pre-specifying effect – the robot's *drive* – to keep itself upright may put an unobservable limit to the computation of TiPI maximization. This makes it impossible to analyze whether this limit is reached because of the balancing controller or because of TiPI maximization. Consequently, the goal of this thesis to analyze the influence of the robot's IM-driven behavior on human perception cannot be fully understood.

This section derives a simple motion model $\mathfrak{M}$, which allows the robot to directly control its servos. An implementation of the model $\mathfrak{M}$ computes robot actions (i.e., servo commands) from the desired speed (i.e., the output from TiPI maximization). The section guides through the analysis of the robot behavior toward the derivation of a simple motion model based on linear equations. The aim is to provide an example of what can be discovered when applying TiPI to a real robot and how to approach possible issues.

2.6.1. Overview

Subsection 2.6.2 provides a more detailed explanation of the notation and units. Subsection 2.6.3 discusses a first naïve approach for $\mathfrak{M}$: simply take the output values from TiPI maximization and map each of the values directly to one of the two servos. The section shows that this did not yield any interesting behavior. Instead, the robot just alternated between spinning left and right without any locomotion.

Subsection 2.6.4 follows up on that and analyzes the problem qualitatively and quantitatively. Essentially, the independence of the control parameters, which is a core requirement of TiPI maximization, is violated because changing the speed of one motor affects the speed of the other under fixed control. Subsection 2.6.5 derives a simple motion model $\mathfrak{M}$. Subsection 2.6.6 then evaluates the motion model. It shows that the requirements of TiPI maximization are fulfilled to the extent that it results in a highly variable robot behavior. Subsection 2.6.7 summarizes the section.

2.6.2. Notation and units

This section derives a motion model which computes robot control commands (i.e., actions a) from the desired servo speed s. More formally, let $A \subset \mathbb{N}$ be the robot's set of actions and let $S \subset \mathbb{N}$ be the possible desired sensations of the robot, then the motion model for a robot with two motors is

$$\mathfrak{M} : S^2 \to A^2 \tag{2.14}$$

Three variables are needed to explain the robot dynamics and describe the motion model in this section. Let $i \in \{l, r\}$ be the index of the left and the right servo, then

$$\begin{aligned} &a_i \in A, \text{ the action, i.e., set servo speed} \\ &s_i \in S, \text{ the sensation, i.e., measured servo speed} \\ &\hat{s}_i \in S, \text{the desired sensation} \end{aligned}$$

The *actions* or *set servo speed* a denotes the set speed for the left and right servo of Sphero as requested by the Sphero API. The documentation says the set speed unit is "pulse-width modulation (PWM) duty cycles", which range from -2048 to 2047. However, the sensor readings for the duty cycles and the actual motor command are not similar, but linearly correlated. For the sake of a straight forward implementation to a real Sphero robot, this section uses the values for a, which is how they are forwarded to the API, without further specifying the units. A similar argument follows for the *measured* or *actual servo speed* s. Sphero provides sensor readings of the servo speed by measuring the servo's back electromotive force (back EMF)[3]. This value is a voltage appearing between the armature and the magnetic field of the motor's field coil. However, after studying the documentation it is unclear how the values map to the true back EMF, therefore the provided sensor readings are presented without a unit.

Furthermore, there are two constants k, which describe the absolute maximum value of the set speed and the desired speed:

[3] The back EMF is related to what is also known as the counter-electromotive force (counter EMF, CEMF).

$$k_a = 120, \text{ the maximum set speed}$$
$$k_s = 320, \text{ the maximum desired speed}$$

The values are the same for the left and right servo, as well as for forward (positive) and backward (negative) speed, e.g., the minimum set speed is $-k_a = -120$. They both have been empirically derived in such a way, that the robot can operate in the experimental environments of this thesis. For example, the maximum speed was determined so that the robot can locomote freely without losing its head when it bumps into an obstacle.

Note: a is the *signed* set speed value. To set this speed to the Sphero API, a tuple of $(speed, mode)$ is needed. Let f be a function which computes the API *speed* and *mode*, then

$$f(a_i) = \begin{cases} (a_i, \quad 1), & \text{if } a_i > 0 \\ (|a_i|, \; 2), & \text{if } a_i < 0 \\ (0, \quad 0), & \text{otherwise} \end{cases}, \forall i \in \{l, r\}.$$

2.6.3. Naïve motion model

This section describes the first attempt to allow the robot to control its servo speed directly based on its IMs. The idea of the motion model was simply to map the output from TiPI maximization directly to the robot's servo commands.

Let $y \in [-1, 1]^2$ describe the output of an implementation of TiPI maximization, then the desired speed can be scaled as follows:

$$\hat{s} = y \cdot k_s \tag{2.15}$$

Then, the servo commands a can be modeled with $\mathfrak{M}$ as follows:

$$\mathfrak{M}(\hat{\mathfrak{s}}) = \frac{\hat{s}}{k_s} \cdot k_a = y \cdot k_a \tag{2.16}$$

Figure 2.2 shows an example of a robot behavior resulting from using Equation 2.16. The figure shows the measured servo speed of each of the two servos. If one servo is moving backward, while the other moves forward with an almost similar absolute speed (i.e., $s_l = -s_r$) the robot simply just spins on the spot. The blue and red shaded areas indicate a left or right spin of the robot around its z-axis.

The gray shaded area is another typically recognized pattern. The robot there leans forward and backward, as if it tries to move forward or backward. However, it is unable to move, and explores different actions. Despite these three unique behavior patterns, Figure 2.2 shows that the resulting behavior of the robot is alternating between spinning right and left, without locomotion. In other words, the robot's behavior is very repetitive and has little

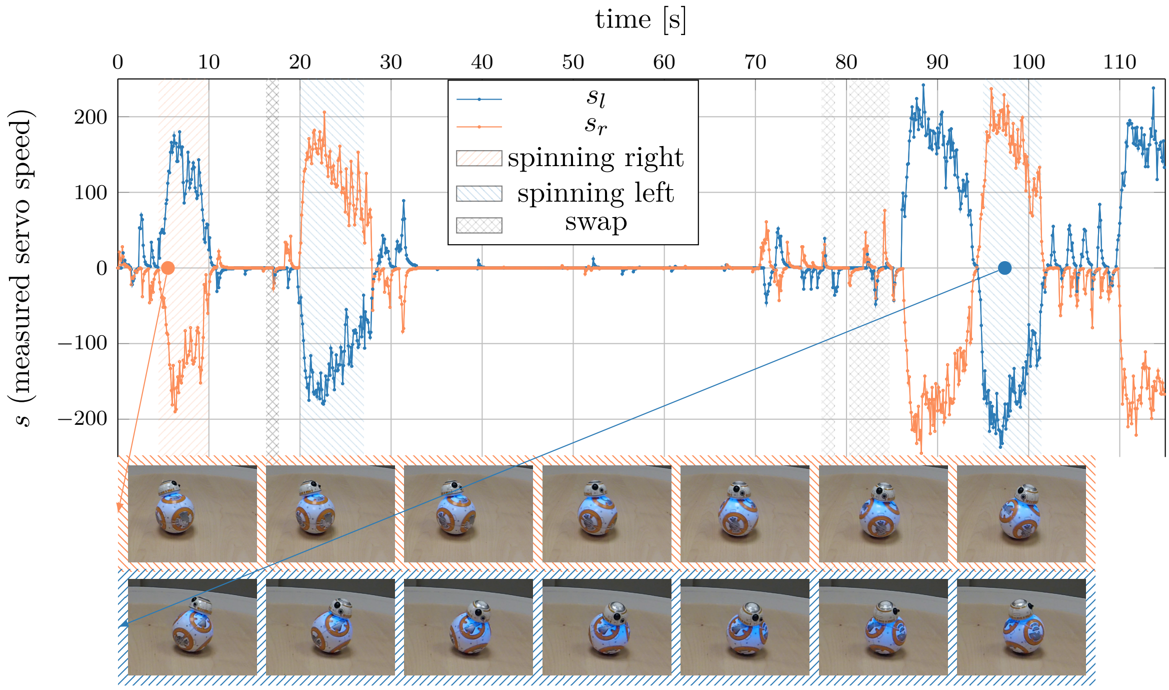

Figure 2.2.: The figure shows the resulting robot behavior if the output of the robot controller is mapped to the robot's speed commands. The robot swaps between spinning left and right at various speed levels, but does not locomote or explore.

variety.

The important question here is whether the robot is not *curious* enough to also explore forward and backward movements. Is there an issue with the TiPI maximization formalism, the controller implementation or is the robot simply not suitable?

A simple experiment confirms that something keeps it from exploring different motion regimes: when the robot was lifted by its head, i.e., the magnetic attachment, so that the main shell did not touch the ground, it started exploring additional straight and turning motions. This shows that the robot behavior depicted in Figure 2.2 is not theoretically limited to only a small set of behaviors, but there are practical limitations.

The central difference between a robot lifted by its head and a robot locomoting on a surface is the additional friction of the surface. The next section investigates the practical limitations in more detail.

2.6.4. Dependency analysis

This section expands on the above argument and analyzes the development of the measured servo speed s_i depending on the set servo commands a.

Data collection

The first method was to build a controller that applies a sine function to each servo, with each function having a slightly different frequency. This way, motor commands in a harmonic fashion would be applied and at the same time, the exploration of different relations between the two set speeds is possible. The issue with that was, however, that the robot would operate for a relatively long period of time with full speed around the extremes of that function. This was an issue for the robot as it then often bumps into obstacles and was likely to loose its head, which in turn changes the dynamics of the robot drastically.

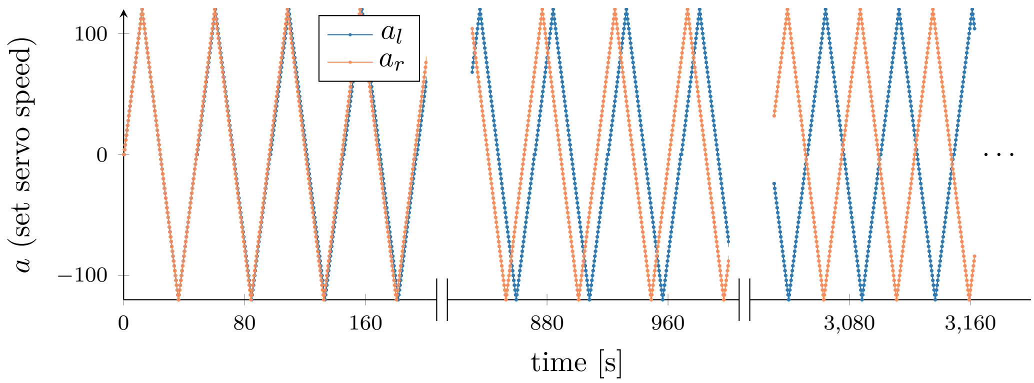

Figure 2.3.: The resulting set motor speeds a_i for $i \in \{l, r\}$ of a simple controller to collect measured speed information.

A controller sends motor commands a to a Sphero robot in constant time steps ($\Delta t = 0.1\,\mathrm{s}$). The controller starts with $a_i = 0$ for both servos ($i \in \{l, r\}$). This means, for time step $t = 0$: $a_i(0) = 0$. After four steps, the command is increased by 4: $a_i(t) - 4 = a_i(t-1) = a_i(t-2) = a_i(t-3) = a_i(t-4)$. This continues until $a_i(t) = k_a$. Instead of further increasing, the servo command is decreased by 4 every fifth time-step until the minimum motor command $-k_a$ is reached. This continues for one full period, i.e., the maximum and the minimum was reached and the motor command is back to 0. Then, a lag is introduced for the motor command a_r. Instead of 4 steps, the servo speed is only increased after 8 steps. The data collection runs for about 45 min, yielding 4 to 12 readings per combination of a_l and a_r. Figure 2.3 shows the resulting control signals for both servos. The controller allows collecting a comprehensive set of control patterns and their effects on the actual motor speed. This spans a wide area of the control/sensory space in a controllable way.

Figure 2.4 shows two scenarios for data collection. In scenario (a), the robot's head is attached and its spherical body can move freely. The only external friction on the system is the head. Scenario (b), on the other hand, exposes the robot to a more realistic scenario. The robot freely locomotes on a circular table, the same table which was used in the studies utilizing the presented model (cf. subsection 4.2.1 and 5.2.1).

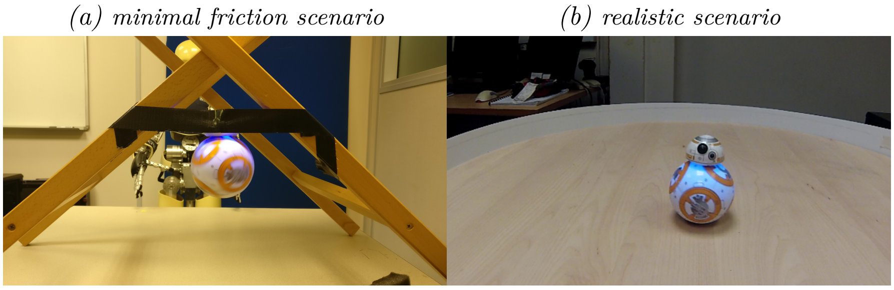

Figure 2.4.: The two scenarios for data collection. In (a) the robot is attached on its head, allowing the shell to freely locomote. In (b) the robot can freely locomote on the table, which was ultimately used for the experiments.

Qualitative analysis

The above data collections allow for an analysis of how the actual servo speed s depends on the set motor speed a. This section discusses this dependency with the help of plots. A quantitative confirmation follows in the next subsection.

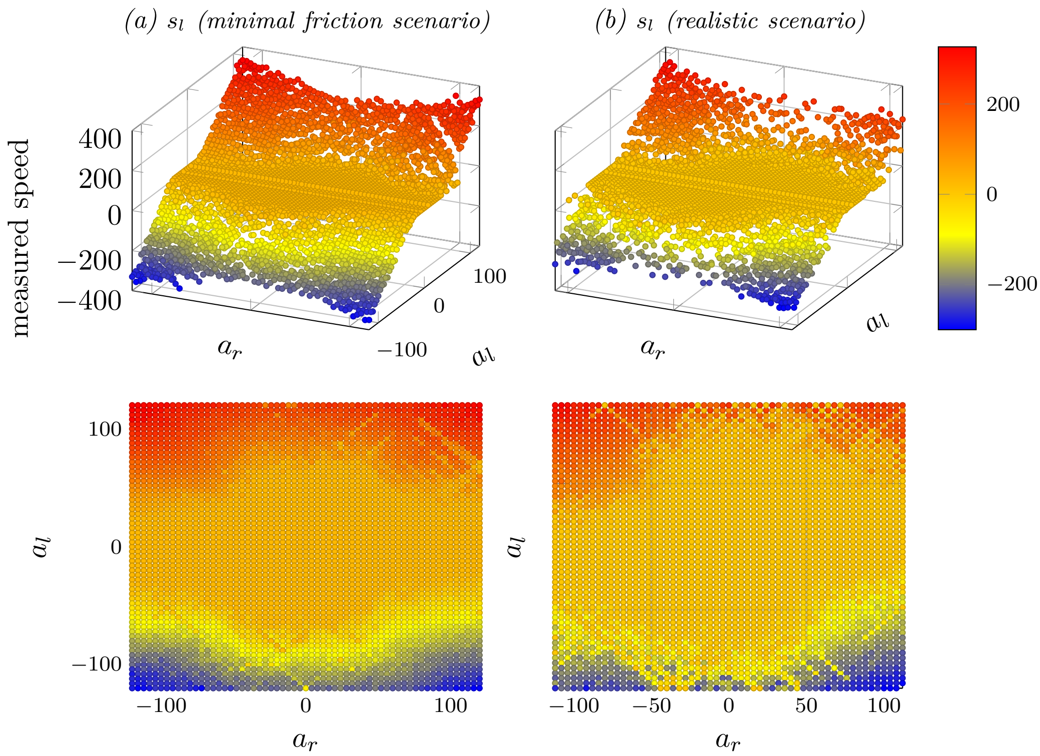

Figure 2.5.: All plots show the median values of the actual speed of the left servo motor (s_l) with respect to the set servo speeds a_l and a_r (first row as 3-D scatter plot, second row as contour plot). The recordings are conducted in an (a) minimal friction scenario (first column) and in a (b) realistic scenario (second column).

Figure 2.5 shows three-dimensional plots, where the measured, actual speed s_l is plotted with respect to both set speed components a. If all readings are plotted, the shape of the data is not easy to grasp. The plot therefore plots the median value per set speed combination.

First of all, in both scenarios, the measured speed s_l decreases the closer the set speed a_r approaches 0. The decreasing appears to be stronger for the realistic scenario (b). In addition, for the realistic scenario, the measured speed is further decreased when the robot attempts to move straight (i.e., $a_l = a_r$). This is seen in the upper right corner and the lower left corner of the contour plots in Figure 2.5b. The figures show that the robot locomoting on the ground in the realistic scenario (b) has a more complex dynamic in comparison to the minimal friction scenario (a). All these observations make it safe to conclude that in either of the two scenarios the actual servo speed s depends on both set servo speeds. In particular, s_l depends on a_r.

Quantitative analysis

This section provides quantitative evidence for the above observations, namely that it is likely that the measured speed s depends on both set speeds a_l and a_r. For the analysis of the dependency, the inverse of the motion model is used:

$$s = \mathfrak{M}^{-1}(a) \tag{2.17}$$

The simplest model can be expressed as a system of linear equations for each servo, which can predict s depending on a as such:

$$\begin{aligned} s_l &= \theta_0^{(l)} + \theta_l^{(l)} a_l + \theta_r^{(l)} a_r = \theta_0 + \theta^{(l)} \cdot a \\ s_r &= \theta_0^{(r)} + \theta_l^{(r)} a_l + \theta_r^{(r)} a_r = \theta_0 + \theta^{(r)} \cdot a \end{aligned} \tag{2.18}$$

The observations above already indicate that there is no full linear dependency. However, for now the simple linear model is used to quantify dependency, and it is refined later.

The first interesting question here is to quantify whether the measured servo speed truly depends on both variables a_l and a_r, or whether one of the variables is enough to explain the measured servo speed. To analyze this, four different models per variable s_l and s_r need to be investigated. For $i \in \{l, r\}$:

$$s_i = \theta_0 + \theta_l a_l + \theta_r a_r \tag{2.19}$$

$$s_i = \theta_0 + \theta_l a_l, \text{ (i.e., } \theta_r = 0) \tag{2.20}$$

$$s_i = \theta_0 + \theta_l a_r, \text{ (i.e., } \theta_l = 0) \tag{2.21}$$

Effect sizes for comparing the model strength are used to analyze which of the above models can predict the dependent variable s best. Table 2.1 reports three effect sizes: the Akaike information criterion (AIC), the rooted mean-square error (RMSE), and R^2. R^2 is the fraction of variance explained by the model. It gives the percentage of how much the model variance can be explained by the variance of the residuals (i.e., estimations)[4] Note that some people argue that R^2 does not have much value as a measure of effect size[5]. However, it is added for completion since some readers might be more familiar with it. For the present data, it provides the same information as AIC and RMSE. AIC is an estimator of out-of-sample prediction error and therefore a goodness-of-fit estimator of a model for a given set of data. The smaller the value for AIC, the better. RMSE provides a quantity showing how much the prediction can differ in units of the dependent variable (i.e., the measured servo speed s).

Table 2.1 shows the analysis of effects[6] for the two dependent variables s_l (a) and s_r (b). It can be seen that the full model (Equation 2.19) shows the best goodness-of-fit (AIC) for the data, has the lowest error (RSME) and describes 63 % of the variance of the data, which is the highest value for R^2. The other models, which use only one of the set speeds to explain s_i, do not produce as good results. This shows that the observations made earlier can be confirmed: the measured speed is indeed dependent on both variables a_l and a_r.

However, when looking at the effects in general it can be seen that they are not very convincing. RMSE can be interpreted as the standard deviation of the residuals or the prediction error. Predicting s wrongly for around 60 units is quite a large error, considering that the maximum absolute value of the measured servo speed is around 320. This means all models poorly describe the data.

Table 2.1.: Linear regression results.

(a) s_l (measured left servo)

Eq.	AIC	RMSE	R^2
2.19	281032	59.3	0.62
2.20	283450	62.2	0.58
2.21	304564	94	0.04

(b) s_r (measured right servo)

Eq.	AIC	RMSE	R^2
2.19	280190	58.3	0.62
2.20	304392	93.7	0.03
2.21	282263	60.8	0.59

Figure 2.6 visualizes the model $\mathfrak{M}^{-1}$. Finding a plane which fits all the values accurately is impossible. More importantly, however, the intercept of the plane does not describe the starting points of true servo movement accurately enough. Applying this model to a robot would cause its behavior to look similar to the initial tests: the robot would only spin right and left.

For a better prediction of the motor speed, the straightforward idea is to describe $\mathfrak{M}^{-1}$ piece-wise per section where the set motor speeds have the same signs. In other words, one part is a linear model for s_l^{++} with only positive predictors a_l and a_r, a linear model for s_l^{+-}

[4]The *adjusted* R^2 explains the same but penalizes the use of more independent variables. The value of the adjusted R^2 for all examples here is the same as the R^2 and is therefore not reported.

[5]An example discussion of why not to use R^2: (Ford 2015).

[6]The models are computed with R's built-in `stats` package and its method `lm()`. For example, Equation 2.19 is computed with `lm(`$s_l \sim a_l * a_r$`)`.

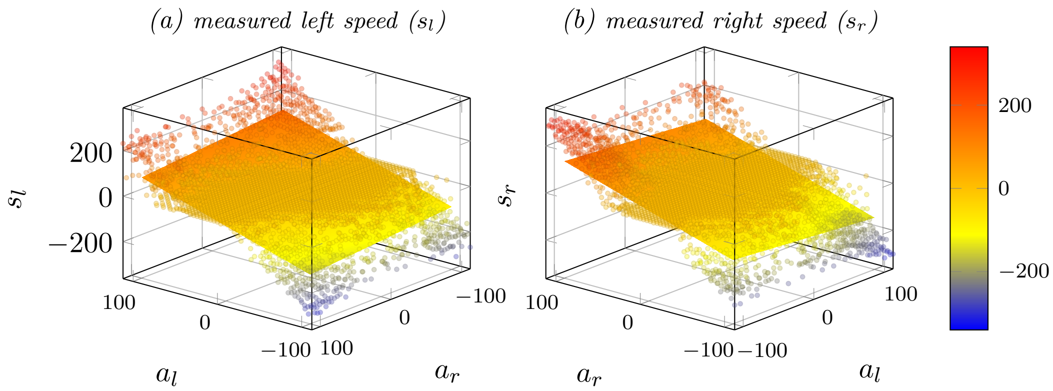

Figure 2.6.: The planes visualize the linear models described by Equation 2.19. The single points are the median values for the measured speed of s_l (a) and s_r (b).

with $a_l > 0$ and $a_r < 0$, and so on.

Table 2.3 shows the results for the sector-specific linear models. As before, it can be seen that s_i can be best described by both set speed values (Equation 2.19). This confirms that s_i is dependent on both motor commands a. The error to predict s_l ranges from 31.5 to 37 when using Equation 2.19. Looking back at Table 2.1, the error for the full model was 59.3, which is significantly higher.

Table 2.3.: Linear regression results based on sectors.

(a) s_l^{-+}

Eq.	AIC	RMSE	R^2
2.19	27102	31.5	0.82
2.20	28896	43.4	0.66
2.21	31793	73.1	0.04

(b) s_l^{++}

Eq.	AIC	RMSE	R^2
2.19	12793	37	0.7
2.20	13003	40.2	0.64
2.21	14293	66.8	0.01

(c) s_l^{--}

Eq.	AIC	RMSE	R^2
2.19	12793	37	0.7
2.20	13003	40.2	0.64
2.21	14293	66.8	0.01

(d) s_l^{+-}

Eq.	AIC	RMSE	R^2
2.19	28256	30.9	0.83
2.20	30316	44	0.65
2.21	33281	73.2	0.04

Figure 2.7 visualizes the models for the left and right measured servo speed. The points on the figure depict the median values of the measured servo speeds. The planes visualize the model $\mathfrak{M}^{-1}$, which fits those points as per the four sectors. When comparing the visualization to the one of Figure 2.6, where only one linear equation is used to fit all these points, it can be seen that the approach to split the model per sector is already increasing the accuracy.

Table 2.5 shows the coefficients per sector, which describes the planes visualized in Figure 2.7. To confirm the necessity of each coefficient, the model is tested with the reported value for the coefficients. The t-test is used with the null-hypothesis that the modeled values

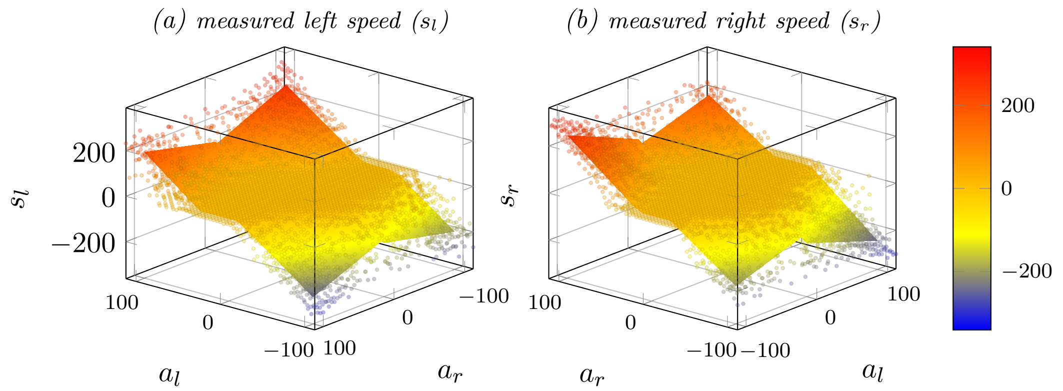

Figure 2.7.: The planes visualize the motion model constructed out of four linear models (Equation 2.19). Each of the models describes the data in one sector, where the signs of a_l and a_r are constant. The single points are the medians of measured servo speeds for s_l (a) and s_r (b).

do not differ when the coefficient is 0. Table 2.5 presents the test results, which show that the null can be rejected for all the coefficients of the model. This confirms that each coefficient makes a statically significant difference.

Table 2.5.: Sector-specific linear regression coefficients for Equation 2.19.

(a) $-+$

	estimate	std. error	t value	p value
$\theta^{(l)}$_0	163.83	3.065	53.4	$< .001$
$\theta^{(l)}$_l	2.73	0.025	110.6	$< .001$
$\theta^{(l)}$_r	-1.34	0.027	-50.2	$< .001$
$\theta^{(r)}$_0	-173.17	2.997	-57.8	$< .001$
$\theta^{(r)}$_l	-1.32	0.024	-54.8	$< .001$
$\theta^{(r)}$_r	2.86	0.026	109.6	$< .001$

(b) $++$

	estimate	std. error	t value	p value
$\theta^{(l)}$_0	-176.46	6.729	-26.2	$< .001$
$\theta^{(l)}$_l	2.95	0.055	53.5	$< .001$
$\theta^{(l)}$_r	0.75	0.049	15.2	$< .001$
$\theta^{(r)}$_0	-154.51	6.475	-23.9	$< .001$
$\theta^{(r)}$_l	1.01	0.053	19	$< .001$
$\theta^{(r)}$_r	2.37	0.048	49.9	$< .001$

(c) $--$

	estimate	std. error	t value	p value
$\theta^{(l)}$_0	149.69	5.297	28.3	$< .001$
$\theta^{(l)}$_l	2.76	0.044	63.3	$< .001$
$\theta^{(l)}$_r	0.69	0.042	16.6	$< .001$
$\theta^{(r)}$_0	169.83	5.146	33	$< .001$
$\theta^{(r)}$_l	1.09	0.042	25.7	$< .001$
$\theta^{(r)}$_r	2.63	0.04	65.1	$< .001$

(d) $+-$

	estimate	std. error	t value	p value
$\theta^{(l)}$_0	-165.11	2.992	-55.2	$< .001$
$\theta^{(l)}$_l	2.7	0.023	115.9	$< .001$
$\theta^{(l)}$_r	-1.44	0.026	-54.7	$< .001$
$\theta^{(r)}$_0	166.88	3.004	55.6	$< .001$
$\theta^{(r)}$_l	-1.05	0.023	-45	$< .001$
$\theta^{(r)}$_r	3.03	0.026	114.7	$< .001$

2.6.5. Deriving a motion model

This subsection derives a motion model $a = \mathfrak{M}(\hat{s})$ to allow the robot to directly control its desired servo speed $\hat{s}$. Instead of fitting $\mathfrak{M}$ directly, I use the coefficients (Table 2.5) from the inverted model $\mathfrak{M}^{-1}$ (Equation 2.19) to compute $\mathfrak{M}$. This way, it can be directly analyzed whether the above observations were sufficient and the resulting model $\mathfrak{M}$ yields interesting behavior.

When $\hat{s}_l$ and $\theta^{(l)}$ are given, a_l can be computed as such:

$$\begin{aligned}\hat{s}_l &= \theta_0^{(l)} + \theta_l^{(l)} a_l + \theta_r^{(l)} a_r \text{ (from EQ 6)}\\ -\theta_l^{(l)} a_l &= \theta_0^{(l)} + \theta_r^{(l)} a_r - \hat{s}_l\\ a_l &= \frac{\hat{s}_l - \theta_0^{(l)} - \theta_r^{(l)} a_r}{\theta_l^{(l)}}\end{aligned} \tag{2.22}$$

Equivalently, starting from the equation for s_r:

$$\begin{aligned}\hat{s}_r &= \theta_0^{(r)} + \theta_l^{(r)} a_l + \theta_r^{(r)} a_r \text{ (from EQ 6)}\\ a_l &= \frac{\hat{s}_r - \theta_0^{(r)} - \theta_r^{(r)} a_r}{\theta_l^{(r)}}\end{aligned} \tag{2.23}$$

As the analysis above shows, the equations of the determined linear system of Equation 2.18 are defined. This allows to linearly combine the two linear equations to express the model for a. Using 2.22 and 2.23 to express a_r:

$$\begin{aligned}\frac{\hat{s}_l - \theta_0^{(l)} - \theta_r^{(l)} a_r}{\theta_l^{(l)}} &= \frac{\hat{s}_r - \theta_0^{(r)} - \theta_r^{(r)} a_r}{\theta_l^{(r)}}\\ a_r &= \frac{\theta_l^{(l)}\left(\hat{s}_r - \theta_0^{(r)}\right) - \theta_l^{(r)}\left(\hat{s}_l - \theta_0^{(l)}\right)}{\theta_l^{(l)}\theta_r^{(r)} - \theta_r^{(l)}\theta_l^{(r)}}\end{aligned} \tag{2.24}$$

Equivalently, using 2.23 and 2.22 to express a_l:

$$a_l = \frac{\theta_r^{(r)}\left(\hat{s}_l - \theta_0^{(l)}\right) - \theta_r^{(l)}\left(\hat{s}_r - \theta_0^{(r)}\right)}{\theta_l^{(l)}\theta_r^{(r)} - \theta_r^{(l)}\theta_l^{(r)}} \tag{2.25}$$

The above linear system of Equation 2.24 and 2.25 together with the computed parameters θ (Table 2.5) describe the direct motion model $\mathfrak{M}$ used in the studies of this thesis. The model can compute robot motion commands a out of the desired motion speed $\hat{s}$. The output $y \in [-1,\ 1]^2$ from TiPI maximization is intended to change the actual servo speed and can serve as the input to the model. For this, y needs to be mapped to the unit space of $\hat{s}$. With k_s which describes the maximal desired absolute speed, then

$$\hat{s} = y \cdot k_s \tag{2.26}$$

A sensible option for the scaling factor k_s is the overall measurable speed of an operable robot in its current environment. This means, that the scaling should be the maximum speed at which the robot can safely operate. For the interaction study environments of this thesis, for example, it has been found empirically that $k_s = 320$ provides good results. The model can now be used to map the output $y \in [-1,\ 1]^2$ (from TiPI maximization) to the necessary motor commands a. The next section evaluates if an implementation of this motion model is sufficient to yield exploratory robot behavior.

2.6.6. Evaluation

This section evaluates the motion model $\mathfrak{M}$ and whether its application results in a more interesting robot behavior with much more behavioral variety. For this purpose, a robot that generates its behavior based on TiPI maximization and the above motion model was placed on the circular table. Similar to the interaction scenarios of, e.g., the second study, a human nudged the robot from time to time. The robot locomoted on the table for 160 s while it was recording its control signals y, the commands a and the measured actual servo speed s.

The question is whether the output for one of the servos, for example s_l, can be explained independently by only one component of y, for example y_l. If that is the case, then the TiPI requirement that the output error is independent of the TiPI controller parameters is given. Consequently, this should result in a robot behavior with much more variety.

The motion model $\mathfrak{M}$ uses the TiPI output y and the parameter θ to compute servo commands for the robot. Similar to the derivation of the motion model $\mathfrak{M}^{-1}$ above, four linear models for both servos are analyzed to answer the question whether the measured speed of each servo can be explained best by one or two components of the control values y. For each servo $i \in \{l, r\}$:

$$s_i = \theta_0 + \theta_l y_l + \theta_r y_r \tag{2.27}$$

$$s_i = \theta_0 + \theta_l y_l + \theta_r y_r + \theta_{l:r} y_l y_r \tag{2.28}$$

$$s_i = \theta_0 + \theta_l y_l \tag{2.29}$$

$$s_i = \theta_0 + \theta_l y_r \tag{2.30}$$

Table 2.7 shows the results for fitting these four linear regression models above. Table 2.7a shows that s_l is best predicted with the equations which contain y_l, i.e., Equation 2.27, 2.28 and 2.29. They all have a comparably low prediction error of around 43, and the goodness-of-fit AIC is very similar ~2340. When observing the results for Equation 2.30 on the other hand, the goodness-of-fit is much lower, and the prediction error is almost three times as high. The equivalent is true for the results from predicting s_r (cf. Table 2.7b). The results are similar, independent of whether both predictors y_l and y_r or only y_r is considered. This shows that a robot which uses the derived motion model $\mathfrak{M}$ can control the desired servo speed almost independently.

Table 2.7.: Effects of predicting s from y.

(a) s_l

Eq.	AIC	RMSE	R^2
2.27	2338	42.9	0.89
2.28	2340	42.9	0.89
2.29	2337	43	0.89
2.30	2821	126.1	0.01

(b) s_r

Eq.	AIC	RMSE	R^2
2.27	2243	34.7	0.89
2.28	2244	34.6	0.89
2.29	2745	106.4	0
2.30	2252	35.6	0.89

Figure 2.8 visualizes the results for (a) the left measured servo speed s_l and (b) the right measured servo speed s_r. What can be seen is that for the case of predicting the left servo speed s_l, the most explanation is provided by y_l. If y_l changes, the measured servo speed s_l changes too. The influence of y_r, on the other hand, is very small and seems to only play a role for very small y_r. The same is true when observing the plot for s_r (cf. Figure 2.8b): the most influence on the actual servo speed s_r is provided by y_r. The variable y_l, on the other hand, does not explain much about the actual speed. Figure 2.8 supports the quantitative analysis in the paragraph above: the actual servo speed s is only dependent on one of the components of y. In other words, the motion model allows the robot to independently control each servo's speed. In theory, the motion model should fulfill the requirements for using TiPI maximization, which should increase the variety of the robot's behavior.

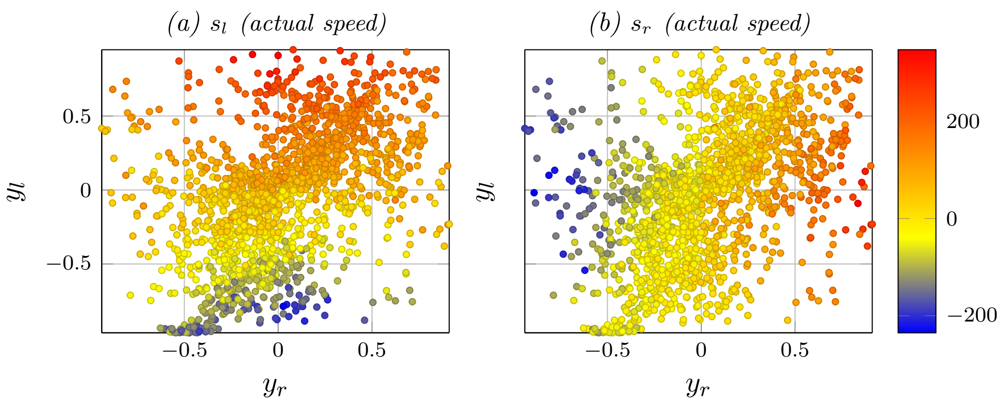

Figure 2.8.: The plot shows the measured servo speed for (a) the left and (b) the right servo with respect to the control signals y. It can be seen that the biggest influence on the measured speed comes from exactly one of the control signals.

Figure 2.9 shows a snapshot of the resulting measured speed values. The robot tries to explore the relationship between motor signals in much more variety when compared to the initial behavior shown in Figure 2.2 at the beginning of this section. At that point, the robot was only spinning left and right. With the use of the motion model $\mathfrak{M}$, the robot now explores other means of locomotion, including forward locomotion or simple turns. This means the robot's resulting behavior is indeed more variable.

It is most likely possible to achieve even better predictability with a different model that uses, e.g., a combination of different kernels or a multi-layered neural network. This, however, is not pursued further here, since the goal was not to accurately predict the servo speed. Instead, the goal was to enable the robot to control its two servos independently, so it can explore the space and try to predict the outcome itself.

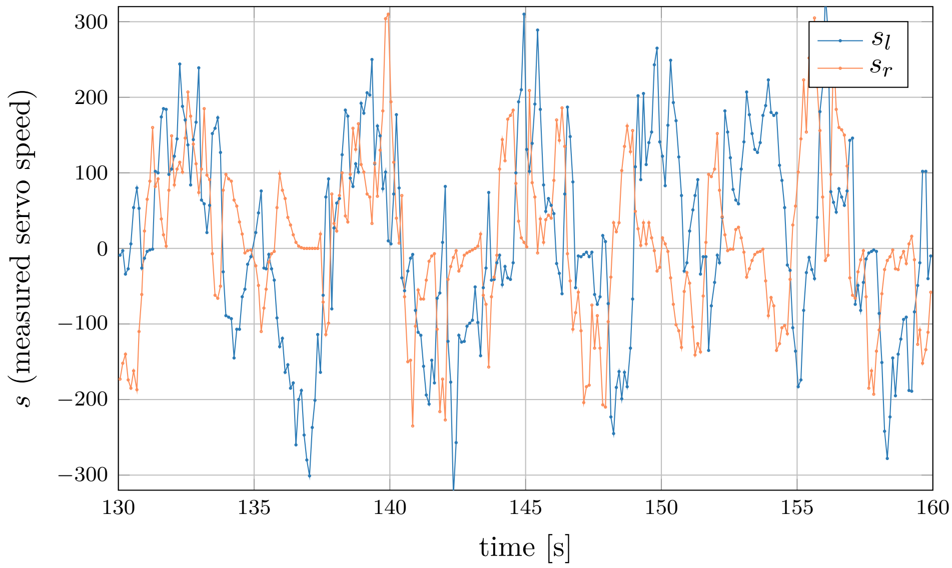

Figure 2.9.: A snapshot of the analyzed experiment. In comparison to the behavior shown at the start of the section (cf. Figure 2.2), the robot now has a higher behavioral variety, including spinning, but also turning and moving straight.

2.6.7. Summary

This section showed that the measured servo speed for each of the two servos of the robot platform Sphero is dependent on both speed commands of both servos. When mapping the resulting TiPI maximization control signals to the robot speed commands, it resulted in repetitive behavior with only little behavioral variety. This was because a crucial requirement of the used formalism to maximize TiPI was not met: the control signals were not independent of the control parameters.

The evaluation shows that the developed motion model allows the robot to successfully directly control its two servos, which, at the same time, increases its behavioral variety based on its IMs.

2.7. Human-robot interaction tool with proximity information

This section describes the wand-shaped human-robot interaction (HRI) tool depicted in Figure 2.10, which was used in the second study and the final study. The tool is a lightweight aluminum tube with a ping-pong ball attached to its tip, and a fabric wristband to its bottom to provide a better grip. The tool served two purposes: (i) it motivated interaction with the robot and (ii) it helped the robot to sense human proximity.

Figure 2.10.: The picture shows the human-robot interaction tool used in the second study and the final study.

The idea for point (i) derived from observations in the first study. There, participants used their hands to interact with the robot. They did not feel very comfortable with this and only interacted with the robot when the situation required it. Providing participants with the wand-shaped tool was meant to mediate this. It was hypothesized that the presence of the tool would already encourage interactions: firstly, because there was a tool in their hand with the introduced purpose to interact with the robot, and, secondly, because the tool allowed interactions without the need to physically touch the robot, even providing a distance between the robot and the human.

The second idea of introducing the tool was because it enabled the means of proximity sensing (point ii). Without the tool, it would be difficult to design an experiment where participants interact in a predictable way, but without giving them precise guidance on how to interact with the robot. For example, some participants may prefer to use their right or left hand, or alternate between them. If the robot could only sense one hand, maybe because of a wristband, the experimenter would have to intervene and ask the participants to only use a specific hand. This would give the participants room to assume the robot's sensing capabilities or imply the robot's goal. This is something the study design of this thesis actively tried to avoid.

The tool structures the experiment and the expected interactions since it was the presented way to interact with the robot. This structure allowed the robot to measure the distance to the very same object, and not, e.g., to the hands of different participants. This section discusses how proximity sensing to humans was enabled.

The question was, however, what kind of technology is best suited for proximity sensing? Subsection 2.7.2 provides an overview of known technologies for proximity sensing and discusses their applicability for the studies of this thesis. It proposes Bluetooth Low Energy (BLE) as a promising candidate for proximity sensing. Subsection 2.7.3 then presents BLE components that were used to develop a proximity sensor based on BLE. A proof-of-concept evaluation of the overall setup is briefly discussed. Subsection 2.7.4 presents how the wand-shaped tool and the robot were modified so the sensor could be used for the final study of this thesis. Subsection 2.7.5 summarizes the section.

2.7.1. Requirements

This section outlines the requirements needed to enable a robot to sense the proximity of human participants in the interactions studies. Firstly, the sensor should be unobtrusive, meaning it should not be obvious to the participant that the robot has the capability to sense proximity. This enables creating studies without biasing the participant for the proximity sensor. If the participants were aware of the robot's capabilities, they may have expected a certain competence or implied goal for the robot. This is in line with the attempts in all the interaction studies of this thesis, where the participants' expectations were kept to a minimum. Furthermore, the system needed to be applicable to a spherical, mobile robot, as the one described in section 2.5. In particular, the robot does not have the possibility to equip additional hardware. In addition, if the sensor output is used as an input for the controller maximizing time-local predictive information (TiPI), the values needed to be provided fast and should not have sudden drops to fulfill the assumptions of the TiPI implementation. And lastly, it was ideal that the sensor system was affordable. This enabled using it in the everyday environment and made replacing it easy. The primary needs can be summed up as follows:

- sense the proximity of a human
- provide sensor readings fast to compute predictive information (PI)
- unobtrusive
- applicability to a small, spherical robot
- self-contained sensor system independent of the robot
- affordable components

While investigating suitable sensor systems, an eye was kept on possible requirements of future studies. For example, it was assumed that the same system can also distinguish between participants. Understanding direct contact (i.e., touch) can also enhance the robot's behavior or the data analysis of the recorded sensor values of the robot. The secondary needs, i.e., the needs hypothesized to be important for future studies beyond the thesis, were therefore:

- recognize touch
- distinguish between multiple humans

2.7.2. Technologies to sense proximity

This section discusses a selection of popular sensor systems for retrieving proximity information between two agents. Each subsection outlines the benefits and shortcomings in relation to the requirements provided above. The section concludes with a summary and proposes BLE as a potential candidate for an unobtrusive and cheap proximity sensor for a mobile, spherical robot, with little configuration needs.

Laser scanners Laser range sensors have a long history of being used for robotic systems, as the sensors provide distance information to surrounding objects with already little computational complexity and few environmental constraints. A common application is to track multiple humans (Schulz et al. 2003) and follow one or more humans in the environment (Lee et al. 2006; Leigh et al. 2015). One of their major advantages is that the sensors can be used in an indoor as well as an outdoor setting, in contrast to, for example, camera systems which is explained later.

However, additional tracking is needed for the robot's self-motion to distinguish between obstacles, such as a wall or an interacting human. With wheeled robots, for example, this is relatively easily applicable because the direction and the position of the sensor are known. However, the spherical robot platform used in the studies of this thesis would make such modeling much more challenging. In addition, these sensors cannot be mounted on the robot without an additional processing unit, which would make the components too heavy and clunky for the small, spherical robot.

External camera systems External cameras for capturing motion are popular in many domains, such as surveillance systems, but they can also be used to track robots (Michel et al. 2006), or enable robots to track interacting humans or their limbs, e.g., their hands (Calinon et al. 2010).

External tracking systems can provide the ability to track robot positions with an accuracy of less than a millimeter[7], by using high-frequency cameras and special markers that mark the object/person/limb being tracked. This provides 2-D positional information or distance information between the human and a robot (Khoramshahi et al. 2016). Such tracking systems enable research despite limited robot sensing capabilities. For example, Khoramshahi et al. (ibid.) investigated gaze cues of a simulated robot looking at a human participant's hand holding a marker.

External camera systems were not applicable to the studies of this thesis for several reasons. If they are used with markers (e.g., providing a glove with markers), then it would become obtrusive and the participants could become aware of the sensing capabilities of the robot (something I wanted to avoid). If only external cameras were used without the markers, the

[7] The company OptiTrack claims that it can track robots with less than 0.3 mm positional error and less than 0.05° rotational error (OptiTrack 2020).

computational complexity of retrieving proximity information would be high. Most likely, some kind of configuration per human would be needed to infer data. Both approaches would need quite a substantial number of cameras to avoid occlusion and, more importantly, the approach would constrain the application area for the study to a room with either a tracking system with markers or one which allowed for ideal lighting conditions. The studies of this thesis all take place in the same room. However, with an eye on future studies, an external camera system would render results incomparable to future studies conducted elsewhere, such as other laboratories or communal spaces.

On-board camera On-board cameras are less obtrusive compared to tracking systems. They are very popular devices for collecting information about humans in the environment. For instance, they can help to understand a human's gaze or facial expressions. Camera images can also help to distinguish objects and/or the distance to those objects (especially if they provide depth information).

However, current approaches, such as using deep learning for object recognition, still needs substantial computational power and many approaches are prone to lightning conditions (Feng et al. 2018). There are approaches that can deal with various lighting conditions (e.g., Dijk and Scheunemann 2019). However, they need knowledge about the environment to infer the distance to objects. In addition, a model of the self-motion (i.e., odometry) is needed to handle occlusions (e.g., Feng et al. 2018).

Vision approaches feel most *natural* and will presumably be implemented in everyday robots eventually, given their biological plausibility when observing humans. Camera devices are also comparably cheap and can be attached to many different robot platforms. However, for the robot platform in the studies of this thesis, a camera was not applicable for various reasons. Firstly, the computational power to compute distance information out of image data was not given, especially not the one needed to achieve that in 100 ms. Secondly, there was no possibility to mount additional hardware on the robot, and, most importantly, the self-motion of the robot was very complex. This makes a reliable odometry calculation difficult, which would have made it very challenging to give meaning to received images.

Radio-frequency identification Radio-frequency identification (RFID) uses electromagnetic fields to identify tags that passively or actively emit frequencies. The most widely deployed RFID tags are off-the-shelf, narrow band tags. They do not need an external power source, they are small, and thus can be applied unobtrusively to everyday objects, which made it a good candidate and was investigated in more detail. Wood et al. (2017) used this technology in an HRI scenario. A robot was equipped with an RFID reader. The human robot operator could then change the robot behavior depending on the RFID tag presented to the robot. This is similar to the way RFID is most predominantly used: granting door access via RFID cards. The human, however, needs to be aware of this technology and control.

A prominent use case of RFID is indoor localization or distance measuring. Localization techniques usually use many tags in an environment or several readers to infer a position accurately (e.g., Truijens et al. 2014; Martinelli 2015). For example, Megalou et al. (2019) presented a robot that navigates in a room using *self-localization and mapping*. Passive reference RFID tags, which were augmenting the environment, helped to infer the position with the help of received signal strength (RSS) of other passive tags which had an unknown location. The approach, however, could not be transferred to the studies presented here. This has several reasons: the accuracy error was around a multiple of 10 cm, which is good for many applications, but not for a robot locomoting on a table with 91 cm in diameter. Furthermore, the environment would need a dense distribution of additional RFID tags which, again, would make the robot less autonomous and constrain future studies to a specific environment. Most notably, it is known that RSS is prone to the antenna orientation of the reader (Martinelli 2015). In the example above, the tags and the reader were allocated in a plane in order to achieve 2-D positioning. This was something the robot in this study would not have been able to deliver.

Ma et al. (2017) proposed a solution for 3-D localization with sub-centimeter accuracy, using only passive, off-the-shelf RFID tags and ultra-high frequency readers, and by exploiting the time-of-flight of the signals. The setup, however, needed, at least, two high-power consuming readers (in different locations). This, again, made it inapplicable for the small robot presented here.

Bluetooth and Bluetooth Low Energy Bluetooth was presented in 1994 to replace wired data connections by using radio transmissions. It became very popular with the core specification Bluetooth 2.1, a standard also known as Bluetooth Basic Rate/Enhanced Data Rate (BR/EDR). It offers a way to pair with peripheral devices (e.g., headphones or mobile phones) and to stream data.

Similar to the RFID technology, some research suggested the use of RSS to retrieve inferring proximity to people wearing Bluetooth devices. For example, this information could then be used to infer social networks (Do and Gatica-Perez 2011), analyze interaction patterns with augmented objects (Siegemund and Florkemeier 2003), change public displays to sustain interaction (José et al. 2008) or for indoor localization (Subhan et al. 2011). These examples made Bluetooth a promising choice. However, this technology has significant shortcomings. For example, scanning for other Bluetooth devices can take up to 10.24 s, making it inapplicable in highly dynamic environments, like the one presented here where a human interacted with a locomoting robot. In addition, the devices have a high energy consumption, due to the protocol's focus on communication between devices which requires large message payloads. This makes it unsuitable for embedded systems or to equip human participants unobtrusively as a larger external power source would be needed.

In 2011, the core specification Bluetooth 4.0 was introduced Bluetooth SIG 2016. Its subsystem BLE, also referred to as Bluetooth Smart, addresses the shortcomings mentioned above. BLE uses small 2 MHz bands over the unlicensed 2.4 GHz radio band. Only three channels are used for advertising[8] messages. These are chosen in such a way that collisions with the most commonly used WiFi channels are unlikely, since WiFi also uses the 2.4 GHz radio band (Faragher and Harle 2015). In contrast to older Bluetooth protocols (since 2.1 is referred to as Classic Bluetooth), BLE uses very short duration messages with small payloads, yielding low power consumption (ibid.). This allows devices that are operating on a coin cell battery with 250 mAh to last 1 to 2 years (Bluetooth SIG 2016; Rault et al. 2014). All this decreases the maximum scanning time from 10.24 s (Bluetooth Classic) to less than 10 ms (Faragher and Harle 2015), which allows a scanning device to see rapid changes in proximity.

Furthermore, BLE devices are small enough to attach them to people unobtrusively. RSS indications between devices are built-in the BLE standard. This makes exploiting the RSS easily applicable. Although the RSS and the distance between devices are associated, there are shortcomings when fully relying on RSS readings.

The biggest disadvantages for human-inhabited environments are partial occlusions (Schwarz et al. 2015; Ahmad et al. 2019) and fluctuations of the RSS (Faragher and Harle 2015). One way to deal with this is by using multiple BLE beacons. For example, Schwarz et al. (2015) presented a robot that could use BLE to find a key ring. The research was conducted in a household environment equipped with stationary BLE devices. A robot had the task to find a key ring, which was equipped with a BLE beacon. The robot achieved this by triangulating its position using three measures of RSS between the stationary devices and the beacon. This technique can be extended by using a pre-processed map. For example, a mobile receiving device collects RSS of the surrounding static sensors. A map of fingerprints is computed, where each cell contains RSS data to all surrounding sensors. With this pre-processed map, an indoor position of a mobile sensor can then be inferred by comparing the stored fingerprint with the current fingerprint (e.g., Faragher and Harle 2015; Ng et al. 2019). Such an approach handles the fluctuations observed in the RSS and makes localization more accurate than established WiFi localization (Faragher and Harle 2015). However, similar to the arguments for other technologies, it is not ideal to constrain future experiments to a specific environment, e.g., an environment with preset BLE beacons and or scanners.

Research suggests that for distances smaller than 25 cm, the actual distance and the RSS are correlated (ibid.). The chosen setup kept this in mind, mounting a BLE device to the robot's head and to the tip of the wand. This allowed the robot to infer the proximity of the approaching wand, and react and adapt accordingly.

Summary The small, spherical robot platform used in the interaction studies of this thesis posed various challenges that only BLE seemed to overcome. BLE correlates to the true

[8] "Advertising" is the term used in the BLE context when a message is sent without knowing who will receive it and if somebody receives it. This is similar to the term broadcasting known from WiFi.

distance for interactions between 0 to 25 cm, it is a small and affordable, system-on-chip technology, which could be mounted on the small, spherical robot used in this thesis. All this makes it a good candidate for a sensor system measuring proximity.

2.7.3. Setup

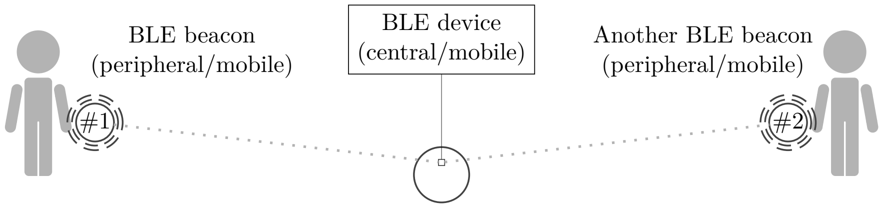

Figure 2.11.: The overall experimental setup from (Scheunemann, Dautenhahn, Salem, et al. 2016b). A central BLE device (mounted on the robot) was used to passively scan for advertisements of peripheral BLE devices (attached to/worn by humans). The central device computed packages with RSS data and ID of received advertisements. The self-contained system-on-chip transmitted these packages wired or wirelessly via BLE.

This section describes the BLE devices and their setup to enable a robot with a proximity sensor for close distances. Scheunemann, Dautenhahn, Salem, et al. (2016b) used BLE to enable a robot to receive information about a human participant's proximity, to recognize touch gestures and to distinguish between human participants. Figure 2.11 visualizes the setup: a mobile, rotating robot was equipped with a central BLE device and peripheral BLE devices were attached to the human participants.

Each of the peripheral devices (i.e. beacons) attached to humans advertised their ID. The central device passively scanned for these advertisements. The central device was a system-on-chip, meaning it could compute this information independently. It collected the RSS and the associated beacon ID, and transmitted this to an external computer which controlled the robot. This way the robot controller received information about the surrounding beacons and their RSS.

Components

Figure 2.12 shows the hardware components: (a) the beacon Gimbal Series 10 and (b) the central device BLED112. Both components are described in the next two subsections.

Peripheral proximity beacons Gimbal Series 10 Figure 2.12a shows the beacon *Gimbal Proximity Beacons (Series 10)* that was used in the setup. It is a small ($40 \times 28 \times 5.5$ mm), inexpensive (US$5) device, which can be powered with a standard coin cell battery of the type CR2032.

The device can be configured with two different proprietary protocols: the Gimbal protocol and Apple's iBeacon (Apple Inc. 2014). The setup here used the iBeacon protocol, simply

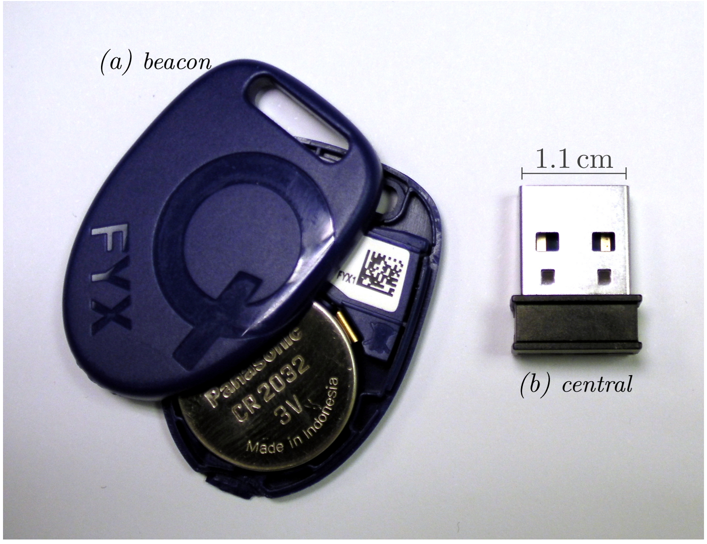

Figure 2.12.: The two BLE components used in the setup. (a) Gimbal Series 10 is a self-powered and configurable advertising beacon, which can easily be attached to humans, robots or objects. (b) Bluegiga's BLED112 is a central BLE device that scans for the RSS of surrounding peripherals.

because many BLE devices and upcoming products utilize the standard, making the integration of other devices in this architecture more comfortable. Apple's iBeacon was introduced in 2013, however, a protocol specification was only provided in 2015 (Apple Inc. 2014; Apple Inc. 2015). The protocol has the disadvantage that it does not report the transmission channel. A BLE scanner then receives RSS from different channels, which results in a *smearing effect* (Faragher and Harle 2015).

Each beacon has to be assigned a 20 byte payload. This is what the iBeacon protocol reserves to identify the beacon purpose and its position. The first 16 bytes are the beacon's UUID. This is a unique identifier of the beacon's application and allowed the scanner to pre-filter advertisements for only the beacons of the described setup. The last 4 bytes can specify information about the position of whatever the beacon is attached to. This can be the identifier of a participant, including more details such as whether the beacon is attached to their left or right wrist.

The iBeacons were configured in such a way that they allowed the maximum perception by the central device. For example, the antenna operated as an omnidirectional antenna, as opposed to a directed antenna. This way the direction from the beacon to the scanner was influenced the least. The beacons had the maximal transmission interval of ~100 ms and a maximum Transmission Power (txPower) of 0 dBm. This way, the central device received the maximum readings possible and the high power increased the accuracy. Both settings also maximized energy consumption. This, however, was acceptable as beacons powered with a coin cell battery will still last for several days.

A central scanning device Bluegiga's BLED112 The embedded firmware of the robot is proprietary, i.e. there is no possibility of changing the software. This made it impossible to use the existing Bluetooth adapter or to add an adapter. Therefore the central module needed to be independent of the other hardware in the robot.

Bluegiga's module BLE112 was used as the central device for scanning the environment for advertisements of other BLE devices such as beacons (Silicon Laboratories Inc. 2015). This device is based on the 6×6 mm chip CC2540F128 from Texas Instruments: a "cost-effective, low-power, true system-on-chip (SoC)" for BLE applications (Texas Instruments 2013). The radio frequency transceiver, the 8051 microcontroller unit (MCU), 8 kB SRAM, 128 kB programmable Flash and a complete BLE stack[9] enables full BLE device capabilities (see Bluetooth SIG 2016).

Figure 2.12b shows the USB dongle *BLED112*, which was used in this setting. With the size of 12.05×18.10 mm it is much bigger than the non-USB version. However, it allows for easier prototyping, as the dongle can be positioned flexibly to different robots/machines without any soldering. Bluegiga provides a protocol *BGAPI* to control the integrated BLE stack, i.e., to send commands and to receive events/responses, in either of the following two options:

- connect an external microcontroller/PC via UART or a USB CDC virtual serial port to control the stack with BGLib, an ANSI C implementation of BGAPI
- use the embedded CC2540 chip on the BLE112 and control it with BGScript, the BGAPI implementation in a BASIC-like scripting language

Here, the second option was used, which allowed the scanner to be used independently of other hardware. The central BLE device was set in a constant passive scanning mode, offering its own GATT Bluetooth profile. The machine which ran the robot client was connected via Bluetooth to this service and waited for changes in its characteristics. Each received advertisement package was sent roughly every ~100 ms, which triggered a two-step response event on the chip:

i. extract the sender ID and the corresponding RSS to the sender

ii. communicate the information to the robot client, i.e., change the profile characteristic

This way, changes in the RSS could influence the robot's behavior. More details of the implementation are publicly available (see Scheunemann 2018e; Scheunemann 2018a).

Proof-of-concept evaluation

Scheunemann, Dautenhahn, Salem, et al. (2016b) presented proof-of-concept evaluations. Two human participants wore wristbands with BLE beacons and a spherical robot was

[9]The full BLE stack includes protocols such as GAP, GATT, L2CAP and SMP.

equipped with a central device. Their results show that the above setup can be used (i) to increase the robot's awareness of participants that are present in its environment/proximity, (ii) to use a BLE device as a touch sensor for deliberate use and (iii) to enable the robot to distinguish between interacting individuals. This means that all capability requirements outlined at the beginning of this section were met.

Most relevant for this thesis is the evaluation of the proximity information. Figure 2.13 shows example data of the measured RSS between the central device and the beacon. What can be seen is that there was an almost linear relationship between the beacon distance and the RSS. For a beacon which operates on maximum power (txPower = 0 dBm), the relationship is linear in the range of 0 to ∼25 cm. This is in line with analysis by Faragher and Harle (2015).

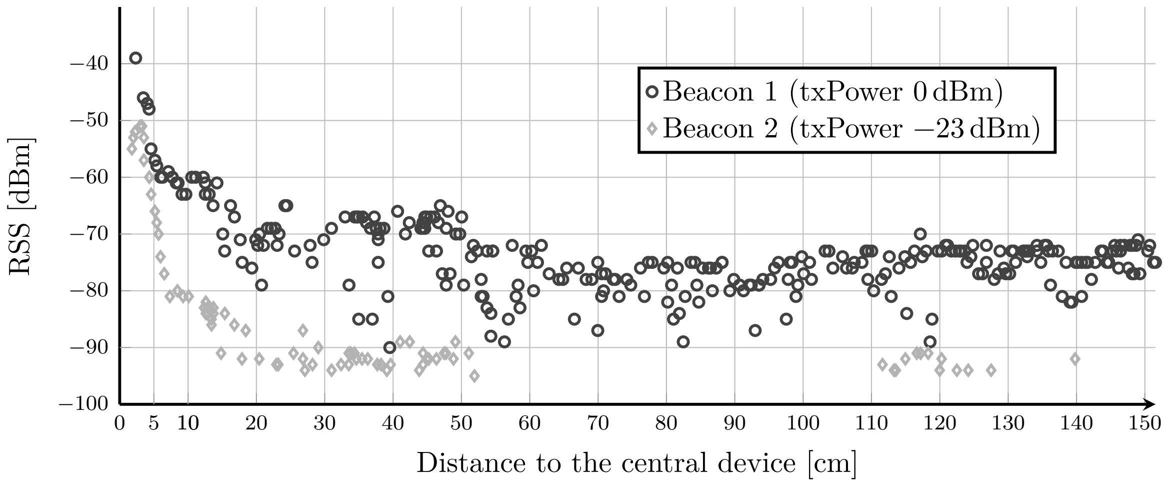

Figure 2.13.: The RSS of two beacons. Beacon 1 was set up with the highest possible transmission power (txPower = 0 dBm) and beacon 2 to the lowest (txPower = −23 dBm). The x-axis shows the distance between the beacon and the central device. Each beacon was moved with a roughly constant speed of $\sim 5\,\frac{cm}{s}$ *away from the central device.*

Setup adaptations for this thesis

The technological requirements were met with the above setup. However, what turned out to be a challenge was to translate the setup to the small robot platform used in this thesis. The robot is closed-source hardware, which rendered it difficult to equip with any additional hardware. This was the reason why a system-on-chip device was chosen in the first place. It could be equipped to a robot by only using its power. However, opening the shell damaged the robot in such a way that the dynamics changed and rolling was not as smooth anymore. This was too huge of an impact on the robot hardware. Initially, an important reason for an off-the-shelf robot was to be able to easily replace the robot with another from the same batch. This way, the robot replacement had similar characteristics. However, opening the robot changed the central feature of the robot: its locomotion, which made this benefit vanish (especially as the damage was not anticipated to be caused reliably).

For the final study, the setup can also be realized by swapping the places of the beacon and the central device. This needs some motivation. An added benefit of the studies in this thesis compared to the proof-of-concept evaluation above is the use of the wand-shaped tool. The human participant solely used the tool to interact with the robot. Therefore, the two devices were placed in the tip of the wand and the robot itself (see the next section). This gave the robot an idea about the proximity of the wand (and therefore the human). The central device could sense the bacon without any occlusions if it was in close distance, and the wand was close enough as it was used for any type of interaction. In particular, the proximity could be sensed by only using one beacon and one central. This in turn meant that both the devices could be swapped: the robot carried the more lightweight, self-powered beacon and the central device was placed in the wand shaped tool. This had two benefits: the most lightweight devices were placed on the robot directly, and the wand, with its pouch, offered more space to place the additional power source unobtrusively.

2.7.4. Assembly

This section describes the assembly of the setup with the components described above. The first step was to equip the robot with the beacon. This started by making the size of the beacon smaller, without harming its functionality (Figure 2.14a). Ideally, to avoid changes in the robot's dynamics, the beacon would then be placed as close to the robot's center of mass as possible. However, as it was discussed earlier, this would mean the shell needs to be cut open. This idea has been rejected because the robot's dynamics would change too much.

(a) prepare beacon case *(b) top mount* *(c) inside mount*

Figure 2.14.: The beacon's weight was reduced by modifying its shell in such a way, that it remains functional (a). In initial trials, the beacon was mounted on the top of the head (b). This changed the robot dynamics drastically, in a way that the robot could not keep the head on top anymore and quickly lost its head. The extra weight of ∼5 g had to be placed closer to the robot's center of mass. Therefore, the robot's head was cut open so it fits the beacon (c).

Instead, the first idea was to mount the beacon on the robot's head (Figure 2.14b). This already allowed the robot to sense the wand proximity. However, it turned out that the extra weight of ∼5 g was already enough to make the robot dynamics unstable: the robot could not keep its head upright and rather *pulls* it or *pushes* it. Therefore, the beacon was mounted as close to the center of mass as possible, meaning the head needed to be cut open and the beacon could then be pushed inside the head (Figure 2.14c). A white tape covered up the hole in the head to hide the additional robot functionality for the participants. Unlike for the shell, opening the head does not change the dynamics of the robot. This is because the prepared head can be reused on a new off-the-shelf robot.

The next step was to assemble the wand-shaped HRI tool. Figure 2.15a shows the components for assembling the tool. A thin metal tube held a USB female plug. The central device was plugged into the USB cable. The tip was shielded with a table tennis ball. Cables with a JST PH 2-pin connector[10] are soldered to the power wires of the USB cables. This enabled the ability to connect an external battery to the plug, which then powered the system-on-chip central device. The battery was hidden in a black pouch which was attached to the end of the wand. This design had two benefits. Firstly, it enabled the battery to be changed easily and without disassembling the wand. Secondly, the weight of components is closest to the point of human contact. This construction kept the additional weight on the tip of the wand to a minimum, which in turn made it more comfortable to handle the wand. Figure 2.15b shows how the assembled wand and robot were used in an experimental run with the author.

(a) Explosion view *(b) Usage*

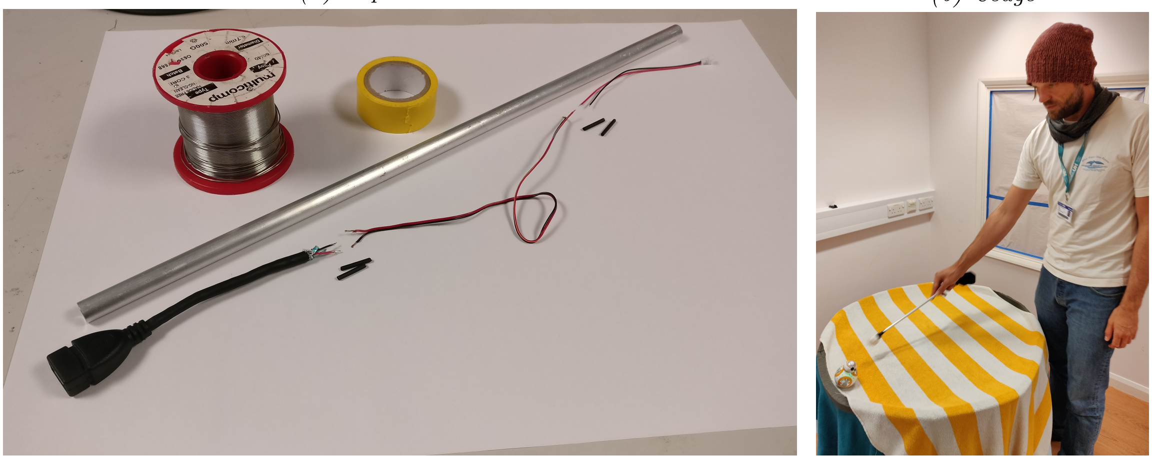

Figure 2.15.: The image shows an explosion view of the components of the HRI tool (a). Cables were soldered to a female USB connector and were placed in the metal tube. This way, an external battery could power the BLED112 and could be placed flexibly. The second image shows the assembled tool in use (b). A ping-pong ball covered the BLED112. A black pouch at the end covered a small battery.

2.7.5. Summary

The technology BLE and the proposed devices fulfilled the requirements outlined in subsection 2.7.1. The BLE devices were small and could be attached unobtrusively, which is important for research where human participants should be unaware of the robot's sensor capabilities. The technology could be self-contained on a chip. This means a central BLE device could scan surrounding beacons, but it did not need to be part of the robot itself. Instead, an external machine could be connected to the central device and gather the scanning

[10] JST describes a range of common standards developed by Japan Solderless Terminal.

results. This was able to extend the robotic sensing capabilities without the need to have access to the robot's firmware.

BLE offered the capability to quickly scan for surrounding signals of beacons. The robot could therefore react quickly to proximity changes to a beacon. BLE could use omnidirectional antennas, something other technologies, such as RFID, were lacking. The technology was cheap at the time of conducting the research, for example, a beacon cost US$5 and a central BLE device (a self-contained system on a chip) cost £11.49.

A limitation of BLE was that it needed to be powered somehow, in contrast to passive RFID tags. However, a beacon with a coin cell battery only increased the weight by ~5 g and a beacon could run on a coin cell battery for 1 to 2 years. This means the technology could still be applied to a small robot platform in most cases.

A further argument for BLE was that many manufacturers used the technology for their devices, such as mobile phones, modern smartwatches and fitness bracelets. This also resulted in research into using smartphones for, e.g., medical workflow tracking (Antunes et al. 2018) and driver and passenger identification (Ahmad et al. 2019). Due to its popularity, it can be expected that the amount of BLE devices in our environment will further increase in the future (Faragher and Harle 2015). The inexpensive availability (prices may even fall due to increasing popularity), the widespread use of BLE in existing devices and the many upcoming additional applications and hardware devices further enhanced the motivation to utilize this technology.

Scheunemann, Dautenhahn, Salem, et al. (2016b) provided a proof-of-concept evaluation that the technology can be applied to retrieve proximity and touch information, as well as to distinguish between participants. The technology is applied to the final study in chapter 5. The tool's usability and its impact on the participant's experience have not been directly evaluated.

Chapter 3.

Study I

3.1. Introduction

This study was the first attempt to answer the main research question of how an autonomously, intrinsically motivated robot is perceived, and whether this may sustain human-robot interaction (HRI) (RQ2). The robot's intrinsic motivations (IMs) were realized by maximizing the information-theoretic quantity of time-local predictive information (TiPI). The interesting property is that this computational model works on the sensor channels of the robot, without providing any meaning to the information. A robot that maximizes its TiPI in the sensor space would try to predict and excite its sensors and would only be constrained by its morphology and the environment. Implementations on a simulated robot have shown that this yields "exploratory and playful" behavior (Der and Martius 2012).

In order to investigate the human perception of this behavior generation, a few other questions arise: (i) what is a suitable baseline behavior for comparison, (ii) how to design a study to measure human perception and (iii) how to measure whether the intrinsically motivated robot can sustain HRI?

These research objectives accompany the whole thesis. The idea this chapter proposes is to have a baseline behavior that is as close as possible to the behavior generation of the autonomous robot. This was achieved by letting the robot adapt its parameters based on IM, and then letting the robot use one set of these parameters for its baseline behavior. This resulted in a fully autonomous robot behavior, which was not adaptive to the environment.

Two methods were realized in order to ensure that the participants interact with the robot. Firstly, the study *enforced* interactions. The environment was designed so that the robot autonomously locomoted on a table with a variety of surfaces and different altitudes. One side of the table was kept open so that the participant had to interact with the robot in order to keep it on the table. Secondly, the participant was further *encouraged* to actively seek interaction in order to fulfill a task: they were asked to understand, whether the two robots had different strategies.

Another question that is motivating the thesis is that of how to measure sustained interaction? Unarguably, the best way would be to measure the interaction time of human participants with a robot. However, any longitudinal study which was to assess the interaction time would simply be too cost-intensive to be rolled out without any prior quantifiable evidence that an intrinsically motivated robot may sustain the interaction. To measure whether intrinsically motivated autonomy in robots has the potential to sustain interaction, this study therefore measured human perception with the help of questionnaires.

The working hypothesis was that robots are more interesting to interact with if they have perceived agency or goal-directedness. This allows the human interaction partner to assign motivations to the robot, support or hinder its goals, or even sympathize with its *joy* when achieving a goal. Once humans identify something as an agent, they are likely to direct their attention toward that agent, trying to understand its goals, intentions, and behavior. Competence is a dimension that can measure agency or is sometimes interchangeably used to describe the concept of agentic (Fiske et al. 2007). There are two questionnaires which offer scale dimensions for *Competence* and related *Perceived Intelligence*: the Robotic Social Attribute Scale (RoSAS) and the Godspeed scale. The Godspeed questionnaire also offers other dimensions, which seem intuitively interesting when it comes to our judgment of robots: Animacy, Anthropomorphism and Likeability.

Animacy is a measure of *aliveness* (Bartneck et al. 2009). There is evidence that humans already perceive objects as animated if the cause of their movement changes is not obvious to the observer (Tremoulet and Feldman 2000). Perceiving something as *animated* seems like a necessary criterion of whether we can perceive agency. On the other hand, there is HRI research which found that the perception of Animacy is dependent on whether the human interacts with a robot or whether the human observes only an HRI scenario: the Animacy perception is hindered if, and only if, a robot's actions are visibly goal-directed, but only if we interact with the robot (Fukuda and Ueda 2010). In contrast, if a human only observes the HRI scenario, they perceive the robot which is following instructions the most animated. This shows two things: in order to understand what sustains HRI, it is important to have physical interaction. In addition, our perception of agency and goal-directedness might not reflect solely on how *alive* we perceive a robot.

Anthropomorphism is a measure often studied to understand our perception of robot appearances (e.g. Walters et al. 2007). Anthropomorphism is a complex and ambiguous term, without a consensus on a definition for the concept. However, the concept of anthropomorphism is believed to measure more than just appearance, but also relationships. Airenti (2015) claims that we anthropomorphize our pets, which we interact and play with, knowing they are not humans. We treat them with similar care as we treat our peers, to the point that we develop similar empathy (ibid.). In order to understand whether to reach similar enjoyment or close relationships as we do with pets, anthropomorphism might be an important dimension to understand more about HRI.

It seems therefore plausible to not just collect participants' responses to Competence and Perceived Intelligence, in order to investigate their perception of agency, but also to accompany these measures with the dimension Animacy, Anthropomorphism and Likeability from the Godspeed scale.

3.1.1. Research questions

Following the arguments above, this study investigated the following research questions:

RQ1 Is an autonomous, intrinsically motivated robot perceived as more animated or anthropomorphized than a reactive baseline behavior?

RQ2 Is an autonomous, intrinsically motivated robot more liked than a reactive baseline behavior?

RQ3 Is the perceived agency of an autonomous, intrinsically motivated robot higher compared to a reactive baseline behavior?

3.1.2. Overview

In this study, an intrinsically motivated robot was compared to a robot with a reactive baseline behavior. Section 3.2 presents the design of the baseline behavior, following a description of the study design in section 3.3.

Firstly, the used minimal, non-humanoid robot platform and the environment is described in subsection 3.3.1. The robot platform Sphero was used, which is a spherical, non-humanoid platform with only two degrees of freedom. This reduces observable complexity and, in addition, should help to decrease a participants' expectation bias of the robot and prevent participants from anthropomorphizing the robot. Subsection 3.3.1 describes the tasks of the robot and the human participants. Subsection 3.3.2 presents the conditions of the study. One condition was the robot with a reactive baseline behavior. The other condition used a fully autonomous adaptive robot behavior based on predictive information (PI) maximization. The behavior patterns are described in subsection 3.3.3 to better understand the different characteristics of the behaviors in the two conditions.Subsection 3.3.4 describes the measures used for the study. These were standardized, popular scales used in HRI, namely the RoSAS and the Godspeed scale. The presentation of the study design is finalized by the description of the procedure (3.3.5) and the sample (3.3.6).

Section 3.4 presents the results. The study did not find any statistically significant results. Contrary to the hypothesis, the reactive baseline robot was more liked and perceived slightly more intelligent and animated. What stood out was the medium effect for the factor Warmth, a universal dimension from social cognition implemented by the RoSAS questionnaire. Section 3.5 discusses these results, followed by a detailed description of the implications for the studies that followed in section 3.6. Then section 3.7 completes the chapter with the conclusion.

3.2. Baseline behavior

A challenge for investigating the human perception of robots is the baseline comparison. The simple act of moving (Dautenhahn 1997; Hoffman and Ju 2014) or the appearance of a robot's *head* (Blow et al. 2006; Hoffman and Ju 2014) can already change a human's perception of a robot. Therefore, to compare the perception of different robot behaviors, it is best to choose the same robot platform.

But what is the best way to generate a baseline *behavior*? The following three alternative means of behavior-generation for serving as a baseline can be considered:

i. human remotely controls the robot

ii. random behavior

iii. pre-adapted reactive behavior

Human remotely controls the robot Ideally, I wanted to see how the algorithm compares to a human remotely controlling the robot. However, human-controlled behavior has a high degree of variance, dependent on the particular human controller. Furthermore, it is unclear how much access the human controller should have to environmental information.

If a human directly observes participants, they may obtain much more information than the robot, giving them an unfair advantage when creating behavior responses. It can be considered a separate means of research to provide the sensor input in any meaningful way for the human while keeping it comparable to the perception of the robot. For example, if the human controller is limited to only the plots of the robots' sensors, then the human controller would likely struggle to make sense of this limited input.

Again, eventually, I would like to test whether an intrinsically motivated robot is perceived as more or similarly agent-like compared to a remote-controlled robot. A Turing test could be used to answer this question. However, creating a meaningful control for the limited platform would need a preceding investigation of how to provide fair, shared sensory input for humans and robots. In addition, PI maximization would need adaptation to be applied to more complex robot platforms with sensor capabilities similar to a human. Therefore, a different baseline behavior would need to be implemented first.

Random behavior The problem with using random behavior as a baseline is that "randomness" actually has a set of parameters that needs to be chosen. For example, how often do values change, or is it the change of value or the value that is being randomized. The issue with mimicking the behavior generated by PI maximization is that the change of behavior is dependent on the input. So, when should that change? An option considered, for example, was a Braitenberg-style vehicle (Braitenberg 1984). In a thought experiment, Braitenberg (ibid.) created vehicles with surprisingly complex behavior generated out of a set of simple internal rules. For example, imagine a vehicle with two proximity sensors in the driving direction. If the vehicle senses an obstacle to the right, it could turn to the left, and vice versa. This yields quite a complex obstacle avoidance behavior, without the use of, e.g., a state machine. The robot platform used here, for example, could drive toward and then bump into an obstacle. Depending on the tilt, it could then change its heading to the left or right and continue as soon as its tilt is small. Sometimes in addition and randomly, it could just change its heading without an obstacle to add extra complexity. I performed some preliminary trials with random values, but I was quickly facing the question of a fair baseline behavior again. The biggest issue with a random baseline behavior is its comparability among research and robot platforms. Having the experimenter choose these values leads to basically designing a certain kind of behavior (chosen from a whole range of behaviors), which makes it problematic as a baseline. If a scientist wants to replicate this experiment in the future, choosing a different robot platform, a systematically created baseline behavior, ideally independent of the robot's morphology, would be ideal, and would also be helpful when comparing results with the current study.

Pre-adapted reactive behavior I decided to use a pre-adapted reactive behavior that has adapted its network weights to the environment, but does not continue to adapt during the experimental trials. The pre-adaptation was done in the same environment, with the same behavior-generation algorithm (i.e. PI maximization) and the same robot platform as the robot in the non-baseline case. Later during the experiments, the baseline robot used all the same sensor inputs, but it only used the network parameters[1] adapted at the end of the pre-adaption. That way, the baseline was reactive to the same sensor input. It was reactive to human and environmental perturbations, i.e., it was sensitive to changes either of the environment, by input from human participants or just from moving itself. However, the baseline robot did not continue the adaptation of its parameters. Its states were thus deterministic, i.e., on the very same sensor input to the robot in the very same state, its next state would always be the very same. However, this was challenging to observe in the real world, due to the huge variety and combinations of sensor inputs and environmental states, which made the behavior a good baseline candidate. Receiving the network parameters was done with three pre-adaptations for 5 minutes. The network parameters were then randomly chosen among the three. The initial parameters were the same for the baseline and the non-baseline behavior.

3.3. Study design

As mentioned earlier, the design of the study was a challenge in itself, due to a lack of previous investigations on the human perception of intrinsically motivated robots. This

[1]This refers to θ, a set of parameters representing the synaptic weights and biases. Their dynamics were described in subsection 2.3.2.

section describes and discusses the study design.

3.3.1. Robot, environment and tasks

There is evidence that the appearance of the robot and its behavior needs to be consistent or, in other words, balanced (Ishiguro 2007; Fukuda and Ueda 2010). This means a very realistic robot may elicit discomfort if its motion does not match expectations. For example, a humanoid robot with the degrees of freedom similar to a human and some technical complexity in its arms and face may elicit discomfort, not just by not being able to gesture or communicate, but simply by not moving in accordance with our expectations.

(a) The robot platform *(b) Schema* *(c) HRI environment*

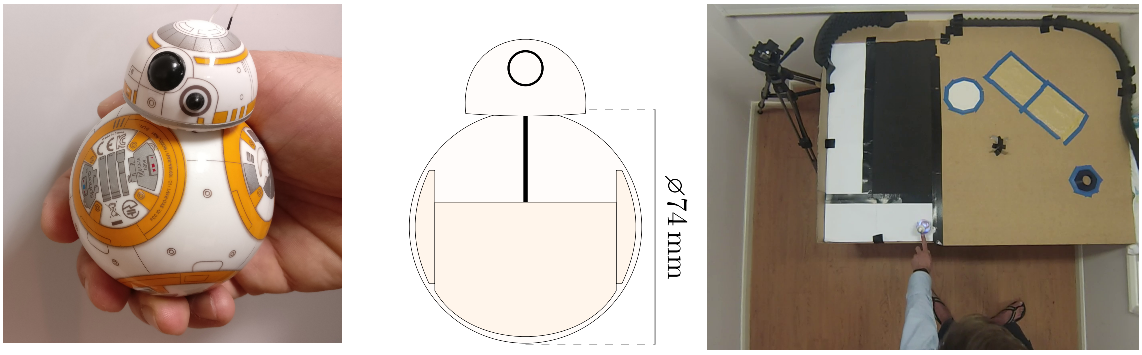

Figure 3.1.: (a) The used robot platform: BB8 from Sphero. (b) A 2-D cross-sectional view of the robot. A two-wheel vehicle, kept in position by a heavy weight, moves the sphere when driving. A magnet attached to the vehicle keeps the head on top of the sphere facing the direction of movement. (c) The environment the robot explores during the trials from a bird's eye perspective. The white area is paper, the black is foam material and the beige-colored area is wood. At the top of the foam material is a hill area and a pit in the lower part. The bottom edge does not have a wall, which forced the participant to interact with the robot.

In this study, a very simple platform with a few degrees of freedom was chosen. This lowered the expectation bias of the participants and allowed me to solely focus on the effects induced by predictive information. Although simple, the robot offers some sensor capabilities to enable rich behavior and adaptation capabilities to the environment. A suitable platform is the off-the-shelf spherical robot from the company Sphero, specifically, the BB8 platform (Sphero, Inc. 2020a; Lucasfilm Ltd. 2015) as depicted in Figure 3.1a. In the BB8 version, the spherical robot Sphero has a head. A magnet keeps the head in driving direction, which gives the user a sense of the heading direction of the robot. This, and the fact that many people know the robot from movies, provides a better impression of a robot than using the original Sphero: a white, solely spherical robot. The robot specifications are described in more detail in section 2.5 (pg. 36).

Figure 3.1c shows the experimental environment. Two tables formed the space where the robot could move around. The area was 180 cm × 120 cm in size. It was open on one side where the participant was supposed to stand and interact with the robot. In Figure 3.1c you can see the author nudging the robot. The surface of the table differed in friction and height. The black foam area had a hill (top) and a pit (bottom). Additionally, the black area and the white paper area were softer and had higher friction compared to the wooden part.

The participants' task was to observe the robot and understand whether it had a strategy for exploring the environment. I directly asked the participants: "Aside from keeping the robot in the area, try to observe the robot's behavior. Also, try to understand the robot's exploration strategy (if any)?" (the full protocol is described in subsection 3.3.5). The idea was that this would encourage the participants to interact with the robot. Additionally, one side was kept open, so participants had to actively interact with the robot to prevent it from falling off the table. The idea was that this enforcement of interaction would provide the participants with a better understanding of the robot's capabilities and behavioral richness.

3.3.2. Conditions and their order

The experiment consisted of two different types of behavior generation (i.e., conditions):

REA_b **(reactive):** participants interact for approximately 10 minutes with a reactive robot and were then asked about what they had seen.

ADA_b **(adaptive):** same as REA_b, but the robot was continuously adapting, based on maximization of PI as a motivation to interact with its environment.

The adaptive robot in the ADA_b condition realized behavior motivated by maximizing PI, and it continuously updated its internal networks based on that gradient during the experiment. The PI formalism is described in detail in section 2.3 of the background chapter. The reactive robot in the REA_b condition started with the same network configurations as the adaptive one (based on pre-trial adaptation). This determined how it reacted to sensor input, but it did not further update its internal network during the experiment.

The type of behavior generation was a within-subjects variable, i.e., each condition was presented to all participants. The order of REA_b and ADA_b was therefore randomly assigned but counterbalanced over the number of participants to avoid interaction effects. Table 3.1 shows the orders of conditions, along with the number of participants.

Table 3.1.: Order of conditions

	order of conditions	participants
A	$ADA_b \rightarrow REA_b$	8
B	$REA_b \rightarrow ADA_b$	8

Figure 3.2 shows how the starting configurations for all networks of both conditions were derived. They were generated in two steps. Firstly, three trials with the robot for 5 minutes in the previously described environment were conducted. At the end of each trial, the robot's

network configurations were saved. In a second step, one of these network configurations was randomly chosen as the starting configuration, i.e., for condition ADA_b and REA_b.

The PI formalism allows having different levels of adaptivity to changing environments and new stimuli. The update rate for ADA_b was determined empirically. I noticed that the robot sometimes got caught in the pit mentioned earlier. When this happened, it would need to adapt in order to leave and continue exploring. The ADA_b adaptation rate was set so that the robot would change its behavior and leave the pit in less than 20 seconds. The hypothesis was that a high adaptation rate yields a higher perceived intelligence, as the robot would continuously adapt to new stimuli and change the way it would react to certain inputs. The robot in the ADA_b condition was assumed to be perceived as more intelligent, as it would be the only one able to leave the pit. This is discussed later in more detail.

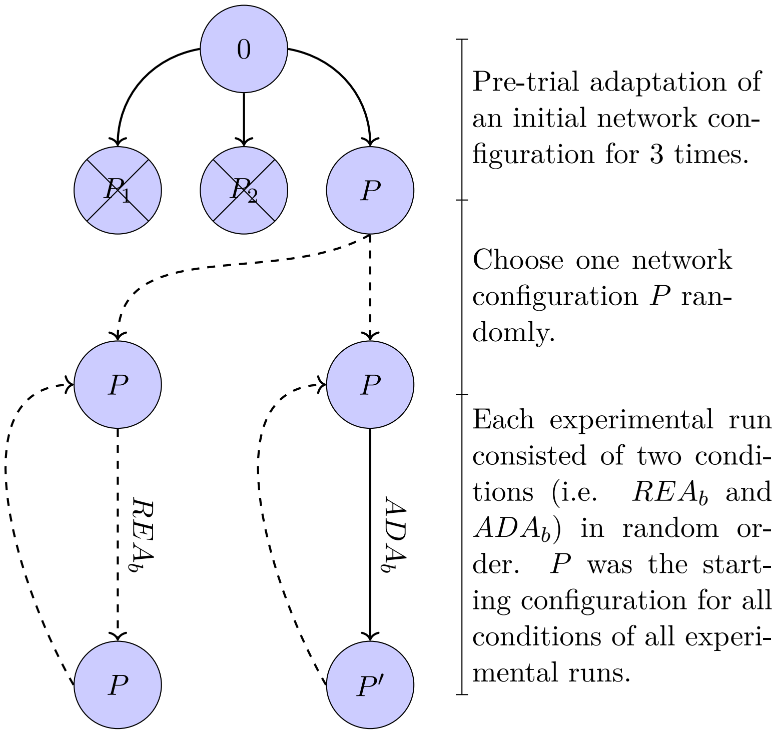

Figure 3.2.: The starting point of each condition was a network configuration P. The network configuration P was chosen randomly from three 5 minute pre-adaptation trials. Only in the ADA_b condition the adaptation continued during the human-robot interaction trial. In condition REA_b, the robot was only reactive but was not adaptive to the environment.

3.3.3. Description of the robot's behavior

The chosen sensors determined the robot's behavior to a large extent, as PI tries to excite sensor input. For example, if one decides to only use the inertial measurement unit (IMU) reading of the yaw angle speed, then the robot only needs to adapt its heading in order to excite the sensor. It is likely that it would not generate any output for rolling forward, as such a movement would not excite the sensors. Empirically, I decided on the following sensors as input:

- pitch and roll angles from the IMU
- x and y component (forward/backward and left/right) of the accelerometer
- z component of the gyrometer, i.e., the angular velocity of the robot when spinning

The aforementioned strategy of using a fixed network for the reactive robot yielded a somewhat predictable behavior for the condition REA_b. The robot preferred left turns in light of environmental perturbations or human interaction, i.e., if it hit a wall, it would often turn left. Its major trajectory was that of circling in different radii. Given this behavior, it seemed likely that people would become bored very quickly. However, in initial trials and according to the observations presented later, almost all participants did not recognize the mentioned pattern.

The adaptive robot (ADA_b) started with the same network configurations as the reactive robot (REA_b). A very high update rate was chosen for its model, as discussed in section 3.2. Its trajectory had a tendency to be straight, and if it reached an obstacle it adapted its heading to be able to continue moving in another direction. However, as soon as a participant interacted with the robot, it was not trivial to understand what the robot would do next to increase sensory stimuli. Example videos for both conditions are publicly available (see Scheunemann 2019; Scheunemann 2017e).

3.3.4. Measures

Two standardized scales were used to measure the participant's perceptions of the robots: the Godspeed scale designed by Bartneck et al. (2009), which has been widely used in many experiments, and the Robotic Social Attribute Scale (RoSAS) designed and evaluated by Carpinella et al. (2017), which is comparably new to HRI. Using standardized questionnaires helps when comparing the results with other experiments. The use of questionnaires was motivated in section 2.4.

Godspeed uses a 5-point semantic differential scale and investigates for the dimensions Anthropomorphism, Animacy, Likeability, Perceived Intelligence and Perceived Safety. The RoSAS collects responses for the dimensions Warmth, Competence and Discomfort. The authors of RoSAS do not recommend a specific scale, but instead recommend having a neutral value, e.g., an uneven number of possible responses. The used questionnaire consists of 7-point Likert-type items.

The focus of this study lied on the dimensions Animacy, Anthropomorphism, Likeability, Perceived Intelligence and Competence to answer the research questions outlined in subsection 3.1.1. However, all the other dimensions were included in the questionnaire and are also reported. This entails Perceived Safety from the Godspeed scale, and Warmth and Discomfort from the RoSAS. This was done, on the one hand, to hide the questionnaire intention and on the other hand to check for other effects or suitable dimensions for later study designs. In addition, it further helps when comparing effects in follow up studies. The dimensions consist of several items. The mean of all items encompassing a dimension is used to calculate the dimensions' response.

Both scales were included in the questionnaire handed out to the participants after each condition. Two additional open-ended questions were asked in the last questionnaire after both conditions had been conducted:

i. "Can you describe the different behaviors of the robot? Did the robot have any particular strategy for exploring?" and

ii. "What were the best and/or worst aspects of the robot's behavior?".

The purpose of these open-ended questions was twofold: on the one hand, answers could reveal more insights into what participants focused on. This in turn could provide ideas of how to amend the study design or whether the baseline robot was perceived particularly differently. On the other hand, they were used to draw the participants' attention to the robot and encouraged their interaction. The pre-test and post-test questionnaires are attached to this work in section B.1 (on page 175 and 178).

3.3.5. Procedure

Participants were welcomed to the experimental room and they were handed an information sheet. They were welcomed 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. It was in the beginning and at this point that it was emphasized that they could leave the study whenever they feel uncomfortable, stressed or bored.

After that, the environment and the robot were presented and briefly described, starting with a description and demonstration of the interaction possibilities. They were shown how to use their hand as a *wall* or nudge the robot to prevent it from falling off the side of the arena that was not enclosed by a wall, or to illicit new behavior through interaction. For each interaction shown, the participants were asked to repeat them:

- "you can use your flat hand to stop it"
- "you can use your pointing finger to poke it"
- "you can use your flat hand and rotate it"

The experimenter then described the main conditions of the experiment:

- "the same robot hardware is used for both 10 minute interactions"
- "at the beginning of each interaction, the robot has almost no knowledge about its body and/or about the world"

Participants were informed of the robot's task to explore the environment.

- "the robot will try to understand/explore the world and itself"

They were asked to observe whether the robot followed a particular strategy to do so, and if they could identify any specific behavior. They were also asked to prevent the robot from rolling over the open edge.

- "you need to take care the robot is not falling off the table"
- "try to observe the robot's behavior and try to understand the robot's exploration strategy (if any)?"
- "you can use all techniques above to further understand the robot's strategy"

The two tasks above aimed to initiate a human-robot interaction so that participants' responses reflect their perception of the interaction. Their task to prevent the robot from falling off the table aimed to *enforce* the interaction, while the question about robot strategies aimed to *encourage* it.

Participants then filled in the pre-questionnaire. This gathered information regarding their sex, age and background. Next, the two conditions were presented to the participants in a randomized order. The resulting allocation to either order A or B was counterbalanced, as shown in Table 3.1. Each condition lasted approximately 10 minutes. They filled in two post-questionnaires containing the two scales and the two additional questions discussed earlier (cf. subsection 3.3.4). The questionnaires can be found in Appendix B. The entire experiment took 50 to 60 min per participant.

3.3.6. Sample

Sixteen participants were recruited (5 females; 11 males) between the ages of 23 to 60 years ($M = 33.4$, $SD = 9.3$). Figure 3.3 shows the distribution of participant's ages.

All participants but one were university staff or students, and all of them had a background in Computer Science. To gain insight into their experience with robots, the participants were asked how familiar they were with interacting with robots, programming robots and the chosen robot platform. A 5-point Likert scale was chosen with the value 1 for "not familiar" and 5 for "very familiar". Figure 3.4 shows all the responses to the self-assessment. The participants' background were reflected in their responses to the pre-test questionnaire. Their self-assessed experience for interacting with robots gave an average of 4.4 (a) and the average

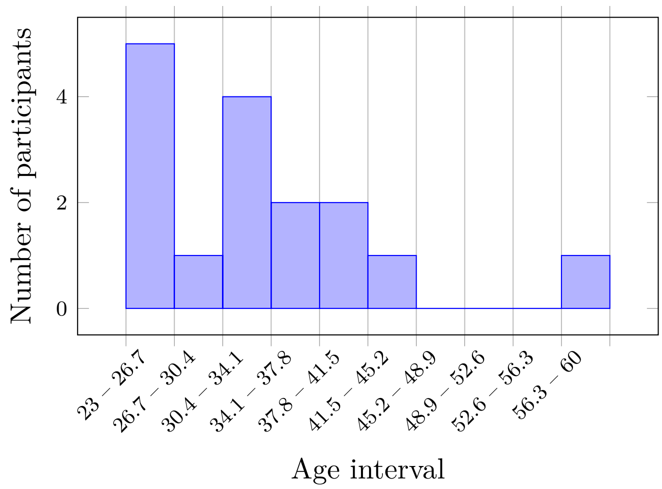

Figure 3.3.: A histogram of the participant's age in the study. It is skewed toward younger participants.

familiarity with programming robots was 3.8 (b). However, their experience with the chosen robot platform was relatively low and rated an average of 1.9 (c). Having participants with higher experience in robot interaction and robot programming was on purpose. I expected that those participants would be more critical toward an HRI experiment. Importantly, all participants were naïve with regard to the purpose of the experiment.

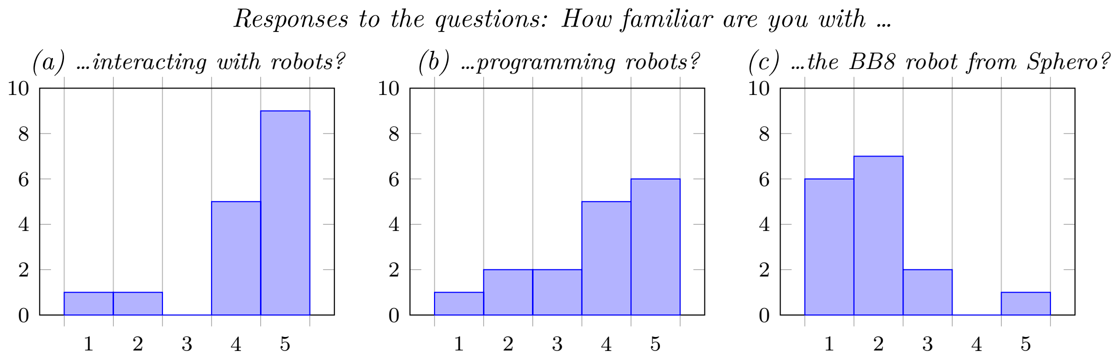

Figure 3.4.: The response distributions of the pre-test questionnaire on 5-point Likert-type items. Participants were asked about their familiarity with (a) interacting with robots, (b) programming robots and (c) the Sphero BB8 version. The majority of participants had little experience with the Sphero robot but had interacted with robots already.

The study was ethically approved by the Health, Science, Engineering & Technology ECDA with protocol number aCOM/PGR/UH/03018(1). The notification of approval is attached to this work in section C.1. The experiments were conducted from May to June 2018 over the course of 13 days. The anonymity and confidentiality of the individuals' data are guaranteed.

3.3.7. Data preparation

Table 3.2 shows the results of testing the internal consistency reliability of the scales of both the standardized questionnaires, analyzed with the use of Cronbach's α. The item *quiescent-surprised* was negatively loaded on the scale dimension Perceived Safety. Even if reversed, the reliability was poor. The item was therefore removed. After preparation, all scale dimensions showed a good reliability, ranging from $\alpha = 0.75$ to $\alpha = 0.92$. This was evidence that the dimensions could be analyzed without any further preparation.

Table 3.2.: Internal consistency reliability scores.

	dimension	items	α
RoSAS	Warmth	6	0.84
	Competence	6	0.87
	Discomfort	6	0.81
Godspeed	Anthropomorphism	5	0.75
	Animacy	6	0.75
	Likeability	5	0.89
	Perceived Intelligence	5	0.87
	Perceived Safety	2	0.92

3.4. Results

This section presents the results of the study in three different subsections. The first two subsections present the results of the quantitative analysis of the questionnaires[2]. The third subsection is concerned with the qualitative analysis, presenting insights from the open-ended questions.

The order of the two conditions was counterbalanced, i.e., the sample was split in half and the two orders A and B had the same amount of participants (cf. subsection 3.3.2). In addition, the order had been randomly assigned. All this was done to avoid interaction effects, i.e., to avoid that observed condition effects are influenced by the order of the conditions. Subsection 3.4.1 shows that there was no evidence for interaction effects. This, in turn, allowed investigating the condition effects, i.e., main effects, independently of their order.

Subsection 3.4.2 presents the main effects between the conditions. An analysis of variances, or a non-parametric type as used above, would not show any direction or size of an effect. Therefore, a two-way paired difference test was used to understand possible effect directions: the Wilcoxon signed-rank test. The results did not indicate any statistical significance, but they revealed medium effects for the dimensions Warmth, Discomfort (both RoSAS) and

[2]All tests were non-parametric. This was mainly because the collected questionnaire responses were ordinal, but also as these tests are more robust for small sample sizes.

Perceived Intelligence (Godspeed). The section is concluded with the results of a qualitative analysis of the open-ended questions (subsection 3.4.3).

3.4.1. Interaction effects

An analysis of variances (ANOVA)[3] is commonly used for investigating interaction effects, i.e., effects that show whether the order of the conditions influences the participants' responses to a condition. The study had two independent variables: one within-subjects variable and one between-subjects variable. The type of *behavior generation* (short: behavior) was an independent within-subjects variable. It consisted of the two levels: the conditions REA_b and ADA_b. The between-subjects variable was the *order* of how the conditions were presented and had the levels A and B (cf. subsection 3.3.2).

Table 3.3 shows the results of the ANOVA-type test for each questionnaire dimension, i.e., for each independent variable or factor. The last column reveals that there was no

Table 3.3.: ANOVA-type test results for the independent variables "type of behavior generation" (level REA_b, ADA_b), their "order" (level A, B) and the interaction of both variables "order:behavior" for the dimensions of the RoSAS and Godspeed scale.

		order			behavior			order:behavior		
	dimension	F	df1	p	F	df1	p	F	df1	p
RoSAS	Warmth	1.411	1	0.235	0.355	1	0.551	1.890	1	0.169
	Competence	0.022	1	0.881	0.001	1	0.981	0.210	1	0.647
	Discomfort	1.889	1	0.169	2.359	1	0.125	0.002	1	0.965
Godspeed	Anthropomorphism	1.606	1	0.205	0.057	1	0.812	0.982	1	0.322
	Animacy	2.367	1	0.124	0.009	1	0.923	0.017	1	0.897
	Likeability	0.035	1	0.852	0.033	1	0.857	0.813	1	0.367
	Perceived Intelligence	0.250	1	0.617	1.649	1	0.199	0.007	1	0.935
	Perceived Safety	0.370	1	0.543	0.742	1	0.389	0.403	1	0.526

interaction between the condition and order, i.e., all p values were bigger than the confidence level of 5 % ($p > 0.05$). This means that the conditions could be analyzed independently of their order [4]. The second column shows that there seemed to be no statistically significant condition effect for any of the dimensions. The next section discusses the main effects in more detail.

[3] For computing the ANOVA-type test the `R` package `nparLD` was used. As the study consisted of one within-subjects variable (behavior) and one between-subjects variable (order), it could be expressed as F1-LD-F1 Model. The `nparLD` package offers the function `f1.ld.f1()` for computing such models.

[4] Note that the same test was used to analyze if the participants' self-reports from the pre-test questionnaire had an effect on the participants' responses to the conditions, but no statistically significant effects were found.

3.4.2. Main effects

This section discusses the main effects of the study. A main effect is the effect of a single independent variable (here: the behavior generation) on a dependent variable (here: the questionnaire dimension) while ignoring all other independent variables. As discussed before, the main effects could be analyzed individually as there were no interaction effects between the conditions and their order.

Table 3.3 indicated that there were no statistically significant effects between the conditions. This section adds an investigation of the underlying effect sizes. Even if there is no statistical significance, an effect size can reveal tendencies for further investigations or adaptation in future studies.

To understand the effects it is worthwhile to investigate the actual change of the within-subjects variable *behavior* for one participant, rather than the total responses to one condition. The question was: does a number of participants respond higher to one condition than the other? This was important to understand whether many participants perceived one condition differently than another condition. For example, if all participants responded to the dimension Perceived Intelligence for the robot in the ADA_b condition a bit higher than in the REA_b condition, this would yield the insight that ADA_b was perceived as more intelligent among participants. This type of analysis allows revealing effects even if participants answered mostly at the extreme ends of the scales. For example, one participant might have been excited to see a robot locomoting and responded in all conditions at the right end of the scale. Another participant might have been more reserved toward a locomoting robot and rather answered on the left side of the scale. The important question of this study was: do both of them, nevertheless, perceive the ADA_b robot as more intelligent than the REA_b robot?

A test investigating for such effects is called a paired difference test. The Wilcoxon signed-rank test[5] is such a paired difference test and a non-parametric alternative to the popular paired t-test. Figure 3.5 shows the results for the two-sided Wilcoxon signed-rank test for comparing the difference $REA_b - ADA_b$. Depicted are the results for all dimensions (y-axis). If the response to ADA_b was higher than REA_b for one participant, then the difference is negative. If the median of all the differences (i.e., the point estimate) is negative (x-axis), then there was a tendency that most participants scored ADA_b higher for that observed dimension.

The 95 % confidence interval is depicted as error bars. If the confidence interval entails 0, then there was no certainty as to whether this difference was truly an effect in one direction, i.e., the effect was not statistically significant and the p value would be larger than 5 %.

Figure 3.5 shows that participants seemed to perceive the robot in the REA_b condition as more intelligent and competent. In addition, participants perceived the robot in REA_b as more safe, but the uncertainty of that effect spanned over a large range and was therefore

[5] The Wilcoxon signed-rank test is part of `R`'s built-in `stats` package and is implemented as `wilcox.test()`.

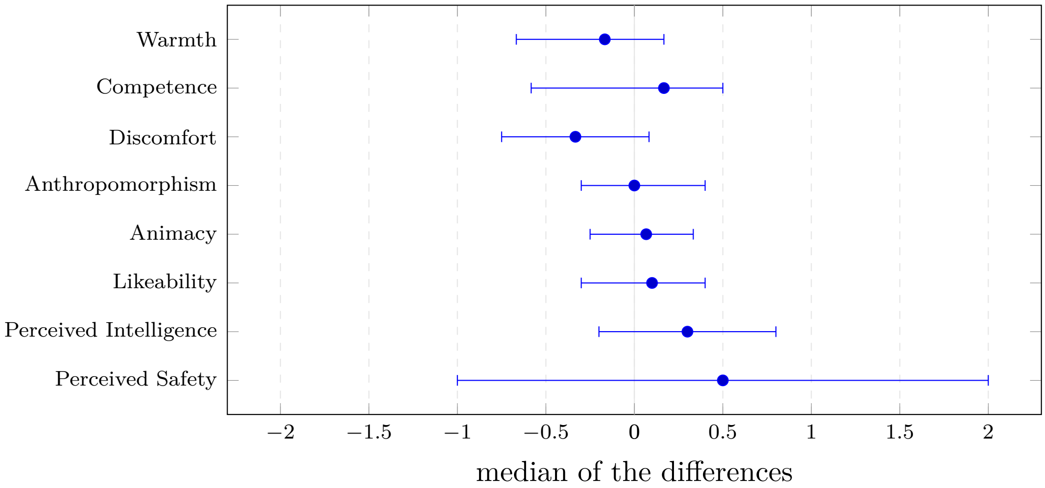

Figure 3.5.: The results of the two-sided Wilcoxon signed-rank test. Depicted are the point estimates (median of the differences) between the condition REA_b and ADA. The 95 % confidence interval is plotted as error bars.

ignored. On the other hand, there was some effect that participants perceived the adaptive, intrinsically motivated robot (ADA_b) as more warm and, in the same way, their Discomfort perception was increased.

However, Figure 3.5 also shows that all these effects were not statistically significant for any of the dimensions. This means there was not enough evidence that the robot's Animacy and Anthropomorphism (RQ1), Likeability (RQ2) and Competence and Intelligence (RQ3) in the ADA_b condition was differently perceived by the participants than for the robot in the REA_b condition.

There were two possible causes. Firstly, there may have been an effect but the chosen sample size was too small to reveal any statistical significance. Secondly, there may have been simply no effect. Finding out the needed sample size is challenging to tackle before conducting the experiment. A power analysis is the tool for revealing the needed sample size prior to a study. However, without a comparable study or any model to predict the effects, this was not possible in advance. Although no statistical significance was present, putting the results in relation to the sample size helps to see effects for further investigation.

Table 3.4 shows the test statistic V and its p value of the Wilcoxon signed-rank test. It also presents the detailed results of the point estimate and its corresponding confidence intervals which were depicted in Figure 3.5. In addition, the table shows the results of the standardized effect size r.

The standardized effect size r is a robust measure for small sample sizes present in this study (Rosenthal et al. 1994; Yatani 2016). The underlying effect is either small ($r \geq 0.1$), medium ($r \geq 0.3$) or large ($r \geq 0.5$)[6]. In a sense, r normalizes the p value by taking the

[6]There is no mutual agreement on how to verbalize the effect size r. However, the subjective interpretation of Pearson's r by Cohen (1992) has been often used. I adjust r upward to interpret the values.

Table 3.4.: Wilcoxon signed-rank test results between REA_b and ADA_b for all dimensions.

	dimension	V	estimate	95 % confidence interval lower bound	95 % confidence interval upper bound	p	r
RoSAS	Warmth	27.0	-0.167	-0.667	0.167	0.346	0.236
	Competence	65.0	0.167	-0.583	0.500	0.776	0.071
	Discomfort	21.0	-0.333	-0.750	0.083	0.087	0.428
Godspeed	Anthropomorphism	47.5	0.000	-0.300	0.400	0.889	0.035
	Animacy	68.0	0.067	-0.250	0.333	0.649	0.114
	Likeability	51.0	0.100	-0.300	0.400	0.700	0.096
	Perceived Intelligence	81.0	0.300	-0.200	0.800	0.233	0.298
	Perceived Safety	19.5	0.500	-1.000	2.000	0.350	0.234

sample size into account. This is good for two reasons: firstly, it allows quantifying the effect, and secondly, it allows comparing the effect between studies with different sample sizes. For the present study, the effect size r provided an indicator for the magnitude of effect that the participants' perception was different.

The results for r showed that there was possibly a medium and large effect for Warmth and Discomfort, respectively, in favor of the adaptive robot (ADA_b). There was also a medium effect for Perceived Intelligence in favor of REA_b. There was no evidence that participants perceived one robot more competent than the other. There were only small or no effects for Likeability, Animacy and Anthropomorphism.

3.4.3. Qualitative analysis

This section presents the analysis of the answers to the open-ended questions. The idea is to better explain the non-statistical significance found above and to discuss the ideas that were considered for the design of the next study.

Participants were not explicitly asked for differences in the seen behavior. However, the answers to the open-ended questions about the robot strategies could be analyzed for a tendency. For example, if a participant addressed any differences. All participants answered the open-ended questions with differing detail and length.

One open-ended question was addressing the robot's strategy and behavior directly: "Can you describe the different behaviours of the robot? Did the robot have any particular strategy for exploring?" A first overview of the answers revealed that some participants described the strategy of the robot. If they found a difference, they either explained the robot as following the wall or having a more circling behavior. Furthermore, 7 participants described the strategy of the robot in ADA_b similarly. Table 3.5 shows the comments they added to the questionnaire. They considered the strategy of the robot as mainly following the edge.

Two of them also thought the robot in REA_b was mostly circling. One pointed out that this was mainly left turns. Their comments are provided in Table 3.6. Subsection 3.3.3 indeed

Table 3.5.: Robot in ADA_b has a tendency to follow the wall.

ID	order	description fraction
502	*B*	"following the edge of the obstacle"
503	*B*	"It seemed to not go away from obstacles but rather tended to go along / around them"
504	*A*	"exploring the environment principally following the boarders of the area"
509	*B*	"seemed to follow alongside edges" and "outter limits well explored"
514	*B*	"the robot was able tot follow walls"
515	*A*	"seemed to follow edges quite a lot"
516	*A*	"the robot was trying to follow the wall"

describes the reactive robot behavior (REA_b) as mostly circling in left turns with various radii. However, only two participants pointed this out, which showed that the baseline behavior was not perceived as too monotone. Table 3.7 shows the comments of 3 participants

Table 3.6.: Robot in REA_b has a circling behavior.

ID	order	description fraction
509	*B*	"always turning left"
516	*A*	"following the circular trajectory"

that considered the robot in the ADA_b as "suicidal", or more so in comparison to the REA_b. Two other participants perceived the robot in both conditions similarly. Only one participant (515) mentioned that the robot "maybe want[s] some human attention". This was surprising as I hypothesized that the approach to the human, and therefore the edge, would be considered more as a wish to approach the human rather than to leave the table.

Table 3.7.: Robot in ADA_b appeared suicidal.

ID	order	description fraction
503	*B*	"the robot required input not to kill itself which gave me more to do"
505	*B*	"a little suicidal"
511	*A*	"The robot ws initated suicidal", the participant later described the robot in REA_b as "less suicidal"

The environment of the study was designed in such a way that the participants could perceive differences between the robots due to their abilities regarding mastering the different surfaces. To my surprise, only 5 participants addressed that at all. Table 3.8 lists their answers. Four participants recognized that the robot in the ADA_b condition could master the "hill" or "valley" better: two gave credit to the robot in the ADA_b condition, and two others recognized that the robot in the REA_b condition was lacking those abilities. One participant, however, was frustrated that the ADA_b condition got stuck in the "slopy area".

Overall, this was evidence that the learning rate of PI was set well. However, the complexity of the environment did not seem to be helpful for distinguishing the robots.

Table 3.8.: Adaptivity of the different robots.

ID	order	condition	description fraction
506	*B*	ADA_b	"can climb the hill"
507	*A*	REA_b	"robot was unable to move smoothly in the valley and mountain area (black)"
508	*B*	REA_b	"but never succeeded to climb heights"
509	*B*	ADA_b	"attempted to go over difficult terrain"
512	*B*	ADA_b(!)	"it was bit sad when robot couldn't get out of the slopy area and kept looking for strategy"

There were also hints about the randomness of the robot's behavior. One participant (513, *A*) considered the robot in both conditions random. Three others considered the ADA_b as random, or as more random than the REA_b robot (501, 505, 510). Two participants mentioned that the robot in both conditions did not seem to remember obstacles or attempted the same movements (509, 510). Two participants, on the other hand, left statements about the robot's intelligence. One considered the robot in the REA_b condition as "move[ing] intelligent[ly]" (516, *A*). Another said that the robot in the ADA_b "seemed much more intelligent" (502, *B*).

3.5. Discussion

The research questions were concerned with investigating whether the intrinsically motivated, fully autonomous and adaptive robot (ADA_b) has a higher perceived Animacy or Anthropomorphism (RQ1), that the robot is more liked (RQ2) and whether the robot is perceived more competent or intelligent (RQ3) compared to the reactive baseline behavior (REA_b). Neither the quantitative analysis of the questionnaire responses (subsection 3.4.2) nor the qualitative analysis of the open-ended question (subsection 3.4.3) provided clear evidence to answer any of these questions. There were no statistically significant or large effects neither in favor of the intrinsically motivated, adaptive robot, nor for the reactive robot. Most importantly, the dimensions which were the focus of the study (Competence, Animacy, Likeability, Perceived Intelligence) all showed, if at all, small or medium effects for the reactive, baseline robot.

What may have played in favor of the reactive robot was the fact that it only rarely approached the edge that was not enclosed by a wall and where the robot could fall off the table. Initially, I thought that people would feel that the robot tries to approach them rather than trying to fall off the table. However, only one participant mentioned "it may have sought attention". I designed the interaction at the edge so that participants could experience the robot's adaptation to interactions. I hypothesize that the participants' perception of the reactive robot was highly influenced by the fact that it approached the edge less often, and that the robot's reactions were more predictive. These may all be reasons for the higher score of, e.g., Perceived Intelligence and Competence for the reactive robot, but also for the higher perception of Discomfort for the intrinsically motivated robot. Participants may have felt particularly uncomfortable with the act of being alert to prevent the robot from falling. In addition, two participants mentioned they felt uncomfortable using their hands to interact with the robot. This feeling may have been amplified by the intrinsically motivated robot which made it more necessary to interact.

This study design aimed to initiate a human-robot interaction so that participants' responses to the questionnaire reflected their perception of the interaction. The idea was that the task to protect the robot from falling off the table *enforced* interactions, while the question about robot strategies aimed to *encourage* it. A positive aspect of this design was that 15 out of 16 participants pointed out robot characteristics or differences in response to the question about the robot's strategy. This was evidence that the idea to *encourage* interaction with a question regarding the robot strategy made the participants more alert toward the robot's behavior. However, it may in turn have also encouraged a critical observation of the robot which falls off the table more often. The result also revealed, that participants were very alert toward the robot falling off the edge, which was probably the central impression of their judgment. A future study design should avoid *biasing* the participants in this way. An idea which is outlined below was to have a more *plane* environment and only encourage the interaction by asking whether the robots are different, as differences do not imply any higher level of intelligence of a strategy.

The study environment was designed in such a way that participants could observe the robot's competence in mastering the variety of terrains. In particular, it was designed so that the robot could be trapped in a pit, which only an adaptive robot could leave. However, only two participants responded positively to the intrinsically motivated robot's adaptivity in "difficult terrain". Having the robot adaptive enough to leave the pit quickly was not as exciting for the participants as I had anticipated. Furthermore, the quick adaptation that was needed to achieve this made the robot too unpredictable, following the answers that three participants considered the robot more random than the baseline. I assumed that a lower update rate for the adaptive robot would make a difference here.

On the positive side, the missing statistically significant results indicated that the baseline behavior was a successful choice. A concern regarding the baseline was that the fixed parameters may result in too much of a monotonic behavior. The robot responded to most of the sensor input, i.e., perturbations by the environment or humans, by giving more speed to the right than to the left wheel, i.e., circling left. However, the motion was mediated by the balancing controller, which aimed to keep the robot upright and therefore also interfered with the robots' reactions. Initial tests prior to this study showed that this resulted in a behavior that was hard to tell apart from the intrinsically motivated robot. The qualitative results confirmed this. In particular, only two participants mentioned in the open-ended questions that the reactive robot's trajectory was *mostly circling*. This was evidence that the baseline behavior was a fair comparison and was a promising candidate for future studies.

This study focused on the perceived agency of the intrinsically motivated robot. It was hypothesized that the intrinsically motivated robot, which could adapt to the environment or human interactions, would be perceived as having more agency. The study results did not confirm this. However, they did reveal that the intrinsically motivated robot was perceived as more warm. The dimension Warmth is created from the items *Happy*, *Feeling*, *Social*, *Organic*, *Compassionate* and *Emotional*. Together with Competence, it is one of the universal dimensions explaining social attribute formation for human-human interaction. It has been argued that Warmth carries more weight in judgments of interpersonal interaction than Competence (Carpinella et al. 2017). In particular, Warmth predicts the valence of social judgments. For example, if we perceive another human as high in Warmth, we likely judge them with more positive social attributes (Fiske et al. 2007). The reason for the focus on agency at the beginning of the study was that it seemed out of the question that participants would perceive one robot higher in a central social concept such as Warmth; a concept also linked to Friendliness and Trustworthiness (Fiske 2018). The findings therefore motivated a shift of the working hypothesis from a focus on agency to a focus on Warmth. The hypothesis was that if a robot is perceived as more warm, it indicates that we may judge the robot more positively, which in turn may sustain the interaction (cf. section 2.4). However, it was crucial that a future study confirms this finding, as the observed effect in this study was not convincing and not part of a research question.

Given the medium effect for Warmth, it was somewhat surprising that the related dimension Likeability (Carpinella et al. 2017) did not show a similar effect. The development of this effect needed careful observation in future studies[7]. This is also true for the dimension Discomfort. At first glance, it was somewhat concerning that the adaptive robot scored high in Discomfort. However, the dimension is not a central dimension to the social attitude judgment from the dimensions Warmth and Competence (ibid.). As discussed earlier, the high perceived Discomfort may be caused by the extra care needed by participants to interact with the adaptive robot, so that it was not rolling over the edge. It could also be an indication that the intrinsically motivated robot elicited an unknown perception in the human participant, maybe due to the unpredictability of its behavior. From behavioral sciences it is known, for example, that unpredictable and aversive stimuli leads to a more sustained level of anxiety when compared to stimuli that were predictable and aversive (Grillon et al. 2004). In contrast, however, agent behavior research by Bickmore et al. (2010) indicates that a degree of unpredictability in the behavior of an animated human might be essential to sustain engagement of the human participants in longitudinal interactions. In other words, if

[7] It is shown in the final study that Likeability did not measure the same concept as Warmth. In particular, it was found that Warmth is linked to the participant's preference to continue interacting with a robot, but Likeability is not.

a longitudinal interaction is too predictive, the human participants might simply get bored. As already pointed out, a follow-up study would have to change the experimental design and avoid forcing the participants to take more care with one robot than the other.

3.6. Implications for next studies

This section presents the implications for the next studies of this thesis deriving from this chapter. The list below provides an overview of them. Each item is described in detail in the following paragraphs.

i. Focusing on the dimension Warmth.

ii. Change the environment: enclose the table and decrease complexity.

iii. A more game-like scenario for encouraging interaction.

iv. Increase motion capabilities with a different motion model.

v. A tool for interacting with the robot.

vi. Increase the number of participants.

Focusing on the dimension Warmth The results of this study indicated that the intrinsically motivated robot was perceived as less intelligent than the reactive baseline behavior. Given that the intrinsically motivated robot had the capability to adapt to new situations and human interactions, this was initially surprising to me. As discussed, one cause was most likely the study design. However, in retrospect, the IM-driven behavior generation could hardly yield an observable goal which shows particular goal-directedness. The goal of the intrinsically motivated robot was to explore how to create rich sensory input (i.e., exciting its sensors), while at the same time trying to predict the sensory outcome. By design, this behavior might have been interesting to watch and enabled the robot to be adaptive to new situations, but it could not elicit a feeling of Competence in the human observer, given that the robot intrinsically tries to expose itself to new states. Therefore, the main focus on the concept of perceived agency of the robot (RQ3) needed a rethink.

With the increasing insight into social cognition, gained with the work from the current study, the dimension Warmth emerged as a promising tool to understand how to sustain HRI. In fact, Warmth is, next to Competence, the second universal dimension for humans to assign social attributes, but it is primary for understanding whether our attitudes toward peers are *positive* (cf. section 2.4). To make HRI sustainable in the long run means the robots would eventually have to score high for Competence and Warmth. However, positive attitudes are usually assigned to humans scoring high in Warmth. The following study therefore focused on whether an intrinsically motivated robot was perceived as more warm.

Change the environment: enclose the table and decrease complexity The idea toward a new robotic environment began by considering a different experiment introduction. Rather than saying "also, your task is to prevent the robot from falling off the table", it could have been said that "we did not enclose the edge by a wall, so you can better interact with the robot when it is seeking attention, but please take care it is not rolling over the edge". Both introductions would have biased the participant in some way. The first one would have let the participants project Competence on the robot which was less likely to approach the edge. The second one may project more social capabilities on the robot which was *seeking attention*, i.e., approaching the edge more often. This introduction was necessary because the environment was designed with an open edge and the robot could not sense the open edge itself. This way, the environment design caused the participants to project a goal onto the robot, which, depending on the robot's behavior, made the robot appear more intelligent than the other. For example, similarly to what is present here, a robot that is mostly slowly circling could be perceived as careful. In contrast, an explorative robot changing its behavior could be perceived as more careless, as it risks falling over the edge.

The environment needed to meet the sensor capabilities of the robots. The robot had no possibility to prevent itself from falling over the edge. Therefore, it was assumed that an environment where none of the robots could fall over the edge would lessen the participants' bias and would make both robots appear similarly competent and intelligent.

Furthermore, it could be said that the capacity of the robot to leave the pit was not a driving factor for participants when judging the robot's competence positively. The update rate that allowed for that skill may have made the adaptive robot appear unnecessarily random. Further studies were needed to find a good update rate for the robot. In addition to enclosing the wall as mentioned above, the environment was simplified to further concentrate on the interaction part only. It was assumed that a round table would make the robot less stuck and decrease the interpretation of the robot goal. Also, a more *plain* environment would further put the investigation solely on the interaction level.

A more game-like scenario for encouraging interaction Enforcing the interaction by asking the participants to prevent the robot from falling off the table may have caused participant's to focus on the robot's capability to stay on the table. In addition, asking participants to address the robot's *strategy* may have caused them to assume that there was a particular strategy.

The idea for the next study was to avoid any interaction *enforcement*. Instead, it concentrated on *encouraging* the interaction in a way that did not imply the high-level intelligence of a strategy. I implemented a more game-like scenario which allowed for the participants to interact with the robot depending on their enjoyment or curiosity to do so, rather than having to interact without them knowing when it is *intended* by the robot. The aim of the game was that the participants had to figure out whether the robot behaviors were *different* or the same. To this end, participants could observe the robot or interact with it when they felt curious about the interaction, or when they simply enjoyed doing so. The game scenario benefited from the changes to the environment as discussed before. It took place in a completely enclosed environment, where the robot could freely locomote without falling off the edge.

The game-like scenario further directed the focus toward the interaction, without giving away the purpose of the study and leaving the development of the interaction solely to the human and the robot.

Increase motion capabilities with a different motion model The study showed that the chosen platform could already create different behavioral patterns. However, the balancing controller, which was used to keep the robot upright, put a layer between the PI controller and the world, enforcing some position that the robot could not truly influence. In a sense, this balancing controller made the environment and the robot's embodiment more constrained and made it less adaptable to them. For example, the robot itself allows spinning and more wobbly locomotion, but could not achieve this by only changing its heading and speed.

Thus, another potentially beneficial change was to allow the robot to *directly* control its servos, rather than controlling the heading and the speed of Sphero's built-in balancing controller. The idea was that this could yield a larger variety in the behavioral regime. It could also enrich the interaction between the human and the robot, as perturbations would have more direct consequences on the robot's behavior. This also further directed the focus on the level of interaction.

A tool for the interaction with the robot Introducing a tool for the interaction was thought to have two positive effects. Firstly, with the presence of a tool, participants would be implicitly encouraged to interact with the robot, as that was the purpose of the tool given to them. This could increase the interactions and in turn let participants focus more on the interaction itself. Secondly, a tool took away the need to use their hands for interaction. Some participants felt unwell with touching the robot. This possibly increased their rating for Discomfort for the adaptive robot, which needed more intervention as it approached the edge more often. Therefore, a wand-shaped tool was presented in the next study. Next to the aforementioned idea of not enforcing interaction, the idea was that this would further encourage the participants to interact more with the robot.

Increase the number of participants The study revealed some effects that were useful for drawing the implications above. However, none of the effects were statistically significant. This could have two potential reasons. Firstly, the presented effects of this study were not effects, but instead, they only occurred by chance. Secondly, the effects shown were relevant and exist, but the sample size was too small to be 95 % confident. If I were to repeat the same study, I could now perform a power analysis that would reveal the needed sample size, because the effect sizes are now known. Given the number of possible changes discussed above however, and given that the center of observation was more focused on another dimension, I did not conduct a power analysis. With keeping an eye on the practicality of conducting such an experiment, increasing the number of participants from 16 to 24, i.e. by 50 %, seemed reasonable.

The next chapter presents the follow-up study, which takes into account the lessons learned when conducting the study described in the current chapter. All the implications discussed in this section were implemented.

3.7. Conclusion

The study investigated whether an intrinsically motivated, fully autonomous robot is anthropomorphized or perceived more animated (RQ1), is more liked (RQ2) or is perceived as more intelligent or competent (RQ3) than a reactive baseline behavior. The quantitative results did not provide an answer to these questions for either of the robots. Although there were no statistically significant or large effects, the results seemed to indicate that the reactive baseline behavior was perceived higher in all the investigated dimensions. In other words, there was a tendency that the baseline was more liked, perceived as more intelligent, and was also more animated. This chapter argues that this was mainly because of the study design. The participants were enforced to interact with the robot in order to keep it on the table. The chapter suggests that a future study should rather encourage interaction. The idea was to introduce an interaction tool, as well as to create a game-like interaction where participants would answer at the end of the study session whether the robots were different.

What was unexpected was that the results suggested an adaptive, intrinsically motivated robot is perceived as more warm. From social cognition, it is known that a high scoring for Warmth can be found for all positive ratings of social attitudes for interpersonal interactions (Fiske et al. 2007). Thus, Warmth plays an important role in going toward sustained HRI.

In addition, the study's baseline behavior turned out to be a fair comparison. The baseline behavior was pre-adapted for the very same sensors and the very same environment as the intrinsically motivated behavior. Future experiments can adapt the underlying baseline generation for their comparisons.

The general idea of investigating robots solely on their behavioral level was very promising. However, this chapter addressed a few changes future studies would benefit from (cf. section 4.5). The following chapters describe studies that followed a similar design approach but were enhanced with the suggested changes.

Chapter 4.

Study II

4.1. Introduction

This study aimed to provide a further step toward answering the main research question: Can an autonomously, intrinsically motivated robot, sustain the interaction with humans? (RQ2).

In the first study (chapter 3), the focus was on investigating the robot's agency. The idea was that a human engages more with a robot if they perceive the robot high in agency. Therefore, the measures focused mainly on the dimensions Animacy, Anthropomorphism, Perceived Intelligence and Competence. In addition, the dimension Likeability was used to understand if the intrinsically motivated robot is more liked than the reactive baseline. The idea was to understand what kind of robotic behavior humans prefer and what may sustain the interaction between robots and humans. The questionnaires Godspeed and Robotic Social Attribute Scale (RoSAS) were used to measure those dimensions. All dimensions provided by the questionnaires were reported along with the ones above, including Warmth, Discomfort (RoSAS) and Perceived Safety (Godspeed). From this, a medium effect on the dimension Warmth was found. This indicated that the intrinsically motivated robot is perceived as more warm than the reactive baseline behavior.

In contrast to the Godspeed dimensions Likeability, Animacy and Anthropomorphism, the dimension Warmth has not found much attention in human-robot interaction (HRI) at the time of conducting the study. The promising effect found in study I triggered more in-depth research of the social cognition between humans. Social psychology research considers the dimension Warmth, together with Competence, as one of the universal dimensions to measure social attitudes (Fiske et al. 2007; Abele, Hauke, et al. 2016). Fiske et al. (2007) argued that humans perceived high in Warmth are judged positively by their peers. Warmth judgments, in contrast to Competence, "carries more weight in interpersonal interaction" (Carpinella et al. 2017). This means, when characterizing other people, we firstly judge their intent (Warmth) before judging their capability (Competence) to enact their intent. Warmth is linked to the measure of trust (e.g. Fiske 2018). A person who is perceived as warm is also perceived as more trustworthy. For example, Kulms and Kopp (2018) used it as an indicator of people's trust in computers. Humans perceived as warm therefore experience more positive interactions (cf. section 2.4).

This knowledge from social cognition changed the focus of this study and how to measure human perception. This study extended the first study and further investigated whether an intrinsically motivated robot is perceived as more warm. Fiske et al. (2007) points out that the dimension Competence can affect the rating for Warmth. Therefore, this study was designed in such a way that none of the robots were perceived as more competent or intelligent compared to the other conditions.

4.1.1. Research questions

As discussed above, this study built on top of the first study presented in chapter 3 and so did the research questions. The previous study found evidence that the intrinsically motivated, fully autonomous robot is perceived more warm compared to the reactive baseline behavior. This study therefore focused on the dimension Warmth and aimed to confirm this observation. However, there were added questions concerning the changes to the study design, which are presented in detail in the next section. A central aim was to have both behaviors – the baseline and the intrinsically motivated one – perceived to be similarly competent and intelligent. This way the focus could be fully on the dimension Warmth, which was then not influenced by observations of the robot's competence. Lastly, participants were directly asked whether they perceived the robot behaviors in the conditions differently. This helped to provide data about the quality of the baseline behavior. The data of perceived difference provided a ground truth for the development of a different baseline behavior for the final study. To summarize, the following research questions were derived:

RQ1 Is an autonomously, intrinsically motivated robot, which directly controls its servos, perceived as higher in Warmth than the reactive, but balanced baseline behavior?

RQ2 Does a simple environment and a more game-like scenario, which does not force the participants to interact with a robot, but rather ask them to look out for differences, allow for a more unbiased perception of the robot's behavior? In other words, is the perceived intelligence and competence of the intrinsically motivated robot similar compared to the reactive baseline robot?

RQ3 Do participants perceive the two robot behaviors as similar?

4.1.2. Overview

In this study, a robot that used predictive information (PI) maximization and directly controlled its servos was compared to the fully reactive baseline robot from the first study.

Section 4.2 presents the study design. Subsection 4.2.1 describes the environment, the robot, the interaction tool, and the participants' task. The robot platform was again Sphero's BB8. The environment for the interaction in this study was redesigned. The complexity of the environment was decreased by having only one surface, no obstacles and the table had no open edge. Subsection 4.2.2 presents the two conditions: one used the baseline robot from the first study, the other used an intrinsically motivated robot that directly controlled its servos (contrary to study I, where the balanced controller of Sphero was used). For the measures that are presented in subsection 4.2.3, the study used the same questionnaires as before: the RoSAS and the Godspeed scale. There were no open-ended questions this time. Instead, participants answered an additional Likert-type question about how different they perceived the robot behaviors. The procedure is described in subsection 4.2.4. Participants were given little information. Their only task was to understand whether the two robots differed in any way. Subsection 4.2.5 then presents the sample of 24 participants for this study.

Section 4.3 presents the results of the study. Most notably, the participants perceived the intrinsically motivated robot as more warm and liked it more than the reactive baseline behavior. However, there was no evidence that the participants perceived one of the robots as more competent, as intended by the re-design of the study. The results also revealed that the participants perceived the two conditions as quite different and that the intrinsically motivated robot was perceived as more animated. Both were not expected and are critically discussed in section 4.4, and shape the implications for the final study which is presented in section 4.5. One of the main implications that were formulated was the need for a more similar baseline behavior, other than the reactive robot behavior used in this study. Section 4.6 concludes the chapter's results and contributions.

4.2. Study design

This section describes the study design. All suggested changes which were discussed in the first study in section 3.6 were implemented in this study. The next sections describe the study design in detail, which includes addressing the changes made in comparison to the first study.

4.2.1. Robot, environment and tasks

Figure 4.1a shows the robot used in this experiment: the BB8 version of the Sphero robot. The robot and its capabilities were described in detail in section 2.5. This was the same platform that was used in the first study, however, there were changes regarding how the intrinsically motivated robot was controlled, which is presented in the next subsection 4.2.2.

Figure 4.1b shows the table that the robot was placed on. In this study, the table had no open edge, which allowed the robot to freely locomote on the table without falling off. This change in the design was motivated by the previous study. The idea was to reduce the mismatch between the implicit goal assignment and the robot behavior, which occurred in study I: people assumed that the robot which falls off the table less often was more competent. The table was circular with 91 cm in diameter and of 72 cm in height. A foam wall of 2.5 cm *(a) Robot BB8*

(b) Table

(c) The HRI tool.

Figure 4.1.: The images show the environment and tools used for the study. (a) The BB8 version of the Sphero robot. (b) The table where the robot was placed to locomote freely. It was covered with a cloth to increase friction. In contrast to study I, there were no obstacles, different surfaces or an open edge. (c) The HRI tool that was used for interacting with the robot by using the white end.

in height and with 4 cm width surrounded the border of the table. These measurements were determined so that the robot could not fall off the table, even when moving with a very high velocity. Three blankets with a total height of 3 to 4 mm covered the surface (including the walls). This applied some friction, which made it easier for the robot to locomote on the otherwise smooth and slippery surface given by the wooden tabletop. The table's distance to the surrounding walls was at least 60 cm, which allowed the participants to freely move around the table. Figure 4.1c shows the HRI tool referred to as a *wand*. The wand was 50 cm long and weighed 78 g. It consisted of a 40 cm long aluminum tube with a diameter of 10 mm. The black end was where a human could hold the wand. The white end was a soft sphere, which was made of an off-the-shelf table tennis ball with a diameter of 40 mm. This end was for contact with the robot. The wand was a new tool that was not present in the first study. Instead, participants were given a task to encourage interaction: they had to keep the robot from falling off the table. In this study, they were given a tool for encouraging interaction with the robot. The idea was that the presence of the tool made the participants more inclined to use it and thus encouraged interaction with the robot. Figure 4.2 shows the author in the full environmental setup. He uses the tool to nudge the robot, which locomotes on top of the table.

Figure 4.2.: The picture shows the author using the interaction tool. He nudged the robot with the white end. It can be seen that the participants were able to freely choose a position around the table to interact with the robot.

4.2.2. Conditions

The study consisted of two 5-minute-long conditions with the following types of behavior generation:

REA_b **(balanced, reactive):** The robot used its balanced mode for locomotion. The parameters controlling the robot had been pre-adapted using time-local predictive information (TiPI) maximization, but the parameters remained constant during the condition. The binary running the robot and the parameters were exactly the same as in the first study. The resulting behavior was therefore very similar to the one described in section 3.2 and the name of the condition was therefore kept the same.

ADA **(unbalanced and directly controlled, adaptive):** In this condition, the robot was intrinsically motivated and was continuously adapting its parameters, based on TiPI maximization (see θ in subsection 2.3.2). Unlike in study I, the robot did not use its balanced mode for locomotion. Instead, it controls both its servos directly (cf.section 2.6).

The intrinsically motivated robot in the ADA condition realized behavior motivated by maximizing TiPI, and it continuously updated its parameters based on that gradient during the experiment. In contrast to the similar robot (ADA_b) from the preliminary study, the robot controlled its two servos directly. This lead to more diverse robot behavior because a larger variety of servo configurations were possible.

The robot input, as with study I, contained 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. In study I, the robot's orientation was used, which was received via its pitch and roll angles from the inertial measurement unit (IMU). In this study, they were replaced with the input speed of each servo by measuring the servo's back electromotive force (back EMF)[1]. This allowed for a coupling between the output of the TiPI controller for the servo speed, and the actual measured servo speed as the input to the controller.

This time, the starting parameters of the TiPI controller were tweaked by hand. As there was direct coupling between the servo speed readings and the controller output (i.e., the set speed for the servos) the parameters 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, the robot started its behavior by moving straight, just like the reactive robot.

The reactive robot in the REA_b condition started with the same parameters as the robot in the first study. The weights were received based on pre-trial adaption. This determined how it would react to sensor input, but it did not further update its internal networks during the experiment. The reactive robot implementation, including the sensory input and the

[1]The back EMF is a voltage appearing between the armature and the magnetic field of the motor's field coil. It is related to what is also known as the counter-electromotive force (counter EMF, CEMF).

parameters, remained unchanged from the one in the first study. The name for the reactive robot condition therefore remained the same: REA_b. The only behavioral difference was grounded in the change of the environment. The friction of the whole environment in this study was less variable compared to the first study, where the friction of the wooden floor and the foam floor made the robot appear faster and slower (respectively). However, the resulting behavior was very similar to the one from the first study. A video supplementing this study shows an example of both conditions (see Scheunemann 2019; Scheunemann 2017e).

The reason for taking the REA_b robot from the first study was two-fold: firstly, the behavior was evidently a good baseline behavior. The robot was interesting to the participants and the behavior was not too simple, so the participants did not see any reoccurring patterns. In fact, in study I they perceived the robot in the REA_b condition as more intelligent (cf.section 3.4). Secondly, keeping the baseline constant, but changing other variables, allowed a better comparison to the previous findings and the previous adaptive robot.

The type of behavior generation was an independent within-subjects variable, meaning that all participants were exposed to both conditions REA_b and ADA. The order of the conditions was randomized but counterbalanced over the number of participants. This resulted in the two orders A and B:

	order of conditions	participants
A	$REA_b \rightarrow ADA$	12
B	$ADA \rightarrow REA_b$	12

4.2.3. Measures

After each of the two conditions (REA_b and ADA) participants were given a post-questionnaire. Similar to study I, the post-questionnaire consisted of the RoSAS and the Godspeed scale. Godspeed uses a 5-point semantic differential scale for the factors Anthropomorphism, Animacy, Likeability, Perceived Intelligence and Perceived Safety. The authors of RoSAS do not suggest a specific scale, but recommend having a neutral value, e.g. uneven number of possible responses for the Likert-type items. As in study I (cf. subsection 3.3.4), the current study used a 7-Likert scale. The scale tests for the factors Warmth, Competence and Discomfort.

To answer the research questions addressed in subsection 4.1.1, the dimensions Warmth, Competence (both RoSAS) and Perceived Intelligence (Godspeed) were needed. The research question RQ1 focused on the effect of the dimension Warmth. The dimensions Perceived Intelligence (Godspeed) and Competence (RoSAS) were needed to answer RQ2 and to understand whether the game-like study design present here helped to perceive both robots similarly in both dimensions. All the other items offered by the RoSAS and Godspeed scale were included in the questionnaire to hide its intention. These items also allowed for finding possible effects that were not part of the initial research questions. Therefore, all factors are presented.

Other than in the first study, the questionnaires did not include any open-ended questions. However, to have a more robust measure of whether the participants perceive the two conditions differently, they were asked directly at the end of the session: "Was the behavior of the robot different in comparison to the previous interaction?". The responses were collected with a 5-point Likert-type item ranging from 1 ("Not at all") to 5 ("Very much so"). The measure was helpful for three reasons:

i. The data collected provided quantitative insights about the perceived differences between the two conditions and helped to answer RQ3.

ii. The study was designed in such a way that the participants had to understand whether the robots in the two conditions were different. Asking this question at the end of the study masked the intent of the study.

iii. The collected data provided a baseline of perceived differences for the evaluation of a new baseline behavior development, such as the one of the final study described in chapter 5.

The pre-test and post-test questionnaires are attached to this work in section B.2.

4.2.4. Procedure

Participants were welcomed to the experimental room and were then handed an information sheet. They were welcomed 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. It was in the beginning and at this point that it was emphasized that they could leave the study whenever they feel uncomfortable, stressed or bored.

They were then assigned randomly, but counterbalanced, to one of the two orders. To achieve this, they drew a folded snippet from an envelope, which contained 24 pieces. They were not informed about the letter they drew. After that, the study environment and the robots were presented to the participants.

Unlike in the first study, participants were not told that the robot had a specific aim and they were not asked to prevent the robot from falling over the edge. Instead, they were only told that their task was to observe whether the two presented robots (one robot per condition) were different. To understand whether the robots were different, they could use the HRI tool. They were allowed to nudge the robot or block it by using the white end of the wand. Both actions were presented to the participants by the experimenter. Again, no other information was provided.

They were informed that they will be asked whether "the two robots' behaviors are any different" at the end of the experiment. The idea was that their intent to answer the question would motivate them to interact with the robot in order to find out the robots' behavioral differences. They were not given any more details but the means of interactions. In particular, the type of possible differences was not revealed.

Participants then filled in the pre-questionnaire. This gathered information regarding their sex, age and background. Next, the two conditions were presented depending on the order that they were assigned to earlier. Each interaction in one of the two conditions lasted approximately 5 minutes. After each of the two interactions, the participants filled in a post-questionnaire containing the two scales and additional questions discussed earlier in subsection 4.2.3. The entire experiment took about 40 to 50 minutes.

4.2.5. Sample

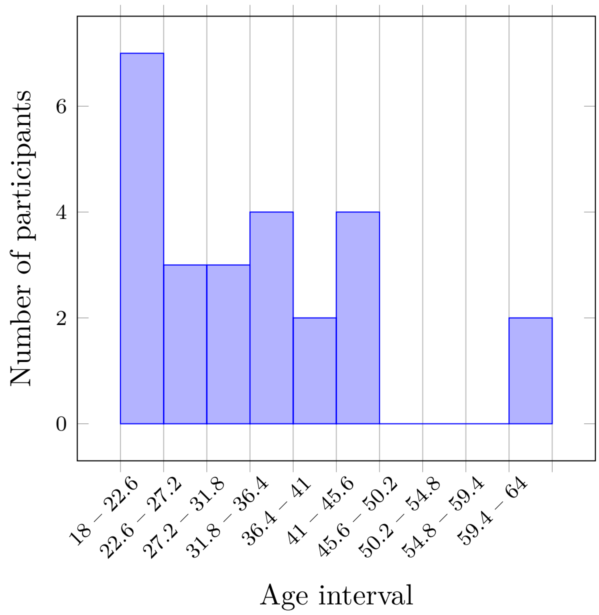

Figure 4.3.: A histogram showing the age of the participants. The participants were evenly distributed over the ages of 22 to 46 years. Overall, the sample was slightly skewed towards younger participants under 22 years old.

I 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). Figure 4.3 shows the age distribution of the participants. Twenty-two participants were students and staff from the university and eight of them had a background in HRI. Seven participants took part in the preliminary study I, whereas nine participants never participated in any prior HRI study. All participants were naïve toward the goal of the experiment. The participants were asked how familiar they were with interacting with robots, programming robots, the chosen robot platform Sphero and the movie series Star Wars. A 5-point Likert scale was chosen with the value 1 for “not familiar” and 5 for “very familiar”. The self-assessed experience for interacting with robots averaged 3.5 (Mode=5). The average familiarity with programming robots was 3.2 (Mode=5) and experience with the chosen robot platform was rated an average of 2.1 (Mode=1). The familiarity with the movie series *Star Wars* was rated 3.2 (Mode=4). Figure 4.4 shows the responses to the questions.

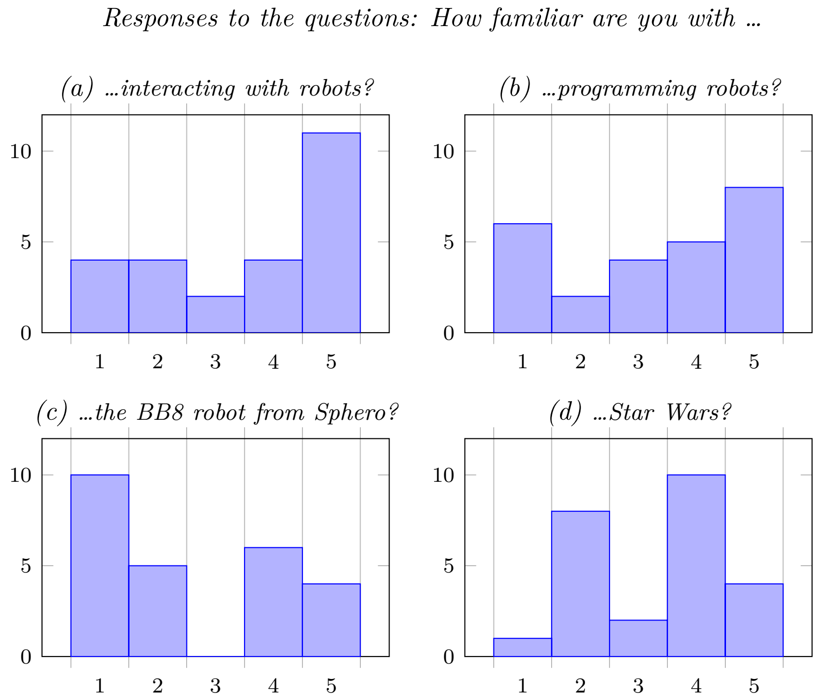

Figure 4.4.: The response distributions of the pre-test questionnaire 5-point Likert-type questions (see section B.2). The scale ranged from value 1 "Not familiar" to value 5 "Very familiar". Participants were asked about their familiarity with (a) interacting with robots, (b) programming robots, (c) Sphero BB8 version and (d) Star Wars. Results indicated that the majority of participants had little experience with the Sphero robot. However, the majority of the people had interacted with robots already.

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 notification of approval is attached to this work in section C.2. The experiments were conducted in February 2019 over the course of 10 days. The anonymity and confidentiality of individuals' data are guaranteed.

4.2.6. Data preparation

The score reliability of the scales of both standardized questionnaires had been tested with the use of Cronbach's α. The item *quiescent-surprised* was negatively loaded on the scale dimension Perceived Safety. Even if reversed, the reliability was poor with $\alpha = 0.54$. The item was therefore removed. Table 4.1 presents all test results, which revealed a good score for reliability ranging from 0.74 to 0.85 and acceptable reliability for the dimension Anthropomorphism: $\alpha = 0.67$. This was evidence that the dimensions could be analyzed without any further preparation.

Table 4.1.: Internal consistency reliability scores.

	dimension	items	α
RoSAS	Warmth	6	0.80
	Competence	6	0.85
	Discomfort	6	0.79
Godspeed	Anthropomorphism	5	0.67
	Animacy	6	0.74
	Likeability	5	0.82
	Perceived Intelligence	5	0.84
	Perceived Safety	2	0.82

4.3. Results

This section presents the results of the study. In contrast to the first study, all results stemmed from a quantitative analysis of the questionnaire responses. The questionnaire responses were analyzed for interaction effects, to understand whether the conditions could be analyzed independently of their presented order. The results that are presented in subsection 4.3.1 showed that there were no interaction effects. Subsection 4.3.2 presents the results of the main effects. They provided evidence that the intrinsically motivated robot was perceived as more warm than the reactive baseline behavior. They further showed that the Perceived Intelligence and Competence of the robots in both conditions were perceived similarly, as intended by the study design.

4.3.1. Interaction effects

An analysis of variances (ANOVA) is commonly used for investigating interaction effects, i.e., effects that show that the order of the conditions influences the responses of participants to a condition. A non-parametric ANOVA-type test was used, due to the relatively small sample size ($N = 24$).

The study had two independent variables: one within-subjects variable and one between-subjects variable. The within-subjects variable, i.e., the independent variable that all participants were exposed to, was the *type of behavior generation* (short: behavior). It consisted of the two condition levels REA_b and ADA (cf. subsection 4.2.2). The independent, between-subjects variable was the *order of the conditions* (short: order). This means that each participant was exposed to only one order. Table 4.2 shows the results of a non-parametric ANOVA-type test[2]. The last column *order:behavior* reveals the likability for an interaction between the conditions and their order. None of the p values (sub-column p) fulfills the

[2]For computing the ANOVA-type test the `R` package `nparLD` was used. As the study consisted of one within-subjects variable (behavior) and one between-subjects variable (order), it could be expressed as F1-LD-F1 Model. The `nparLD` package offers the function `f1.ld.f1()` for computing such models.

Table 4.2.: ANOVA-type test results for the independent variables "type of behavior generation" (level REA_b, ADA), their "order" (level A, B) and the interaction of both variables "order:behavior" for the dimensions of the RoSAS and Godspeed scale. Note that the dimensions Anthropomorphism, Perceived Intelligence and Perceived Safety are abbreviated.

		order			behavior			order:behavior		
	dimension	F	df1	p	F	df1	p	F	df1	p
RoSAS	Warmth	0.098	1	0.755	11.733	1	0.001	0.001	1	0.976
	Competence	0.163	1	0.687	0.047	1	0.828	1.473	1	0.225
	Discomfort	1.365	1	0.243	1.143	1	0.285	1.787	1	0.181
Godspeed	Anthropom.	0.168	1	0.682	14.074	1	<0.001	0.164	1	0.685
	Animacy	0.928	1	0.335	14.565	1	<0.001	0.665	1	0.415
	Likeability	0.786	1	0.375	1.856	1	0.173	1.455	1	0.228
	Perc.Intelligence	0.193	1	0.660	0.043	1	0.836	0.000	1	0.984
	Perc. Safety	10.477	1	0.001	1.343	1	0.246	1.381	1	0.240

criteria $p < .05$. This means that for a 5 % significance level there was no statistical significance and therefore there was not enough evidence for an interaction effect for any of the dimensions. This was particularly true when looking at the dimension Warmth, as the p value was the largest and almost equaled one. This means, the presence of an interaction effect was highly unlikely. This allowed for investigating the main effects between the conditions independently of their order, i.e., the responses to both conditions could be compared independently of whether the participants were exposed to, e.g. ADA, in the beginning of the experiment or at the end[3]. The results are presented in subsection 4.3.2.

A first impression of the main effects could be retrieved by studying the second column *behavior* of Table 4.2 and its three sub-columns. This shows evidence that there were statistically significant effects for the dimensions Warmth, Anthropomorphism and Animacy. The next subsection presents a more detailed analysis of these main effects.

4.3.2. Main effects

The above results showed some interesting effects for the two conditions for the dimensions Warmth, Anthropomorphism and Animacy. A paired difference test could be used to understand the direction, i.e., was the perceived Warmth higher for the condition ADA or the condition REA_b.

The Wilcoxon signed-rank test is a non-parametric candidate known to be robust for small sample sizes. It tests for the null-hypothesis that the two conditions do not differ, i.e., the two-sided test version was used and possible effects in both directions were shown. The test statistic V, the p value, a point estimate, and its corresponding confidence intervals are

[3]Note that the same approach was used to analyze if the participants' self-reports from the pre-test questionnaire or their involvement in a previous experiment had an effect on their responses to the conditions, but no statistically significant effects were found.

reported.

The point estimate (short: estimate) is the median of the differences of the rank. It provides a size and a direction for how much the participants preferred one condition. For example, if the median of the differences for the comparison of ADA and REA_b equaled -0.833, this would mean that on average the participants responded to Warmth with 0.833 units higher in the ADA than in REA_b. The units were the responses to the Likert-type items ranging from 1 to 7 (RoSAS) or the differential scale ranging from 1 to 5 (Godspeed).

Along with the point estimate, another effect size is reported: the standardized effect size r (Rosenthal et al. 1994; Yatani 2016). In a way, it is a normalized p value. The p value is dependent on the sample size. For example, if the effect of something is known to be small, the p value can be further decreased by increasing the sample size. r allows for investigating the size of a potential effect independently of the sample size. Its effect is either small ($r \geq 0.1$), medium ($r \geq 0.3$) or large ($r \geq 0.5$)[4].

Table 4.3.: Main effects for all dimensions of the RoSAS and Godspeed scale for the comparison of REA_b and ADA.

				95 % confidence interval			
	dimension	V	estimate	lower bound	upper bound	p	r
RoSAS	Warmth	27.5	-0.833	-1.333	-0.333	0.007	0.555
	Competence	138.5	0.000	-0.833	0.583	0.988	0.003
	Discomfort	54.0	-0.250	-1.250	0.417	0.287	0.217
Godspeed	Anthropomorphism	26.0	-0.900	-1.200	-0.500	0.002	0.635
	Animacy	30.5	-0.833	-1.250	-0.417	0.002	0.636
	Likeability	38.0	-0.400	-0.700	-0.100	0.038	0.424
	Perceived Intelligence	126.0	-0.100	-0.600	0.500	0.715	0.075
	Perceived Safety	49.0	0.500	-0.500	1.250	0.422	0.164

Table 4.3 shows the results of the two-sided Wilcoxon signed-rank test for all dimensions of the RoSAS and the Godspeed scale. The test compared the condition ADA and REA_b. Using a 5 % significance level, a large, statistically significant effect for the dimension Warmth could be seen ($r = 0.555$, $p = 0.007$). The point estimate, or the median of the differences, was negative. This was because, on average, participants responded higher to the robot in the ADA condition, which made the difference of $REA_b - ADA$ negative. In other words, most participants perceived the robot in the ADA condition as more warm than the robot in the REA_b condition. This directly answered RQ1: an intrinsically motivated robot (as the one in the ADA condition) is perceived as more warm.

Figure 4.5 visualizes the magnitude of the effect, as a means for an alternative result presentation. The magnitude of the effect increased with an increasing distance of the point

[4] There is no mutual agreement on how to verbalize the effect size r. However, the subjective interpretation of Pearson's r by Cohen (1992) has been often used.

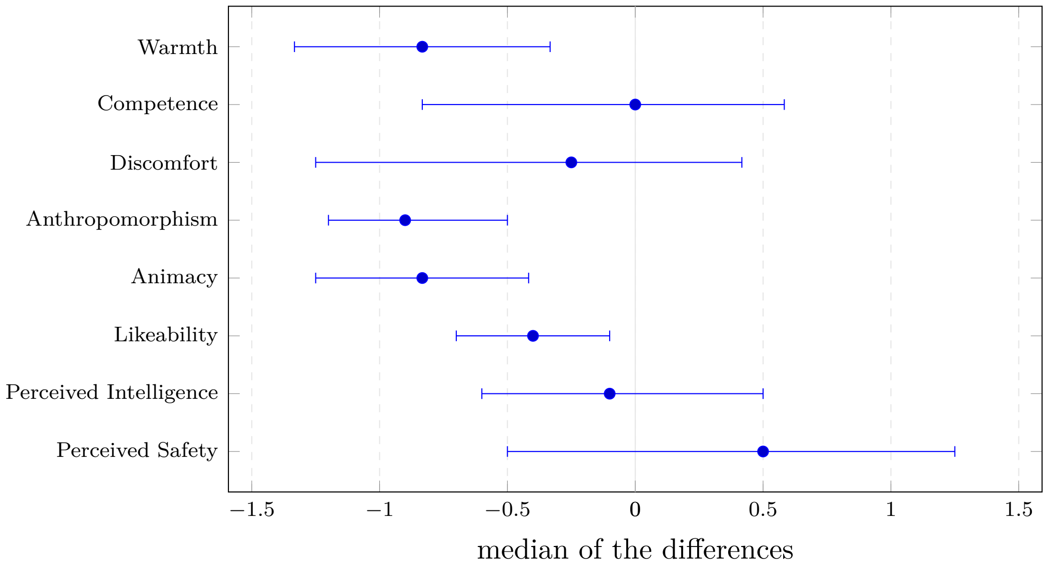

Figure 4.5.: Depicted are the point estimates (median of the differences) between the condition REA_b and ADA. If the estimate is on the left, it indicates that more participants rated the dimension higher in the ADA condition compared to the REA_b condition. The error bars visualize the confidence intervals. It can be seen that there were statistically significant effects for the dimensions Warmth, Anthropomorphism, Animacy and Likeability.

estimate (i.e., the median of the differences) to zero. It also visualizes the certainty of the point estimate. The error bars around the estimate visualize the confidence interval. The narrower the interval is, the more certain it is that the estimate was the true effect. Figure 4.5 confirms that there was a large effect for Warmth in favor of the *ADA* condition.

RQ2 asks: *Does a simple environment and a more game-like scenario, which does not force the participants to interact with a robot, but rather ask them to look out for differences, allow for a more unbiased perception of the robot's behavior? In other words, is the perceived intelligence and competence of the intrinsically motivated robot similar compared to the reactive baseline robot?* The answer to this question could not be given from the results directly and is discussed later. However, given the results shown in Table 4.3, there was no statistical significance for either of the two dimensions Perceived Intelligence or Competence. More importantly, the standardized effect size r was even almost zero and there was no magnitude of an effect. Figure 4.5 visualizes that the median of the differences was close to zero for both of the dimensions. The confidence interval was almost equally distributed around zero and was quite large, indicating that there was no certainty for an effect in any direction. This meant that none of the robots were perceived higher in Competence or Intelligence in agreement by a large enough number of participants, i.e., the participants could not distinguish between the robots on these dimensions. As pointed out above, the implications for RQ2 are discussed in the next section.

The results also indicated a large, statistically significant effect for the two dimensions Animacy ($r = 0.636$ $p = 0.002$) and Anthropomorphism ($r = 0.635,\ p = 0.002$). The point estimate indicated that this large effect was in favor of the *ADA* condition.

There was also a medium, statistically significant effect for the dimension Likeability ($r = 0.424$, $p = 0.038$) indicating that participants liked the robot in the *ADA* more than the baseline. Interestingly, there was also a small effect on the dimension of Discomfort ($r = 0.217$), although the robot in *ADA* was perceived as more warm and was more liked. It felt contradicting at first, but participants could respond high for Warmth and Discomfort at the same time (Carpinella et al. 2017). This was already seen in the first study presented in chapter 3.

4.3.3. Perception of difference

In the first study, the perceived difference was analyzed using the answers to open-ended questions. In this study, the perception of differences between the two conditions was collected with the help of a Likert-type question (cf. subsection 4.2.3). Participants were asked after the second condition: "Was the behavior of the robot different in comparison to the previous interaction?". For the response, a 5-point Likert-type item was provided, ranging from 1 ("Not at all") to 5 ("Very much so").

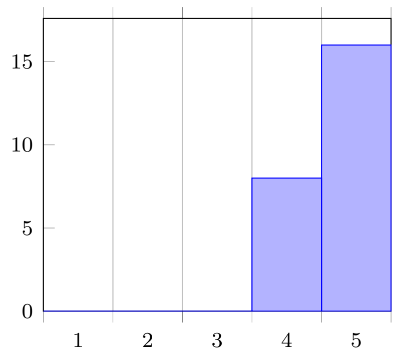

Figure 4.6.: At the end of the study, i.e., after the second condition, participants were asked: "Was the behavior of the robot different in comparison to the previous interaction?". This figure shows the responses to a 5-point Likert-type item, with 1 ("Not at all") and 5 ("Very much so"). It can be seen that the participants perceived the robot behaviors of both conditions very differently.

Figure 4.6 visualizes the responses to the Likert-type item. It can be seen that most participants considered both behaviors very different. The median and the mode of the data were 5, which was associated with the answer: "Very much so". The research question RQ3 concerned the perceived differences between both conditions: Do participants perceive the two robot behaviors as similar? The data was only collected with one Likert-type question and the results had to be interpreted with care. However, almost all participants answered similarly. It could thus be safely said that the behaviors were perceived very differently.

The results indicated that a new, more similar baseline behavior could be useful to confirm the other results of the study. The data here could help with developing and evaluating a new, more similar baseline behavior. The next section discusses the results and their implications for the next study.

4.4. Discussion

The study provided evidence that the intrinsically motivated robot is perceived as more warm than the reactive baseline robot. This is an indicator that a PI-driven behavior may prove relevant for human-robot interaction, as the dimension Warmth is one of the universal dimensions for humans judging social attributes on other humans (cf. section 2.4). Notably, the more we perceive another human as warm, the more we judge them positively and are more likely to interact with them. The results left no doubt that the intrinsically motivated robot is perceived more positively than the reactive baseline robot.

The changes undertaken in this study design toward a game-like scenario helped to focus on the Warmth dimension. Neither the Competence nor the Perceived Intelligence dimension scored high for any of the conditions. This was evidence that the participants did not know if the robot had any goals. Although Competence and Warmth are mainly considered unique dimensions, some interferences between them have been pointed out (Fiske et al. 2007; Abele, Hauke, et al. 2016). The lack of an effect for neither Perceived Intelligence nor Competence was, therefore, an important feature of the study design, which allowed for an isolated observation of the influences of intrinsic motivation (IM) on the perception of Warmth.

In addition, it was – unexpectedly – observed that participants perceived the intrinsically motivated robot as more animated and they anthropomorphized it more. There is evidence that humans perceive a robot higher in Animacy when the robot moves more "naturalistic" (Castro-González et al. 2016). In fact, any object is considered animated if it changes speed and direction without visible influences (Tremoulet and Feldman 2000). Another influence of the perception of Animacy is the reactivity of the robot (Fukuda and Ueda 2010). The baseline behavior was designed to provide both similar movement variety and reaction to sensor input, which allowed for a fair comparison and a focus on the effects of IM. The baseline behavior showed to be feasible in the first study.

However, in this study, the control mechanism for the intrinsically motivated robot was changed (cf. subsection 4.2.2). In contrast to the reactive baseline behavior, where the robot could only move forward and was kept mostly upright due to the balancing controller, the intrinsically motivated 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 such as wobbling locomotion. Therefore, there were three possible explanations for the baseline behavior being perceived as less animated: (i) its different motion patterns, (ii) its reduced reactivity, or (iii) the lack of IM.

With the data presented here, the question could not be answered sufficiently, but I tend to be skeptical and I did not want to argue for (point iii) before carefully observing the baseline behavior. Therefore, the baseline behavior was redesigned for the final study. The changes to create a similarly animated baseline behavior as the intrinsically motivated robot are discussed in the next section (4.5).

However, it needs to be noted that the robot with the baseline behavior was not perceived as inanimate. Instead, participants simply perceived the intrinsically motivated robot as more animated as the baseline. Although this was an indication for the baseline behavior to have less naturalistic movements (as discussed), there is no evidence in the literature that the rating for Warmth is significantly influenced. In the first study, 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, I argue that there is evidence that the different participant responses for Warmth between the two behavior conditions are mainly caused by the robot's IM.

4.5. Implications for the next study

This chapter presented a study design that made the robots in each condition perceived similarly intelligent and competent. This was a major goal and it was important in order to fully focus on the dimension Warmth. It could be said that the changes were successful and they were, therefore, kept constant for the follow-up interaction study. This means, in the final study, the environment, the interaction tool and the game-like scenario were kept the same.

However, the discussion above also presented the limitations of the study design and raised questions that needed addressing in the final study of this thesis. The intrinsically motivated robot was perceived as more animated and as very different compared to the reactive baseline behavior. As discussed above, there were three possible explanations for the baseline behavior being perceived as less animated: (i) its different motion patterns, (ii) its reduced reactivity to perturbations, or (iii) the lack of IM.

In order to strengthen the evidence that IM caused the perception of Warmth (point iii), this section presents two changes to control for the other two possible influences. Firstly, a baseline behavior needed to be developed which was more similar to the behavior generated by the intrinsically motivated robot (point i). Secondly, allowing the robot to distinguish between human and environmental perturbations could show how much its reactivity toward human perturbations influences the perception of Warmth (point ii).

A baseline behavior more similar to the adaptive robot The baseline behavior in this study was carried over from the first study. The reactive baseline behavior and the intrinsically motivated robot were perceived differently in the first study already, and yet, there were no statistically significant effects. This resulted in the argument that the baseline behavior was a good candidate and, therefore, it was the reason that the full baseline behavior was kept and all parameters were constant.

However, the intrinsically motivated robot behavior changed. The implementation of the IM formalism was kept, but some sensors were changed and, more importantly, the motion model was modified (cf. subsection 4.2.2). The intrinsically motivated robot directly controlled its servo speed, while the reactive baseline behavior still used the balanced motion controller. This could have resulted in the differing perception of Animacy. A naïve first solution to get a more similar baseline behavior could be to simply change the motion model for the baseline to the direct control. This would mean that the reactive baseline behavior would directly control its servos by using pre-adapted weights. This however would cause the robot to have a very monotone, repetitive motion pattern, most likely just circling or spinning. Note that this was present already in the current baseline behavior, although it has not been noted (cf. subsection 3.4.3). As discussed, the robot had a tendency to drive in anti-clockwise circles, but the built-in balancing controller made this pattern less prevalent. Another idea was to have a behavior that was still reactive to the user input but had changing random weights. Again, this yielded the challenge of how to design the randomness as discussed in section 3.2.

The idea which was pushed forward was to replay the networks' weight updates of a randomly picked, previous run. Thus, the robot could use the same motion control, the same sensor input and the same initial model parameters, and produce a variable, reactive behavior. More importantly, it would provide the participant with a sense of adaptation. This would not be a true adaptation, as the network updates would not be influenced by any means of current interactions. In a sense, it could be coined *fake adaptation*. I hypothesized that participants would perceive such a baseline behavior as more similar to the intrinsically motivated robot. This could be inferred by two quantitative measures: (i) the perception of Animacy being the same and (ii) the assessed perception of difference being smaller compared to the data collected in this study. This way, if the effect for Warmth persisted for the intrinsically motivated robot compared to such a baseline, the evidence provided by this study could be further underlined, namely that IM yielded the high perceived Warmth.

Use sensors to distinguish between human and environmental perturbations Another important difference between the current baseline behavior, and the intrinsically motivated, adaptive robot, was the reactivity of the robot toward perturbations. As discussed before, the baseline robot kept the weights constant. It reacted to the human perturbations, but its reactions were constrained by the balancing controller. On the other hand, the intrinsically motivated robot directly changed its servos, and perturbations by the human participant yielded immediate changes in its behavior. It could be possible that this played a role in the positive perception of Animacy, Anthropomorphism and – more importantly – even for Warmth.

The final study ruled out that this described "sense of control", meaning the participant seeing direct behavior changes in the robot, played a role for their perception of Warmth. To investigate this, the idea was to equip the robot with sensors which enabled it to perceive human participants in such a way that their perturbations caused different actions than the perturbations induced by the environment. Such a sensor would provide the participants with a similar feeling of direct control and would enable the robot to behave differently to human perturbations than to environmental perturbations. Importantly, such a sensor would allow for controlling whether the direct robot response was the cause for the high perception of Animacy and Warmth.

In the current study, only proprioceptive sensors updated the parameters of the IM. Therefore, from the robot's perspective it was not directly clear whether the perturbations were induced by the human participants or an obstacle in the environment, such as a wall. The solution which was investigated in section 2.7 was to exploit Bluetooth Low Energy (BLE), adding proximity information to a small, mobile, spherical robot. BLE can scan fast for signals and, most importantly, the sensor system could be added unobtrusively to the robot in the final study. This way, the robot's capability to scan for human participants' proximity was hidden from the participants. This is important because providing as little information as possible about the robot's capabilities to the participants was a central design paradigm for all studies of this thesis.

Each of the above two implications were addressed in the final study of this thesis. This way it could be answered whether an intrinsically motivated robot (like the one presented in the current study) is perceived as more warm because (i) of it being intrinsically motivated, or (ii) its pattern of locomotion or (iii) its capability to directly react to human perturbations.

4.6. Conclusion

This study extended the first study and was likewise motivated by the question of whether intrinsically motivated autonomous robots can be beneficial for designing engaging HRI. However, the focus of the investigation was changed to the dimension Warmth, which further triggered design changes for the current study. This chapter presented a within-subjects study ($N = 24$). The participants interacted with a fully autonomous Sphero BB8 robot with two conditions with different behavioral regimes: one realized an adaptive, intrinsically motivated behavior and the other was reactive, but not adaptive. TiPI maximization was used as one candidate measure for IM. Of particular interest was the high similarity between both conditions in Perceived Intelligence ($r = 0.032,\ p = 0.875$) and Competence ($r = 0.003,\ p = 0.988$), which gave support to the non-task-oriented interaction design. This was particularly important as Competence ratings can influence the perception of Warmth, which is the dimension that the study focused on.

The main result was that the perception of Warmth by human participants was high for the adaptive, intrinsically motivated robot ($r = 0.555,\ p = 0.007$). This was in comparison to a baseline behavior that included both similar movement and reaction to sensor inputs – meaning that the difference in perception arose from the robot's adaptation to the physical interaction. This effect was also robust to physical interaction, i.e. it was present even though the robot was physically nudged by the human participants.

The open questions which motivated the design of the final study were the following. Firstly, can the results for Warmth be confirmed with the use of a more similar, fakely adaptive baseline behavior? Secondly, were the direct responses of the intrinsically motivated robot, which directly controls its servos, the reason for the high perceived Warmth, or was it due to the robot being intrinsically motivated?

Chapter 5.

Study III

5.1. Introduction

The last studies showed that a robot which is intrinsically motivated yields an increase in the participants' perception of the robot's Warmth. This started with medium effects that were seen in the first study and those effects were confirmed in the second study with statistical significance. These promising results, however, were accompanied by open questions about what caused the effects and corresponding suggestions to better analyze their origin. This chapter presents the third and concluding study of this thesis, addressing the suggestions outlined in the previous chapter.

One observation from the previous study was that the intrinsically motivated robot was not just perceived as more warm than the reactive baseline behavior, but it was also perceived as more animated and more anthropomorphized. All these effects were statistically significant and, at first glance, evidence that intrinsically motivated autonomy affects human perception.

However, it has been shown that robots, even objects, are perceived high in Animacy if they are moving, and the influence of their locomotion changes are not visible to the observer (Tremoulet and Feldman 2000). This was why the statistically significant results, which were at first glance very promising, raised skepticism. Maybe the baseline behavior of the second study was too different from the intrinsically motivated robot's behavior? The results to the question of whether or not participants considered the two robot behaviors to be different supported this concern since most participants perceived the robots as maximal different.

The participants' sole task of study II was to understand whether the robots behaved differently in the two conditions. This may have caused participants to be particularly alerted or biased towards minor differences between the robot behaviors. More likely, however, was that the effect of perceived Warmth was caused by the use of two different motion controls. The reactive baseline behavior used Sphero's built-in balancing controller together with the reactive controller: the robot updated its heading and speed based on previously explored, constant parameters, and the current sensory input. This resulted in speed and heading information which the robot tried to reach while the balanced controller kept the robot upright. In contrast, the intrinsically motivated robot applied computed parameter updates (the output values from maximizing predictive information (PI)) directly to the servo speed and direction. This way, perturbations by the human participants had a more direct effect on the behavior of the robot, which in turn may have influenced the perception of Warmth by the participants. The previous chapter therefore encouraged further investigations to understand whether the promising results remain in a study when using a similar baseline behavior.

In summary, the aforementioned observations raised the question of whether an intrinsically motivated robot was perceived as more warm compared to the baseline, because (i) it was intrinsically motivated, (ii) due to its pattern of locomotion, or (iii) based on its capability to directly react to human perturbations. In light of the main research question, the hypothesis was that the perception of Warmth was influenced by the robot's intrinsic motivation to explore the spatial relationship to humans and objects in its environment (i.e., point i). This chapter presents the last study of the thesis, which had an adapted study design to answer these questions.

The robot platform, the tool, the environment, the procedure, and the scales that were used remained unchanged compared to the second study. However, two major changes were introduced. Firstly, the baseline behavior was designed in such a way that the motion patterns were similar to the intrinsically motivated robot. The idea was for the baseline to use the same direct control model as the intrinsically motivated robot. For this to work, the baseline robot had to update its parameters based on updates of a previous, randomly chosen interaction. These updates were changes to the robot parameters, but they were not based on maximizing time-local predictive information (TiPI) and therefore resulted in a robot that was not goal-directed or adaptive to the environment or human. However, the goal of the intrinsically motivated robot that used TiPI maximization was not obvious to the participant, since there was no known observable task that the robot needed to fulfill. This was why the difference between the truly goal-directed, intrinsically motivated robot and the robot *replaying* parameter updates of a previous interaction were very subtle, in particular for interactions of short duration[1]. I therefore refer to the *replaying behavior* as being *fakely* adaptive. The hypothesis was that these changes caused the baseline to be perceived as more similar – in particular more similarly animated – when compared to the intrinsically motivated robot. In other words, the baseline became more challenging to distinguish. If the positive effect of Warmth persisted, despite these efforts, then point (ii) from above could be ruled out, namely that the Warmth perception was due to the different motion patterns and because the two robot behaviors were perceived very differently.

The next major change addressed the observation that the perception of Warmth could have been influenced by the capability of the robot to directly respond to human perturbations. To address this concern, this study had an additional condition. The study used the

[1] A video supplementing this study shows an example of all three conditions conducted by one, randomly chosen participant (e.g. Scheunemann 2017e).

same intrinsically motivated robot as in the second study, but had an additional intrinsically motivated robot, which had the added capability of perceiving human proximity. That way, point (iii) could be addressed: was it more important that participants were closely entrained with the robot and therefore saw a direct response from the robot to their actions, or did the perception of Warmth truly stem from the goal-directed, adaptive behavior enabled by the robot's intrinsic motivation (IM)?

If the study could find evidence that the major drive behind the positive effect for Warmth perception from the second study stemmed from the very fact that the robot was intrinsically motivated, it would give more weight and meaning to the results of the previous studies. The working hypothesis throughout the thesis is that the theory about Warmth from social cognition holds true for human-robot interactions. Thus, the humans perceived as high in Warmth experience more positive social interaction. If this transfers to human-robot interaction (HRI), then the findings show that IM can be a key to sustain HRI,

This directly raised the last question that is addressed in this chapter: does the theoretical knowledge about Warmth transfer to HRI? The concept of Warmth and its implications is still an active area of research in social cognition, but there has not been any evaluation yet of whether it transfers to physical *interaction* between humans and robots. This study showed the first step toward understanding whether it does or not. In this study, participants were directly asked after the session that if they could interact with a robot again, which robot (if any) would they prefer. This study then analyzed whether the responses were related to any of the scale dimensions. The study showed that Warmth and Competence – the two central dimensions for human attitude formation – were both related to the participant's preference, and that, interestingly, the dimension Likeability failed to deliver this relation.

In summary, this study concluded the series of interaction studies of this thesis. Its results showed that the intrinsically motivated robots are perceived as more warm, because they are intrinsically motivated, and not because of different motion controls or a specific capability to directly react on human input. In particular, the chapter rules out the alternative explanations for the effect on Warmth and therefore provides more meaning to the previous studies, where it was shown consistently that an intrinsically motivated robot is perceived as more warm. Furthermore, the study results provided the first evidence that the knowledge of Warmth, which discriminates whether humans perceive positive social interaction, transfers to physical human-robot *interaction*.

5.1.1. Research question

The research questions motivated by the introductory section are as follows:

RQ1 Is a reactive robot which is replaying weight updates from a previous run, (i.e. fake adaptivity) a more challenging baseline behavior compared to the one used before in study I and study II (cf. 3.2)? In other words, is it perceived more similar and similarly animated?

RQ2 Does the study design allow for an unbiased perception of the robot's Warmth, i.e. is the robot's Perceived Intelligence and Competence the same between each pair of conditions (similar to RQ2 from study II)?

RQ3 Is an intrinsically motivated robot perceived higher in Warmth than a reactive and seemingly adaptive baseline behavior (similar to RQ1 from study II)?

RQ4 Is a robot perceived as more warm because of being intrinsically motivated, or because its capacity to directly respond to human interaction (with the help of a proximity sensor)?

RQ5 Is the behavior a participant perceives highest in Warmth the same behavior they prefer to interact with again?

5.1.2. Overview

This study compared three different robots: a robot with a baseline behavior and two intrinsically motivated robots. The intrinsically motivated robots were the same, except one used an additional proximity sensor. The baseline behavior was more *challenging* compared to the one from the second study: it used the same direct motion control and its motion patterns were similar compared to the intrinsically motivated robots.

Section 5.2 introduces the design of the study. The environment and the robot (5.2.1), the measures (5.2.3) and the procedure (5.2.4) stayed the same compared to the second study. The two changes to the study design were the additional proximity sensor which was used by one of the intrinsically motivated robots and the baseline robot. These changes are described in 5.2.2 and 5.2.1 respectively.

Section 5.3 presents the results. This study showed that an intrinsically motivated robot was perceived as more warm compared to a similarly locomoting, reactive baseline behavior. The major driving factor behind this was indeed the IM formalism, and not the potential capability of the robot to respond to human proximity directly. Furthermore, this study presented evidence that the participants' responses to the dimension Warmth were related to their preferred condition, which in turn is a first step to understand if the theoretical knowledge from social cognition transfers to HRI. Section 5.4 discusses the results and confirms that the findings provided more meaning to the results of the previous studies and that they helped to answer the main research questions of this thesis. Section 5.5 discusses the limitations of this study and provides directions for future work. Section 5.6 then concludes the study.

5.2. Study design

The following subsections describe the design of the study. Most of the study design remained the same compared to the second study. This section focuses on the differences, e.g.,

Figure 5.1.: The experimental environment showing the robot platform Sphero in its BB8 version and a person using a tool (the wand) to interact with the robot. The robot could freely locomote on the table. The participant could move around the table to observe and interact with the robot. In contrast to the second study, the robot had the capability to sense the proximity to the tip of the wand using Bluetooth Low Energy (BLE).

describing the reasoning and implementation of a new, fakely adaptive baseline behavior, the use of an additional proximity sensor, and the measuring of participants' preference.

5.2.1. Robot and environment

From an observer perspective the robot platform, the environment and the interaction tool remained exactly the same as in the second study (cf. subsection 4.2.1). As a reminder, Figure 5.1 shows all the components and the author (portraying the role of a participant) interacting with the robot.

In this study, the robot was equipped with an additional sensor to receive proximity information of interacting human participants. The sensor used BLE and derived proximity estimates based on measuring the received signal strength (RSS) between two BLE devices. This sensor was introduced to address RQ4 and was described in detail in section 2.7. In this study, the robot carried a BLE beacon in its head which increased the head's weight by ∼5 g. The wand was equipped with a BLE scanner, which scanned for the beacon of the robot. The scanner retrieves the RSS and sends it to the robot controller.

However, feeding the signal directly into the robot as an additional sensor input was not possible, because RSS values are prone to (i) partial occlusions resulting in sudden drops (Schwarz et al. 2015; Ahmad et al. 2019) and (ii) strong fluctuations of the measurements (Faragher and Harle 2015). Both characteristics violate two requirements of TiPI maximization discussed in subsection 2.3.3. The sensor and the receiver were placed in the robot's head and in the tip of the wand-shaped HRI tool. This way, sudden drops were reduced because occlusions were prevented when the human participant interacted with the robot. In addition, the sensor readings were pre-processed to avoid strong fluctuations.

It is known that RSS measurements are rarely overestimated but rather underestimated (Schwarz et al. 2015). The pre-processing therefore consisted of a filter, which passed only the maximum value of a time window of past readings, which smoothed the RSS reading values and filtered false, underestimated readings. Empirically, it was found that a time window of 300 ms was a good candidate: sudden drops were reduced but the readings were still provided in a timely and continuous manner to the robot.

5.2.2. Conditions

This section describes the condition differences, the realization of the baseline behavior and the order of conditions. As in the second study, the same robot platform was used in all conditions. The conditions were the levels of the independent variable which described the type of *behavior generation*. The robots' behavior generation differed either by the realization of adaptivity (e.g. using an IM formalism or fake adaptation) or by the sensory input (e.g. using a proximity sensor or not). Note that the adaptivity generation and the sensory input could have formed two individual independent variables. However, this would have resulted in an additional condition which does not help to answer any of the research questions, but instead increases the session time for participants by 20 to 30 minutes.

Table 5.1.: Overview of the three experimental conditions

condition	sensor input	parameter update
ADA	no proximity sensor	online adaptation with TiPI
REP_{LE}	proximity sensor	adaptation based on replay
ADA_{LE}	proximity sensor	online adaptation with TiPI

Table 5.1 provides an overview of the differences as per condition. The study consisted of three conditions called ADA, REP_{LE}, and ADA_{LE}. The ADA condition was exactly the same as the one used in the second study presented in chapter 4. The only difference was that the robot in this condition was also equipped with, but did not use, the proximity sensor. This allowed the experimenter to treat all the robots the same and, at the same time, keep the condition hidden from both the experimenter and the participant. ADA_{LE} was almost the same as the ADA condition, the difference being that it had an additional input for proximity information using BLE. This means, both robots were intrinsically motivated and their behaviors were realized with maximizing TiPI. REP_{LE} was the baseline behavior of this study. It had exactly the same input as the intrinsically motivated robot in the ADA_{LE},

however, the behavior was generated by repeating weight updates from an earlier, randomly chosen run. A video supplementing this study shows an example of all three conditions conducted by one, randomly chosen participant (see Scheunemann 2017e). From observing the three conditions alone, it can be seen that the different behaviors looked quite similar (as intended). The following subsections describe the changes in the sensor input and the parameter updates in more detail.

Sensor input

All controllers received sensor readings from an accelerometer, a gyroscope, and the two servo motors. The gyroscope provided the angular velocity around the z-axis. If the robot was upright, this axis was perpendicular to the surface. The accelerometer provided the linear acceleration along the forward and sideward axes. Each of the two servos provided their current speed, which were negative for backward motions. In condition REP_{LE} and ADA_{LE}, the controller had an additional input: a one-dimensional proximity sensor which corresponded to the distance of the interaction wand. This way the robot was able to *distinguish* between perturbations by the environment or by the participant. Proximity changes could be induced quite rapidly by the participants. This was different to some of the other sensors. If the robot was nudged, there would be a peak on the accelerometer reading, but that would also only result in a short behavioral change. For proximity, this was different. A change in the wand's position could change the robot's behavior more visibly and rapidly. In a sense, the robot was pushed into a different behavior regime more easily. This could have different effects, for example, with the observable changes the participant could feel that the robot was responding more directly to their input. The effect however was hypothesized to be smaller than the one observable by comparing an intrinsically motivated robot to a baseline behavior. This design helped to address the research question RQ4.

Parameter updates

In condition ADA and ADA_{LE}, the robot was equipped with a computational model of IM. The update rules for the parameters (weights and biases of a neural network) were implemented by TiPI (cf. section 2.3). In a way, the robot tried to excite different sensors through the generation of a variety of motion regimes, but in a predictable way. For example, the robot spun around to excite the gyroscope or accelerated to excite the forward acceleration measured by the accelerometer. On the other hand, in the baseline condition REP_{LE} the controller of the robot was not updated by TiPI, but by replaying parameter updates of an earlier run with a predictive information controller. This means that it changed its network weights and biases, but it was not adaptive toward the current environment or the current participant, i.e., it did not adapt to maximize its PI based on its experience.

This means the robot was reactive towards the current sensory input in all conditions, but the update of the network weights happened either by TiPI (ADA and ADA_{LE}) or by replaying (REP_{LE}) the adaptation of parameters that occurred in a previous experiment. Overall, the regimes of generated behaviors were hypothesized to be very similar (RQ1), alas not adaptive toward the environment in REP_{LE}. This design helped to address RQ3.

Order of conditions

The type of behavior generation was an independent within-subjects variable, i.e. all participants were exposed to the three conditions ADA, REP_{LE} and ADA_{LE}. The order of the conditions was randomly assigned and counterbalanced. The study had 36 participants. This meant that all possible 6 permutations of ordering the 3 conditions occurred 6 times. Table 5.2 shows the order of the conditions and the assigned number of participants.

Table 5.2.: Order of conditions

	order of conditions			participants
A	ADA	$\rightarrow REP_{LE}$	$\rightarrow ADA_{LE}$	6
B	ADA	$\rightarrow ADA_{LE}$	$\rightarrow REP_{LE}$	6
C	REP_{LE}	$\rightarrow ADA$	$\rightarrow ADA_{LE}$	6
D	REP_{LE}	$\rightarrow ADA_{LE}$	$\rightarrow ADA$	6
E	ADA_{LE}	$\rightarrow ADA$	$\rightarrow REP_{LE}$	6
F	ADA_{LE}	$\rightarrow REP_{LE}$	$\rightarrow ADA$	6

5.2.3. Measures

The measures in this study were very similar to the ones used in the previous interaction study. A questionnaire was handed out to the participants after each of the three conditions. The questionnaires of the study are attached in section B.3. They encompass all items of the Robotic Social Attribute Scale (RoSAS) and the Godspeed scale in randomized order. The research questions addressed in subsection 5.1.1 partially relied on the responses to the scale dimensions Warmth and Competence (both RoSAS), as well as on the scale dimensions Animacy and Perceived Intelligence (both Godspeed). However, all items of the two scales were used and the results to all scale dimensions are reported. This helped to hide the questionnaire intention and, at the same time, the dimension results may be helpful to design future studies or enable future studies to compare their results with the results of this study.

There were two changes to the questionnaire compared to study II: one question modification and one additional question. Firstly, like in the second study, the post-questionnaire contained a question to directly assess the perceived difference between conditions: “Was the behavior of the robot different in comparison to the previous interaction?”. The responses were collected with a 5-point Likert-type item ranging from 1 (“Not at all”) to 5 (“Very much so”). The measure was introduced in study II to have a comparison for the development of a new baseline behavior (like the one in this study) and to understand if the baseline was perceived more similar. In contrast to the previous interaction study with two conditions, this study had three conditions. Therefore, after the third condition, participants were asked about their perception of the difference between the previous *and* the first interaction.

Secondly, the questionnaire has one additional question in contrast to the previous study, which directly assessed which condition participants preferred to interact with again: "If you could interact with one of the robots again, which one would you chose?". The participants could choose between one of the three interactions 1, 2, and 3, but they could also tick "no preference". This question was primarily introduced to control whether the dimension Warmth and the participant's preference was dependent. If so, this would give more weight to the results of previous studies where an intrinsically motivated robot was perceived as more warm (a detailed discussion follows).

5.2.4. Procedure

This section describes the study procedure. Participants were welcomed to the experimental room and were then handed an information sheet. They were welcomed 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. It was in the beginning and at this point that it was emphasized that they could leave the study whenever they feel uncomfortable, stressed or bored.

They were then assigned randomly, but counterbalanced, to one of the six possible orders. To achieve this, they drew a folded snippet from an envelope which contained 36 pieces. They were not informed about the number they drew. The participants then filled in the pre-questionnaire. This gathered information regarding their gender, age and background. In parallel, the experimenter prepared the order of conditions and transferred the drawn number to the controlling computer. The numbers had been pre-mapped to a condition order and the experimenter was unaware of the particular order. This avoided the experimenter biasing the participants because participants could not guess the experimenter's intention from explanations or subconscious expressions[2].

After that, the study environment and the robots were presented to the participants. The participant's task followed the one from study II exactly: they were told that their task was to understand whether the three presented robots (one robot per condition) were different. To understand whether the robots were different or not, they could use the HRI tool. The participants were told they were allowed to nudge the robot, and the action was presented to the participants by the experimenter. The participants were then asked to take the interaction tool and try out the action with a passive robot. This study's main purpose was not to investigate the differences in the robots' behaviors. Instead, the main motivation

[2]Orne (1962) described this participant's bias, which he coined *demand characteristics*. It describes that participants may try to perform particularly *well* in order to satisfy the perceived needs of the researcher.

of this task was to encourage interaction with the robot.

No other information or details were provided. In particular, the type of possible behavioral differences of the robots were not revealed nor was any hint of the robot's underlying algorithm given. This was particularly contrasting to the first study, where participants were told that the robot had the aim to explore the environment and they were asked to prevent the robot from rolling over the edge and falling off the table. In this study, the participants could freely decide the amount of interactions they have, and the style of interactions (for example using the tool to nudge or push the robot). There was no specific task provided, like in study II. Therefore, this study confirmed, when answering RQ2, that participants were not encouraged to distinguish the robots by their competence.

Next, the three conditions were presented to the participants depending on the order they were assigned to earlier (i.e. in a randomized order). Again, the order of conditions was also unknown to the experimenter. Each interaction lasted approximately 5 minutes. After each of the three interactions, the participants filled in a post-questionnaire containing the two scales and additional questions presented earlier in subsection 5.2.3. The entire experiment took about 60 to 75 minutes per participant.

5.2.5. Sample

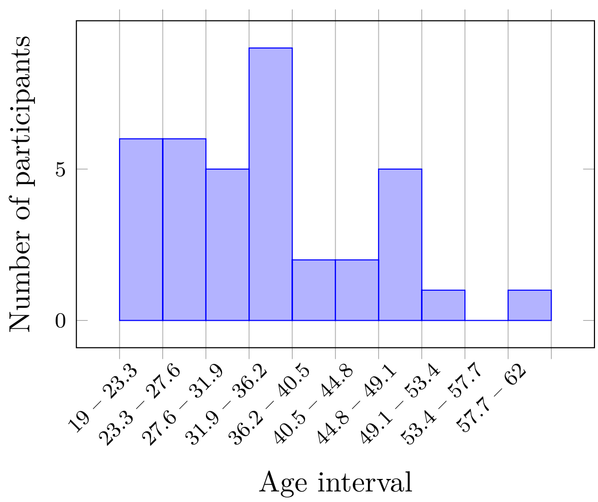

Figure 5.2.: A histogram of the participants' age. The majority of participants were between 19 to 36.2 years old. The figure also shows there was a wide variety of ages.

Thirty-six participants (11 female; 24 male) were recruited for the study, with ages ranging from 19 to 62 years ($M = 33.6$, $SD = 10.2$). Figure 5.2 shows the distribution of participants' ages. The participants were mostly recruited from university staff and students (24) and the majority had a background related to Computer Science (22). Nine participants took part in one of the two previous studies (Six in study I; Nine in study II).

To gain insight into their experience with robots, the participants were asked how familiar they were with interacting with robots, programming robots and the chosen robot platform. A 5-point Likert scale was chosen with the value 1 for "Not familiar" and 5 for "Very familiar". Figure 5.3 shows all the responses to the self-assessment. The majority of participants were familiar with the *Star Wars* movies (Lucasfilm Ltd. 2015), i.e. the response distribution is skewed to the right (d). The flyers that were handed out to the participants may explain this distribution. There was a picture of the used robot platform on them, which resembles a character from the movie series, and participation was maybe most appealing to people who were familiar with the two movies that were released in 2015 and 2017 (both in December). Another skewed distribution can be noted for the responses to the familiarity with the robot used in the study: most people were not familiar at all (c). The responses to the participants' experience with programming robots showed that participants' self-assessed experience was almost balanced (b). There is a slight skewness towards the participants' familiarity with interacting with robots (a). However, when leaving out the participants who felt very familiar with interacting with robots, the participant responses were balanced. Overall, the participants could be considered technophile, but the majority did not work with robots, nor did they have experience with the specific robot platform.

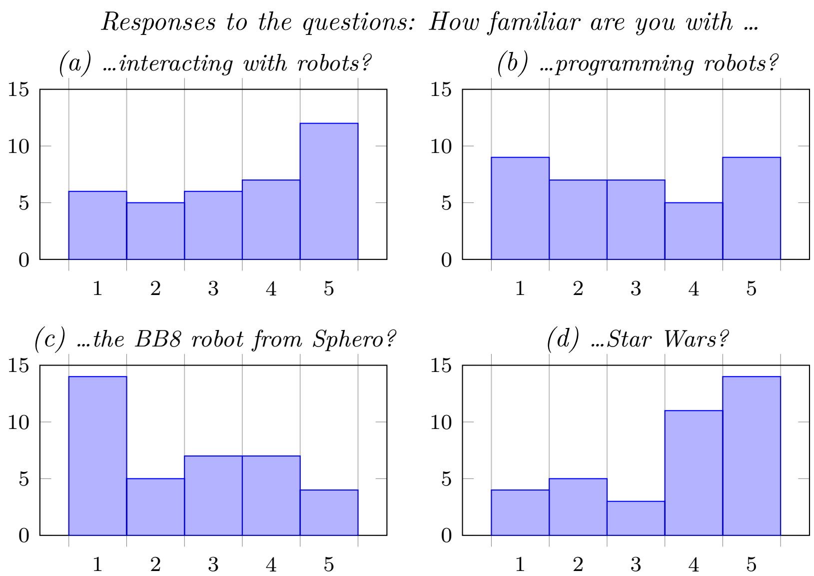

Figure 5.3.: The response distributions of the pre-test questionnaire 5-point Likert-type questions. Participants were asked about their familiarity with (a) interacting with robots, (b) programming robots, (c) Sphero BB8 version and (d) Star Wars. The majority of participants had little experience with the Sphero robot. However, the majority of people had interacted with robots already.

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(4). The experiments were conducted from March to June 2019 over the course of 88 days. The anonymity and confidentiality of the participants' data are guaranteed.

5.2.6. Data preparation

The score reliability of the scales of both standardized questionnaires was tested with the use of Cronbach's α. The item *quiescent-surprised* was negatively loaded on the scale dimension Perceived Safety. Even if reversed, the reliability was poor. The item was therefore removed. Table 5.3 presents all test results, which revealed a good score reliability ranging from 0.79 to 0.92. This was evidence that the dimensions could be analyzed without any further preparation.

Table 5.3.: Internal consistency reliability scores.

	dimension	items	α
RoSAS	Warmth	6	0.86
	Competence	6	0.92
	Discomfort	6	0.79
Godspeed	Anthropomorphism	5	0.82
	Animacy	6	0.82
	Likeability	5	0.87
	Perceived Intelligence	5	0.88
	Perceived Safety	2	0.85

5.3. Results

This section presents the results of the study. Similar to the second study of this thesis, but in contrast to the first study, all results stemmed from quantitative analyses of the questionnaire responses to standardized scales and a few extra questions about participants' preferred interaction and their perceived differences in the robot.

Subsection 5.3.1 presents the results of how different the participants perceived the three conditions. The results provided evidence that all conditions were perceived quite similar to each other. In particular, the results showed that the new baseline behavior was perceived as more similar when compared to the second study.

Subsection 5.3.2 presents the effects on the questionnaire dimensions between conditions. The results showed that the Perceived Intelligence and Competence of the robots in both conditions were perceived similarly, which was intended by the study design and confirmed the results from the second study. The results also showed that there was no effect for Animacy or Anthropomorphism between any of the conditions, which indicated that all conditions, in particular the new baseline behavior, were perceived as similarly animated. Despite the similarity, the results provided evidence that the intrinsically motivated robot was perceived as more warm than the reactive and seemingly adaptive baseline behavior.

Subsection 5.3.3 presents the results collected with the new item of the questionnaire which asked for the preferred condition of the participants. The results indicated that the preferences were spread throughout the conditions, with a tendency toward people who preferred the intrinsically motivated robots. Interestingly, the results showed that the reported participant's preference was statistically significantly dependent on the condition with their highest response to the Warmth dimension.

5.3.1. Perception of difference

This section presents the results of the collected responses to the participants' perceived differences in the robot behaviors. The perceptions of differences between the three conditions were collected using the Likert-type question which asked about the perceived difference to the previous interactions (cf. subsection 5.2.3). Figure 5.4 shows the frequency distributions of the responses. Note that the responses to the differences are depicted independently of their order. This means, for example, Figure 5.4a contains all responses for the perceived difference between the conditions ADA and ADA_{LE}, irrespective of which condition was presented to the participants first or last.

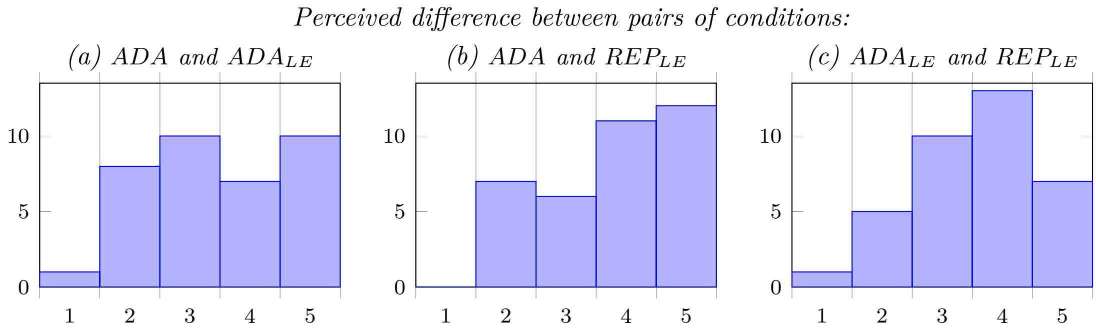

Figure 5.4.: Histograms of the responses to the Likert-question about the difference to previous interactions. The responses (x-axis) were 1 ("Not at all") and 5 ("Very much so"). The response counts are depicted independently of the order, i.e., it does not matter whether the participants conducted ADA before or after ADA_{LE}. The figure indicates that ADA and ADA_{LE} were perceived as the least different.

At first glance, it could be seen that only two responses of two participants[3] considered a pair of conditions to be "not at all" different. This means that all other participants saw some differences between the conditions. This was possibly due to the task given to the participants: they had to answer whether the robots differed or not. To answer this, they were shown how to use the tool. This task and the showcasing of nudging the robot was all the information that was given to the participants. The task was provided to encourage the participants to interact with the robots and to hide the intention of the study. However, this question about the perceived difference may have biased them towards (i) believing the robots were different, or (ii) made them more alert to minor differences. This meant that the overall associated level of difference may have limited explanatory power. Instead, comparing the response distribution between pairs may offer further insights.

[3]The two responses with 1, i.e. "not at all", belonged to two different participants.

From all the pairs depicted in Figure 5.4 the conditions ADA and ADA_{LE} (a) appeared to be perceived as the least different since the responses were more evenly distributed in comparison to all other pairs. This indicated that the two intrinsically motivated robots were perceived most similar, despite their differences in their embodiment (i.e. whether they had a proximity sensor or not). This was a first indicator in answering RQ4, which asks whether the capability of the robot to respond directly to human perturbations (i.e. having a proximity sensor) has more influence on the perception of Warmth compared to the robot's IMs.

To quantify the observations, a Kruskal-Wallis rank-sum test[4] was conducted to examine the perceived differences according to the pairs of conditions. No significant differences ($\chi^2 = 1.38,\ p = .5,\ df = 2$) were found among the three pairs (a), (b), and (c). Therefore, the main result from the responses to the one Likert-type question was that there was no evidence that any of the conditions were perceived as more different than another condition. In particular, there was no evidence that participants perceived the baseline behavior similar to the behavior of the intrinsically motivated robots.

To address RQ1, the question which remained unanswered was whether participants perceived the baseline as *less* different to the intrinsically motivated robots, compared to the difference of the second study. Figure 5.5 helped to answer this question by depicting the relative frequency distribution. This way, the frequencies of this study with 36 participants was comparable to the frequencies of the second study with 24 participants.

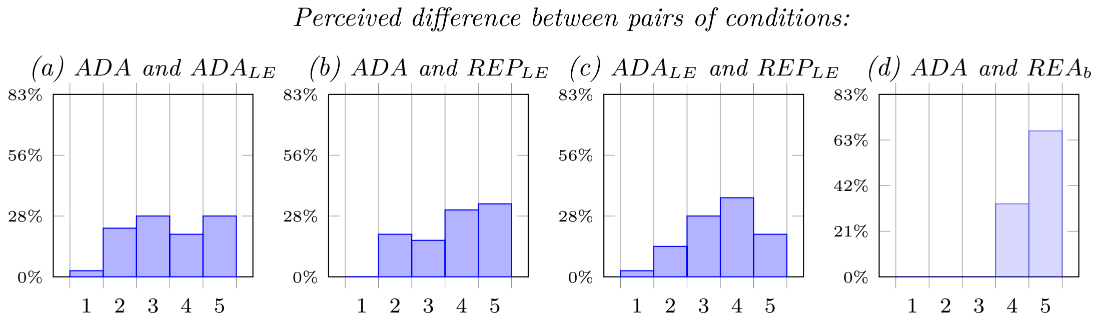

Figure 5.5.: A comparison of the perceived differences between conditions of this current study (first three figures) to the one of study II (d). Depicted are the relative frequency histograms, which allowed for a comparison of the studies with different sample sizes. It can be seen that participants perceived all three condition combinations of this study as more equal compared to the two conditions of study II. In particular, participants perceived the condition ADA more similar to the current baseline REP_{LE} (b) than compared to the baseline behavior REA_b of study II (d).

The first three plots show the frequency distributions of the unordered pairs: (a) (ADA, ADA_{LE}), (b) (ADA, REP_{LE}), and (c) (ADA_{LE}, REP_{LE}). Note that the labeling for the pairs is consistent with the ones in Figure 5.4. The last plot (d) is the histogram of the two conditions ADA and REA_b from the second study. The condition ADA from this study

[4]The Kruskal-Wallis rank-sum test is part of base `R`'s built-in `stats` package and was implemented as `kruskal.test()`.

and the second study remained the same, except that the robot was additionally equipped with the (unused) proximity sensor (cf. subsection 5.2.1). REA_b was the condition with the reactive baseline behavior used in study II[5].

Figure 5.5d shows that the distribution is skewed to the far right, which means that participants of the second study responded mainly with "very much so" to the question of whether the baseline behavior in REA_b was different compared to the intrinsically motivated robot in the *ADA* condition. For all possible pairs of this study, the responses were more evenly distributed between the answers from 2 to 5. This means that participants perceived any pair of this study more similar compared to the pair of study II. In particular, the *ADA* condition was perceived more similar to the baseline condition REP_{LE} than compared to the baseline condition REA_b.

The Kruskal-Wallis rank-sum test (as above) was conducted to quantify the aforementioned observation and to see if there was any difference among the perceived differences between pairs of this study and the pair of the previous study. The test results ($\chi^2 = 20.28,\ p < 0.001,\ df = 3$) showed that there was at least one statistically significant difference between the four pairs (a), (b), (c), and (d).

A Wilcoxon rank-sum test with continuity correction is a special case of the above test for two samples, and it was used to understand further if (d) was statistically significantly different to all other pairs of the current study. The test is a non-parametric version of the popular *t*-test. It can operate on ordinal data and is useful for small sample sizes. The results of the Wilcoxon rank-sum test[6] indicated that the median of ranks of (d) were statistically significantly higher than the median of (a) ($W = 676,\ p = 0.0001$), (b) ($W = 628,\ p = 0.0015$), and (c) ($W = 700,\ p < 0.0001$).

In summary, the results indicated that all conditions of this study were perceived more similar to each other than the two conditions of study II. In particular, when comparing the condition *ADA* to the current baseline behavior REP_{LE}, the conditions were perceived statistically significantly more similar than when comparing *ADA* to the baseline behavior REA_b from study II.

Given that the study design (e.g. the environment, the robot platform, the procedure and the task) remained the same as in study II, this was strong evidence that the current baseline behavior was more similar to the intrinsically motivated robot behavior and therefore even more *challenging*. The results directly answered the first part of RQ1: the new baseline behavior, in contrast to the baseline behavior from study II, was perceived more similar to the behavior of the intrinsically motivated robots.

[5]The condition REA_b with the reactive baseline behavior had also been used in study I, but the perceived difference was derived qualitatively in study I rather than quantitatively in study II and the current one, which made the results incomparable in a straightforward fashion.

[6]The Wilcoxon signed-rank test is part of `R`'s built-in `stats` package and is implemented as `wilcox.test()`.

5.3.2. Effects of questionnaire dimensions

This section presents the results for investigating the main effects of the questionnaire dimensions. This mainly concerned the dimension Warmth (RQ3), Perceived Intelligence and Competence (RQ2) and Animacy (RQ1). All other dimensions of the RoSAS and Godspeed scale are also reported to allow comparisons between possible future studies and this work.

Interaction effects

An analysis of variances (ANOVA) is commonly used for investigating interaction effects, i.e., effects that show that the order of the conditions influences the responses of participants to a condition. The study had two independent variables: one within-subjects variable and one between-subjects variable. The type of *behavior generation* (short: behavior) was the independent within-subjects variable. It consisted of the three levels: the conditions ADA, REP_{LE}, and ADA_{LE}. The between-subjects variable was the *order* of how the three conditions were presented (cf. subsection 5.2.2).

Table 5.4.: ANOVA-type test results for the "type of behavior generation" (level ADA, ADA_{LE}, REP_{LE}) and "order" (level A–F) for all the dimensions of the RoSAS and Godspeed scale. Note that the dimensions Anthropomorphism, Perceived Intelligence and Perceived Safety are abbreviated.

		order			behavior			order:behavior		
	dimension	F	df1	p	F	df1	p	F	df1	p
RoSAS	Warmth	0.151	4.592	0.974	3.454	1.972	0.032	1.683	6.018	0.120
	Competence	0.581	4.026	0.678	0.474	1.994	0.622	1.444	5.466	0.199
	Discomfort	0.462	4.752	0.795	3.334	1.944	0.037	2.363	6.195	0.026
Godspeed	Anthropom.	0.858	4.373	0.496	0.807	1.863	0.439	1.456	6.588	0.182
	Animacy	1.113	4.574	0.350	1.211	1.727	0.294	1.277	6.340	0.262
	Likeability	1.070	4.751	0.374	6.809	1.842	0.002	2.012	6.389	0.056
	P. Intelligence	0.835	4.136	0.506	0.333	1.924	0.708	1.665	5.686	0.129
	P. Safety	0.328	4.651	0.885	0.546	1.867	0.567	1.889	5.421	0.086

Here, a non-parametric ANOVA-type[7] test was used, due to the relatively small sample size ($N = 36$). Table 5.4 shows the results of the test. The last column *order:behavior* reveals the likability for an interaction between the conditions and their order. It shows that there was a statistically significant interaction effect for the dimension Discomfort ($p = 0.026$). The dimension was not part of any research question, therefore a post-hoc analysis of the interaction effect was omitted and the chapter does not discuss the dimension further.

For all other dimensions, and for a 5 % significance level, there was no statistical significance

[7]For computing the ANOVA-type test the `R` package `nparLD` was used. As the study consisted of one within-subjects variable (behavior) and one between-subjects variable (order), it can be expressed as F1-LD-F1 Model. The `nparLD` package offers the function `f1.ld.f1()` for computing such models.

for an interaction effect[8]. Therefore, the results allowed investigating the main effects between the conditions independently of their order, i.e., the responses to both conditions could be compared independently of whether the participants were exposed to, e.g. ADA_{LE}, in the beginning of the experiment or at the end. The results are presented in the next subsection.

A first impression of these main effects can be retrieved by studying the second column *behavior* of Table 5.4 and its three subcolumns. The dimensions Competence, Perceived Intelligence and Animacy were important for answering the research questions. They all do not show any hint for a potential effect (which was intended by the study design). On the other hand, the column reveals a statistically significant condition effect ($p = 0.032$) on Warmth, one of the dimensions this study focused on. This was a very promising indicator that parts of the study design had an effect on this central dimension. The next subsection presents a more detailed analysis of the main effects. In particular, an analysis of which condition created the statistically significant effect for the dimension Warmth.

Main effects

This section presents the main effects between pairs of conditions. The main focus to answer the research questions lied on Warmth and Competence (both RoSAS), as well as Animacy and Perceived Intelligence (both Godspeed). A paired difference test could be used to understand between which pairs of conditions was the effect present and to understand the direction of that effect. For example, it could answer whether the perceived Warmth was higher for the condition ADA_{LE} compared to the condition REP_{LE}.

The *Wilcoxon signed-rank test* is a non-parametric candidate for a paired difference test. In contrast to the popular t-test, it is known to be robust for small sample sizes and can operate on ordinal data. It tests for the null hypothesis that the two conditions do not differ, i.e., the two-sided test version was used and possible effects in both directions were revealed. The test statistic V, a point estimate, and its corresponding 95 % confidence interval are reported, along with a p value and the standardized effect size r (Rosenthal et al. 1994; Yatani 2016). The point estimate (short: estimate) is the median of the differences. It provided a size and a direction for how much the participants preferred one condition. For example, if the median of the differences for the comparison of ADA_{LE} and REP_{LE} equaled 0.417, this means that on average the participants responded to Warmth with 0.417 units higher in the ADA_{LE} condition than in the REP_{LE} condition. The units here are the responses to the Likert-type items that ranged 1 to 7 (RoSAS) or the differential scale which ranged from 1 to 5 (Godspeed). The standardized effect size r is either small ($r \geq 0.1$), medium ($r \geq 0.3$) or large ($r \geq 0.5$)[9].

[8] Note that the same approach was used to analyze if the participants' self-reports from the pre-test questionnaire or their involvement in a previous experiment had an effect on their responses to the conditions, but no statistically significant effects were found.

[9] There is no mutual agreement on how to verbalize the effect size r. However, the subjective interpretation of Pearson's r by Cohen (1992) has been used often.

Table 5.5.: Wilcoxon signed-rank test results and effect sizes between all pairs of conditions.

(a) ADA_{LE} and REP_{LE}

	dimension	V	estimate	95 % confidence interval: lower bound	95 % confidence interval: upper bound	p	r
RoSAS	Warmth	412.5	0.417	0	0.667	0.049	0.328
RoSAS	Competence	330.0	0.167	-0.383	0.667	0.578	0.093
Godspeed	Anthropomorphism	371.5	0.100	-0.100	0.400	0.354	0.154
Godspeed	Animacy	300.5	0	-0.333	0.333	0.959	0.009
Godspeed	Likeability	256.5	0.200	-0.200	0.400	0.399	0.141
Godspeed	Perceived Intelligence	273.0	0	-0.400	0.400	0.893	0.022
Godspeed	Perceived Safety	91.0	0	-0.750	0.750	0.809	0.040

(b) ADA and REP_{LE}

	dimension	V	estimate	95 % confidence interval: lower bound	95 % confidence interval: upper bound	p	r
RoSAS	Warmth	310.0	0.167	-0.250	0.500	0.389	0.144
RoSAS	Competence	290.5	0.167	-0.333	0.667	0.620	0.083
Godspeed	Anthropomorphism	298.5	0	-0.200	0.300	0.747	0.054
Godspeed	Animacy	208.0	-0.083	-0.417	0.167	0.432	0.131
Godspeed	Likeability	386.5	0.500	0.300	0.800	0.002	0.529
Godspeed	Perceived Intelligence	322.5	0.100	-0.200	0.400	0.452	0.125
Godspeed	Perceived Safety	97.0	0.500	-0.250	0.750	0.326	0.164

(c) ADA and ADA_{LE}

	dimension	V	estimate	95 % confidence interval: lower bound	95 % confidence interval: upper bound	p	r
RoSAS	Warmth	196.0	-0.250	-0.583	0.083	0.131	0.252
RoSAS	Competence	299.5	0.083	-0.417	0.417	0.734	0.057
Godspeed	Anthropomorphism	144.5	-0.100	-0.500	0.300	0.628	0.081
Godspeed	Animacy	220.5	-0.100	-0.500	0.167	0.416	0.136
Godspeed	Likeability	345.0	0.500	0.200	0.800	0.006	0.460
Godspeed	Perceived Intelligence	333.0	0.200	-0.200	0.600	0.196	0.215
Godspeed	Perceived Safety	94.0	0.250	-0.750	1.250	0.405	0.139

Table 5.5 shows the results of the two-sided Wilcoxon signed-rank test for the condition pairs (a) ADA_{LE} and REP_{LE}, (b) ADA and REP_{LE}, and (c) ADA and ADA_{LE}. The r column reveals that there were fewer large and medium effects as compared to the second study (cf. 4.3.2). This did not come as a surprise given that the baseline behavior was expected to be perceived much more similar to the other behaviors (the previous section provided some evidence for this expectation). Figure 5.6 combines the three plots from Table 5.5 together in one figure. All three pairs of conditions are plotted per dimension in order to provide a better overview of the present effects.

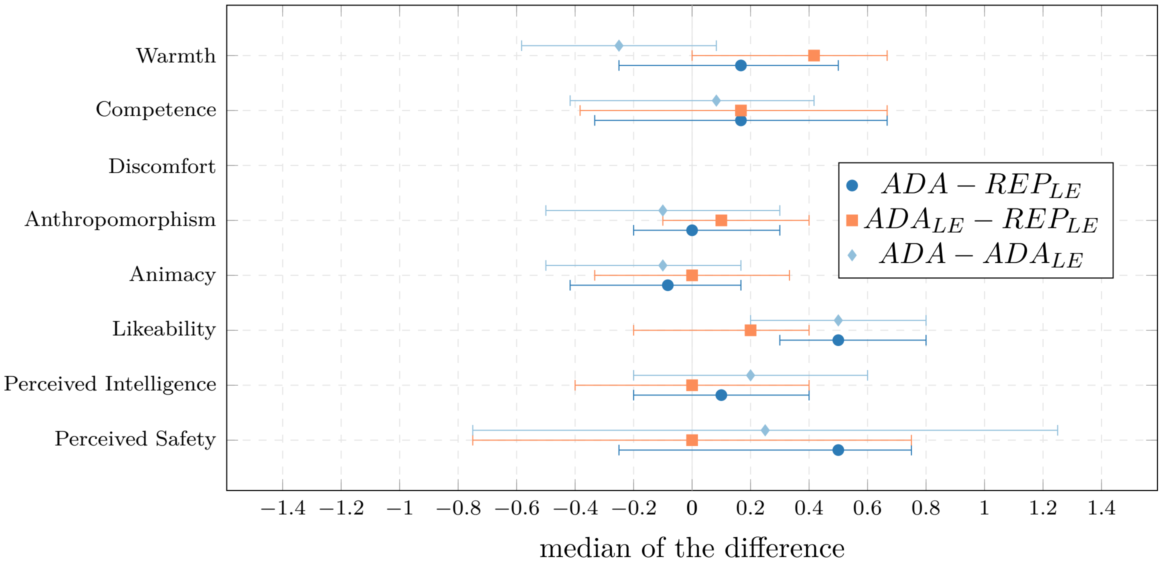

Figure 5.6.: The point estimate (median of the difference) between all three condition pairs is depicted per dimension. This shows the direction of the effect. For example, $ADA_{LE} - REP_{LE} > 0$ (red) indicates that more participants rated the dimension higher in ADA_{LE} compared to REP_{LE}. The error bars visualize the 95 % confidence interval. If 0 is not included, then a statistically significant effect was present (Warmth and Likeability).

All research questions relied to some extent on the results of the questionnaire dimensions. The research questions are addressed individually for overview purposes.

RQ1 *Is a reactive robot which is replaying weight updates from a previous run, (i.e. fake adaptivity) a more challenging baseline behavior compared to the one used before in study I and study II (cf. 3.2)? In other words, is it perceived more similar and similarly animated?*

The results presented in subsection 5.3.1 showed that participants perceived the current baseline behavior as much more similar compared to study II. The results for the dimension Animacy answered the second part of the question. Figure 5.6 shows that, in contrast to study II, there was no statistically significant effect for Animacy. The participants did not seem to perceive one robot as being more animated than the other. As a reminder, in study II the directly controlled, intrinsically motivated robot (ADA) was perceived (with statistical significance) as more animated compared to the balanced, controlled and reactive baseline behavior (REA_b). The result was unexpected and one assumption was that the high perceived Animacy might have stemmed from the use of two different motion controls. This also raised the question of whether the two different motion controls could have had an effect on the likewise promising results on the warmth dimension. This motivated the design of a new baseline behavior, which used the same motion control as the intrinsically motivated robot. The results of this study showed that the new baseline behavior was perceived similarly animated compared to the intrinsically motivated robots. Together with the observed difference by the participants (5.3.1), the results answered RQ1 directly: the new baseline behavior was more challenging compared to the one used in the second study.

However, it was noticeable that if there were small effects for Animacy for any of the pairs, the robot with the proximity sensor was perceived more animated. If both robots had the proximity sensor (a), there was virtually no effect ($estimate = 0,\ r = 0.01$). This indicated that the proximity sensor had some influence on the perception of Animacy, which is evidence that any change of the robot's morphology needs to be considered carefully (I come back to this later).

RQ2 *Does the study design allow for an unbiased perception of the robot's Warmth, i.e. is the robot's Perceived Intelligence and Competence the same between each pair of conditions (similar to RQ2 from study II)?*

Figure 5.6 provides a quick answer for the dimension Competence. For all three pairs of conditions the confidence intervals spanned a similar area of the positive and negative side. From Table 5.5 it can be further seen that there were no effects for any of the pairs ($r < 0.1$). This meant that there was no evidence that participants perceived any of the robots as more competent than the other. The same could be observed for the dimension Perceived Intelligence. However, Table 5.5 reveals a small effect (with no statistical significance) that participants may have perceived ADA as more intelligent compared to the conditions with robots that used the proximity sensor. On the other hand, there was virtually no effect when comparing ADA_{LE} and REP_{LE} (a) on this dimension ($estimate = 0,\ r = 0.02$). Interestingly, this was similar to the above paragraph about the dimension Animacy, but with a changed direction.

However, altogether the results showed that the participants did not perceive any of the robots as more intelligent or competent than the other, which answered RQ2: the study design made the robots appear similarly competent and intelligent, and it allowed for an unbiased perception of the dimension Warmth.

RQ3 *Is an intrinsically motivated robot perceived higher in Warmth than a reactive and seemingly adaptive baseline behavior (similar to RQ1 from study II)?*

There were two conditions with an intrinsically motivated robot. In one of the conditions, the robot had proximity information of the human participant (ADA_{LE}) and in the other it did not (ADA). When comparing the robot in the ADA_{LE} condition to the reactive, fakely adaptive baseline behavior (which also had proximity information of the human participant), there was a statistically significant, medium effect that the intrinsically motivated robot was perceived as more warm ($estimate = 0.417,\ r = 0.328,\ p = 0.049$). What was quite striking was that this effect was by far the largest of any of the dimensions reported in Table 5.6a. However, comparing ADA to REP_{LE} did not provide a similarly convincing result. There was only a small and non-statistically significant effect for the robot in the ADA condition having been perceived as more warm ($estimate = 0.167,\ r = 0.144,\ p = 0.389$).

The results showed that the intrinsically motivated robot was perceived as statistically significantly more warm compared to the more challenging baseline behavior, given that the morphology of the robots was similar (i.e. both robots used the proximity sensor). When the intrinsically motivated robot was not able to sense the human proximity, there was only a small effect that this robot was perceived as more warm.

RQ4 *Is a robot perceived as more warm because of being intrinsically motivated, or because its capacity to directly respond to human interaction (with the help of a proximity sensor)?*

The results presented earlier indicated that the robot's capability to perceive the participant's proximity had small effects on its Perceived Intelligence and Animacy: the robot with proximity information was perceived as less intelligent, but more animated. Both effects were small and not statistically significant, but interestingly, they both seemed dependent on the proximity sensor.

This could not be confirmed for the dimension Warmth. Instead, both intrinsically motivated robots (either with or without the proximity sensor) were perceived as more warm. However, the proximity information seemed to play some role here too. Firstly, the effect for Warmth in favor of the intrinsically motivated robot was larger (and statistically significant) if the proximity information was present. Secondly, there was a small, non-statistically significant effect that ADA_{LE} was perceived as more warm than ADA.

These observations allowed to answer RQ4. They showed that even such a small embodiment change like adding one additional sensor to a robot had an effect. Here, the sensor capability of a robot that could derive information on whether a human was approaching or not was affecting participants' perception. However, independently of the proximity sensor, both intrinsically motivated robots were perceived as more warm when compared to the baseline. This is evidence that the major drive of the Warmth perception is grounded in the robot being intrinsically motivated and not in using a proximity sensor.

5.3.3. Preferences

This section presents the results of the participants' responses to their preferred condition. The data was collected with a Likert-type question at the end of all three conditions. The participants were asked which interaction they would choose if they were to interact with the robot again. They could either answer with one specific interaction or they could say that they had no preference (cf. 5.2.3).

The responses were meant to address two research questions: firstly, it should provide further evidence that the baseline behavior was very similar to the behavior of the intrinsically motivated robots (RQ1). Secondly, it allowed investigating whether a high perception of Warmth was an indicator for people's preference (RQ5).

Table 5.7.: Preferred condition.

ADA	ADA_{LE}	REP_{LE}	No preference
11	13	7	5

Table 5.7 shows the response frequencies from all 36 participants. It can be seen that the majority of participants (31) preferred one specific condition. Seven participants even preferred to interact again with the baseline behavior[10].

At first glance, it seemed that participants preferred interacting with either of the intrinsically motivated robots in either the ADA or the ADA_{LE} condition. This would be in itself a very supporting result for the main research question of this thesis. A very popular quantitative test to analyze categorical data for their goodness of fit is the Chi-squared test. It tests the null hypothesis that there is no difference between the observed frequencies (Table 5.7) and the expected ones. The expected frequencies here were that all conditions were preferred with the same proportion. This way the test helped to find out whether participants answered randomly or did not have a preferred condition.

An issue here was the possibility to answer "no preference". This response could mean that the participant was indecisive, i.e., they had more than one preference. It could also mean that they did not wish to interact with any of the robots again. The responses were therefore ambivalent. However, the decision to include that choice of response was to give indecisive participants an option to respond rather than choosing a distinctive interaction, which in turn would have blurred the results. Participants who were indecisive towards one interaction were left out.

The Chi-squared test results ($\chi^2 = 1.81,\ df = 2,\ p = 0.41$) showed that there was no evidence that the observed and expected frequencies differed, or in other words, the response frequencies from Table 5.7 could well have stemmed from a random distribution, or from a distribution where the participants simply preferred any of the conditions equally.

However, the results supported the findings for RQ1 from earlier, namely that the baseline behavior was similarly perceived as the intrinsically motivated robots. The very fact that participants did actually prefer the baseline further confirmed these findings. In addition, the results also confirmed the previously presented results that the proximity sensor did not

[10]The question directly asked participants for their preference, which may have caused them to want to answer for a specific condition in order to please the experimenter. This limitation is discussed later.

play a significant role in shaping participants' preference (RQ4).

It needs to be noted that the test results had to be considered with care, as Chi-squared test results are dependent on the number of participants. For example, if 3.4 times more participants would answer in the same proportions as in Table 5.7, the above test would have yielded statistically significant results. This limitation was known prior to conducting the study. However, a sample size of an estimated 80 to 100 participants was not feasible. Therefore, by design, the results mentioned so far only provided additional support for the results of the aforementioned subsections.

The remaining question now was whether the preferred conditions of the participants were likewise the conditions where they responded the highest to the dimension Warmth. If such a dependency existed, this would directly answer RQ5: *Is the behavior a participant perceives highest in Warmth the same behavior they prefer to interact with again?*

Model participant's preference

A few new variables are introduced to better address the problem. Let $\mathcal{C}$ be the set of all three conditions and $\mathcal{D}$ be the set of all scale dimensions. Then there are two variables dependent on the participant's responses:

$$o \in \mathcal{C}, \text{the observed preferred condition, self-reported by participant}$$
$$r_{d,c} \in \mathbb{R}, \text{the scale response to a dimension } d \in \mathcal{D} \text{ in condition } c \in \mathcal{C}$$

The *observed preferred condition* o is the participant's answer to the question about their preference. The scale response r is the value computed from the questionnaire responses.

Now the question is how to understand whether the (observed) preferred condition o is dependent on the scale responses r? This needs some motivation. The analyses of the main effects of the scale dimensions have been conducted with a paired test: the Wilcoxon signed-rank test. This means, the change of participants' answers between conditions is analyzed, rather than comparing *all* answers of one condition to another one. The main reason for choosing this test was rooted in the research questions and the corresponding study design. The participants were not given any context but the task to explore whether the robots in the conditions were different. This was done to not bias participants' expectations. However, their responses could therefore have been very different, depending on their assumptions of the robot's capabilities. For example, a person who expected the robot to behave and speak like the Star Wars character might have been disappointed by the robot's actual behavior and had mainly responded on the lower end of the scale. In contrast, a person with fewer expectations might have been excited about the robot's behavior and always answered on the opposite side. Averaging all their answers to one condition and comparing them to another would be less conclusive than comparing whether participants usually rated the one condition higher than the other. This was exactly why the use of paired tests was chosen for analyzing the main effects. For the same argument, investigating whether o is dependent on the response value r might not yield much information, as it is unlikely that the participant's response values can be compared.

Therefore, to understand whether the observed preferred condition o depends on the participant's response r, an auxiliary variable is introduced: the *expected preferred condition* e. This variable categorizes the responses r of the participants and computes a preferred condition out of their responses. A simple way is to use the maximum response value to retrieve an expected preferred condition.

More formally, let $\hat{d} \in \mathcal{D}$ be the dependent scale dimension and let $\boldsymbol{R}_{\hat{d}}$ be a sequence of all three scale responses to dimension $\hat{d}$, constructed as:

$$\boldsymbol{R}_{\hat{d}} = \{r_{\hat{d},ADA},\ r_{\hat{d},ADA_{LE}},\ r_{\hat{d},REP_{LE}}\}.$$

Then, the expected preferred condition e computes with:

$$e_{\hat{d}} = c,\ \text{if } r_{c,\hat{d}} = \max_{c} \{\boldsymbol{R}_{\hat{d}}\} \text{ and } r_{c,\hat{d}} \neq r_{c,d} \forall d \in \mathcal{D} \setminus \hat{d} \tag{5.1}$$

Intuitively, the expected preferred condition is the condition that returns the highest participant's scale response value for a specific dimension. An example: let $\hat{d}$ = Warmth be the dimension that is used as a discriminator to predict the participant's preference. The responses of one participant to the three conditions are given as $\boldsymbol{R}_{\hat{d}} = \{3,\ 4.3,\ 2\}$. Since $\max\{\boldsymbol{R}_{\hat{d}}\} = 4.3$, the expected preferred condition e_{Warmth} is the condition ADA_{LE}. In other words, the participant is expected to prefer the intrinsically motivated robot with the proximity sensor.

Equation 5.1 allowed for computing the expected preferred condition based on all given dimensions. This way the expected preferred condition e could be analyzed for its dependency on the observed preferred condition o.

Qualitative analysis

Figure 5.7 gives a first impression of the data frequencies. The balloon plots are a visualization of a contingency table for each of the dimensions $\hat{d}$ of the questionnaires. The rows are the levels of the observed preferred condition o, the columns are the levels of the expected preferred condition e. A perfect association of e and o would result in high frequencies on the main diagonal and zeros elsewhere.

There were a few qualitative observations possible when studying Figure 5.7. It can be seen that for the dimensions Warmth and Competence the frequencies along the main diagonal are high. As presented earlier in Table 5.5, Warmth showed medium and large main effects for the ADA_{LE} condition. Figure 5.7 confirms that participants preferred the ADA_{LE} condition, and, even more important, this is associated with the expected preferred condition, which was computed from their responses to the dimension. On the other hand, there were no

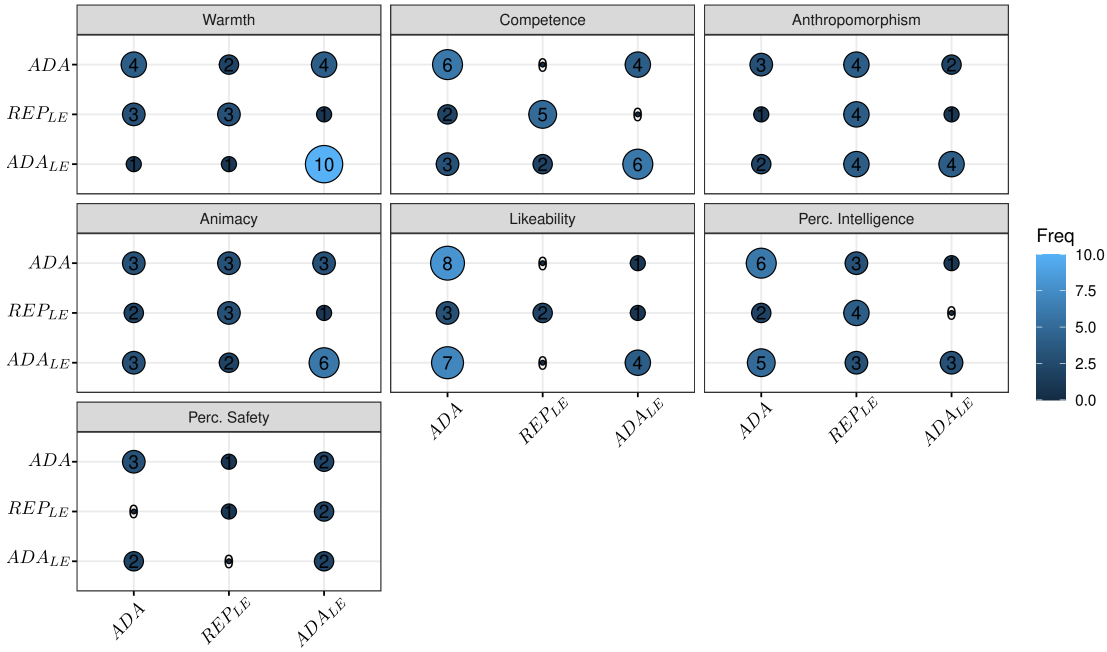

Figure 5.7.: The contingency tables of each dimension d for the observed preferred condition o (rows) and the expected preferred condition e_d (columns) as balloon plots. The larger the size and the lighter the blue, the higher the cell frequency.

main effects for the Competence dimension presented (cf. Table 5.5), which is reflected in the figure, as there was no unique high frequency as for the Warmth dimension. However, it can be seen that there is an almost equal distribution of frequencies along the main diagonal. This means that despite the absence of a main effect, participants' preferences seemed to be associated with their responses to the Competence dimension. This means that the highest responses to the dimensions Warmth and Competence from the RoSAS reflects the participants' preferences. The observations for the other dimensions were not as promising. In particular, Likeability did not show such an association.

Quantitative analysis

This subsection presents quantitative analyses. It first presents the results of an analysis of whether there was *any* dependency between the expected condition e and the observed condition o. It then presents the magnitude of effect, in this case, how consistent was the expected preference e at predicting the observed condition o. The final test was a correspondent analysis that showed how the levels of e and o were associated with each other. For example, were participants who preferred the intrinsically motivated robot $o = ADA$ likewise responding highest to Warmth in the same condition $e_{\text{Warmth}} = ADA$?

Dependency A possible quantitative tool to investigate whether both categorical variables o and e were dependent is the Chi-squared test, which was described at the beginning of the section. However, both variables e and o had three levels which resulted in a contingency table with 4 degrees of freedom. Together with the sample size of 36 participants, the resulting *expected cell values* were smaller than 5, which violates the requirements to use the Chi-squared test. An alternative in such cases is the *Fisher's exact test*[11].

Table 5.8.: Fisher's exact test for count data.

	$\hat{d}$ (predictor)	p	N
RoSAS	Warmth	0.033	29
	Competence	0.007	28
Godspeed	Anthropomorphism	0.803	25
	Animacy	0.593	26
	Likeability	0.118	26
	Perceived Intelligence	0.487	27
	Perceived Safety	0.635	13

Table 5.8 shows the resulting p values of the exact test and the number of considered participants N. Participants were not considered when they either did not have a preference or when e_d was not defined (i.e. there was more than one condition with a maximum scale response value). The Fisher's exact test tests the null hypothesis that both variables were independent. Table 5.8 shows that the null could be rejected for Warmth and Competence, i.e. the participants' responses to these two dimensions were dependent on the observed preferred condition reported directly by the participant, and vice versa.

These statistically significant results showed that participants' responses to Warmth and Competence were associated with their reported preferred conditions. What they did not say, however, was *how* meaningful the association was. Therefore, a measure of the strength of that association (i.e. effect size) was needed, which is presented in the next paragraph.

Effect Size The uncertainty coefficient $U(o \mid e)$ quantifies the magnitude of above effects[12]. It describes how consistent the expected preference e can predict the observed condition o. The uncertainty coefficient U measures the *strength* between categorical[13] association using the conditional entropy, i.e. the proportion of uncertainty (Nehmzow 2006).

U is commonly used to evaluate the effectiveness of cluster algorithms. An interesting property is that it does not take into account any correspondence assumptions, so it did not matter how the levels of e and o were hypothesized to be related. This is a joint property with the Fisher's exact test (or Chi-squared test), which makes U a good choice for an effect size. Note that U is independent of the number of levels of the variables (i.e., the size of the contingency table) or the sample size of the study. This allows comparing the *strength* of the

[11] The Fisher's exact test computed with the implementation `fisher.test` of base `R`'s built-in `stats` package.
[12] Sometimes, it is also ambiguously referred to as Theil's U. Although Theil (1970) derived a considerable part of the uncertainty coefficient, the term Theil's U usually refers to the U statistics used in finance.
[13] U can be extended to continuous variables.

association between this and future studies. U is a directed effect. The interesting question for this study was: how much does the highest scale response to a dimension (e_d) tell us about the observed preferences (o), i.e., the participants' self-reported preferences. More formally: what fraction of the remaining uncertainty of o can be predicted given e: $U(o|e)$?

Table 5.9.: Uncertainty coefficient U.

	$\hat{d}$ (predictor)	N	$U(o \mid e)$
RoSAS	Warmth	29	0.178
	Competence	28	0.135
Godspeed	Anthropomorphism	25	0.036
	Animacy	26	0.054
	Likeability	26	0.063
	Perceived Intelligence	27	0.049
	Perceived Safety	13	0.020

The results[14] in Table 5.9 show that U was by far the highest for the dimensions Warmth and Competence – indicating that by itself, those two dimensions were by far the best predictors of self-reported human preference for this study.

A value of 1 would indicate that a given dimension reduces all remaining uncertainty in the prediction. The value of U always lies between 0 and 1, which allowed comparing how much each dimension predicts the self-reported preference o. It can be seen that Warmth and Competence provided several times as much uncertainty reduction as the other dimensions. Note that this was just the reduction of uncertainty by knowing which conditions had the maximal response for one singular dimension. If Equation 5.1 would combine dimensions, or consider the scalar values, even better predictive power could potentially be achieved.

Correspondence analysis The results of the Fisher's exact test revealed that there was a statistically significant association between e and o for the dimensions Warmth and Competence. Then the uncertainty coefficient U provided a measure of the strength of that association. It was then important to understand the correspondence between the levels of e and o to answer RQ5: do participants prefer the condition with the robot they perceive highest in Warmth?

Observing the distribution in Figure 5.7 already indicated that the expected conditions e correspond mostly with the observed conditions o. This could be seen by the main diagonal which has higher frequencies than the other cells. For example, if the use of Competence results in the expected condition REP_{LE}, then it is very likely that Competence has a high response for that condition too. For Warmth on the other hand, if the expected condition is ADA_{LE} then it corresponds to the same condition as the observed condition. However, for

[14] U was computed with `UncertCoef(table(o,e), direction=c("column"), p.zero.correction=T)` from the `R` package `DescTools`.

the other conditions this is less obvious.

Correspondence analysis (CA) is a candidate to understand whether the levels (i.e. conditions) of two categorical variables, such as e and o, correspond to each other. It is a descriptive statistic and an extension of the Principal Component Analysis (PCA) to categorical variables.[15] In short, the analysis breaks down a higher dimensional table into fewer dimensions, resulting in a plot. The plot then allows analyzing which of the levels of e and o correspond to each other.

Figure 5.8 shows the plotted results of CA for the scale dimensions (a) Warmth and (b) Competence. The first important information is that two dimensions suffice to explain the correspondence between levels.

A strong correspondence between levels of o (blue) and e (red) is indicated by a small angle between the arrows pointing from the coordinate origin to these levels. An angle larger than 90° indicates no association. The distance from the coordinate origin indicates the strength.

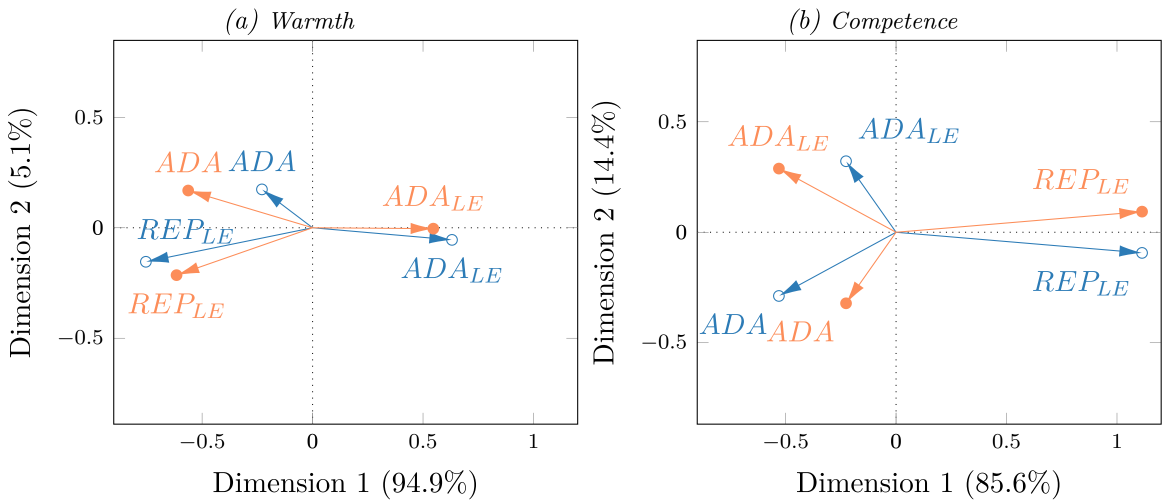

Figure 5.8.: The results of the correspondence analysis (CA) for the levels of o (blue) and the levels of e (red) for the dimensions (a) Warmth and (b) Competence. For example, for the dimension Warmth the condition ADA_{LE} corresponds between e and o.

Figure 5.8a shows that the levels which corresponded most to each other were the same conditions. For example, the level ADA of o corresponded most to level ADA of e. In other words, participants who preferred ADA responded highest to Warmth in ADA. This was true for all conditions, but the strongest correspondence was revealed when observing ADA_{LE}. The same was true for the scale dimension Competence (see Figure 5.8b). The levels of e and o corresponded strongest to each other when they represented the same condition. This correspondence was even stronger for the scale dimension Competence, with the condition REP_{LE} being most strongly correspondent.

These results underlined the qualitative observations discussed earlier in this section. The

[15] Computed with `ca(table(o,e), arrows=c(T,T))` with the `R` package `ca`.

correspondence was strongest for REP_{LE} in Competence and ADA_{LE} in Warmth. The latter has to be taken with care. It has been discussed earlier that Warmth was perceived highest in the ADA_{LE} condition, with statistical significance. This also explains why there are more data points for that condition, which made the correspondence for the other conditions weaker.

Summary The results showed that the expected preferred condition e and the observed preferred condition o were statistically significantly dependent for the dimensions Warmth and Competence, but not for any of the other dimensions. Furthermore, the results showed that the levels of each of the variables (the conditions) corresponded to each other.

What stood out was that these findings of statistically significant dependency and correspondence were prevalent despite the generally small main effects discussed earlier.

An interesting side observation was that there were no such results for any of the other dimensions, in particular not for the dimension Likeability. Given the concept that the dimension Likeability represents, it is very interesting that the dimension failed to show dependency for e and o. Contrary to Competence and Warmth, Likeability did show the highest main effect of all dimensions.

The interesting point about the results here is that the two central dimensions of social attitude formation – Warmth and Competence – were the best candidates to understand the preference of participants. This was because using them to compute the expected preferred condition, which was that there was statistical significance for the dependency test between the observed o and the expected preferences e, showed without a doubt the highest effect for predicting the participant's preference (U). Additionally, it was because the levels of o and e corresponded the most if they represented the same condition.

This directly answered RQ5: yes, the participants' reported preference corresponds to the condition with their highest response to the dimension Warmth.

5.4. Discussion

This study addressed the suggestions and ideas drawn from the second study: (i) to design a more similar baseline behavior to understand whether the effect for Warmth persists, despite similar perception of the robot's behavior and Animacy, (ii) to use a proximity sensor to investigate whether a robot, which is able to directly respond to the input of the participants, has a stronger effect on the dimension Warmth compared to a robot behavior based on IMs and, (iii) to analyze the relationship between the participant's reported preferred condition and the condition with their highest response to Warmth, in order to better understand whether the knowledge about Warmth transfers to HRI.

As a reminder, the second study found that an intrinsically motivated robot was perceived significantly more warm compared to a reactive baseline behavior. This was a very promising result but it needed further investigation to determine the cause of the effect. This was due to the unexpected result that the participants perceived the intrinsically motivated robot as more animated (with statistical significance), and participants also reported that the intrinsically motivated robot and the baseline were mostly very different. The latter was not of too much concern, since the overall task of the participant was to judge whether the two robots were different. This might have biased their judgment of differences or made them more alert to small differences. However, the unexpected, high difference in the perception of Animacy could not be explained. Other studies found that robots, or even objects, are perceived high in Animacy if they are moving and, at the same time, the influences to their locomotion changes are not visible to the observer (Tremoulet and Feldman 2000). This knowledge and the different perception of Animacy between the two behaviors was a strong indicator that the robot behaviors were possibly different.

This conclusion directed the focus to the used baseline behavior and possibilities to enhance it. As a reminder, in the second study the baseline behavior was based on pre-adapted weights of a previous trial, which remained constant during the interaction with the participants. Sensory input was, depending on the constant parameters, transformed into heading and speed variables, which were then applied to the built-in balanced controller of Sphero. The balanced controller kept the robot upright, while simultaneously trying to achieve the requested heading and speed. In contrast, the intrinsically motivated robot directly changed its servos' speed based on the results of TiPI maximization, which also resulted in immediate responses to perturbations.

This resulted in the concern that (i) the use of different motion controls and (ii) the capability of reacting directly towards perturbations was the major influence on the Warmth perception, instead of the IMs of the robot.

To address these two concerns this study was adapted twofold. Firstly, the baseline behavior used the same direct motion control and was fakely adaptive to address point (i). Secondly, a proximity sensor was introduced, so that one robot was able to respond directly toward human perturbations compared to another one, which aimed to address point (ii).

The baseline behavior of this study was changed to allow it to use direct motion control and to be fakely adaptive. Switching to the direct motion control was a matter of not using the balancing controller and directly applying the output of TiPI maximization to the servos (cf. section 2.6). This however made the use of fake adaptivity necessary. If the parameters of the robot were kept constant (like in the previous study), this would have caused a very repetitive behavior (mostly spinning), because there was no interference of a balancing controller. To enable fake adaptivity, the present baseline used parameter adaptations which were stored during a previous interaction with an intrinsically motivated robot. When the baseline behavior was presented to a participant in this study, the robot *replayed* these updates as they occurred previously. This made the behavior of the robot more variable, as the parameters changed, but the robot did not actually adapt to its experiences when interacting with the human. The results of this study showed that these changes successfully resulted in a baseline behavior that participants perceived similarly animated to the intrinsically motivated robot (RQ1). To contextualize the results the reader can get an impression by watching the supplementary video material (see Scheunemann 2017e). It can be seen that the behavior generation across conditions was similar in many features, which led to robot behaviors that were hard to tell apart – even for the experimenter.

The other major change compared to the previous study was the introduction of a proximity sensor based on BLE, which allowed the robot to sense the participant's proximity. The sensor enabled a robot (i) to react based on that very proximity and (ii) to distinguish between human interactions and perturbations induced by the environment. Both these capabilities again enabled the robot to react quicker and directly to human perturbations, which was similar to the differences in the robot motion controls and their reactivity of study II. This study had two conditions with intrinsically motivated robots, which were exactly the same, but in one condition the robot used the proximity sensor and in the other it was not using it. This allowed the analysis of how much the robot's capability to react differently to participants' and environmental perturbations influenced the perception of Warmth. Both these conditions were compared to the third condition: the previously described baseline robot, which also used proximity information.

The results of this study provided evidence that a robot equipped with a sensor that allows unambiguous behavior depending on the human proximity appeared slightly more animated, but less intelligent. Furthermore, if two intrinsically motivated robots were compared (one with and the other without such a sensor), humans seemed to favor the robot which was closely entrained with the human by sensing their proximity. This shows that the choice of the embodiment of a robot needs careful consideration, as even a single sensor (like the one here) can have an effect. However, although the sensor had some effect, the results also showed that the robot's IM is the strongest influence on the perception of Warmth (RQ4).

From what has been discussed so far – a high similarity of the robot behaviors – it may seem unlikely that the initial hypothesis that an intrinsically motivated robot is perceived as more warm than the fakely adaptive baseline behavior could hold true. And yet, the experience directed adaptation played an important role, since using only fake adaptation led to a statistically significant drop in the perception of Warmth. The results of this study revealed that despite the similarity between the behaviors, an intrinsically motivated robot that is closely entrained with a human elicits a feeling of Warmth in an embodied social cognition scenario (RQ3).

This also addressed the two major concerns outlined in the second study: could the largest impact on the positive effect for Warmth stem from the robot's motion control or its capability to react to human perturbations directly? This study alleviated these concerns and showed that the main impact on Warmth is dedicated to the robot's behavior resulting from its IMs.

Up to this point of the thesis, the studies concentrated on the effect on the dimension

Warmth, because in the disciplines around social cognition it is considered central for understanding social attitude formation (Fiske et al. 2007; Abele, Hauke, et al. 2016). Moreover, it has been found that humans experience more positive social interactions when they are perceived high in Warmth (cf. section 2.4). Therefore, it was previously argued in this thesis that if this knowledge transfers to interaction between humans and robots, a robot that is perceived as more warm could likewise sustain HRI.

Throughout this thesis the RoSAS was used: a questionnaire that implements that very concept of Warmth in the domain of HRI. Carpinella et al. (2017) evaluated the scale, but based solely on standstill images, or in other words, only based on visual one-shot impressions. This is certainly an important evaluation, especially given that our impression of Warmth unfolds within milliseconds (Fiske et al. 2007). However, it needs further evidence whether the knowledge about the dimension Warmth from social cognition (i.e., the science for understanding human-human interaction) transfers to physical *interaction* of HRI. Only then is Warmth a good candidate for measuring human perception of physical, interactive robots.

This study investigated the relation between participant's reported preferred condition and their responses to the Warmth dimension. The idea was that if Warmth is indeed a good candidate to predict positive interactions, participants would most likely prefer the condition where the robot scored highest in Warmth. The results of this study showed exactly that: participants prefer to interact with the robot they perceive most warm (RQ5).

This of course can only be considered a first step toward linking the concept of Warmth to HRI. And yet, the findings successfully underpinned the (so far) working hypothesis which in turn gives further weight and meaning to previous findings that an intrinsically motivated robot is perceived as more warm.

5.5. Limitations and future work

This section discusses the study's limitations and links them to potential future studies. Most studies of this thesis were concerned with incrementally developing a study design to best understand the effects of IM in robots on the dimension Warmth. Naturally, most limitations specific to this study design have been addressed before.

A new limitation, however, resulted from observing the analysis for interaction effects. There was only evidence for interaction effects for the dimension Discomfort in this study. However, the analysis showed an overall high likability for possible interaction effects compared to all previous studies. This indicates that the last study could reveal interaction effects if a larger sample size would be chosen. Post-hoc tests could further analyze the data of all dimensions with an eye on understanding how much the ordering of conditions played a role and provide indications for similar future studies with more participants. Other than that, the major part of this section is dedicated to the novel approach to link social cognition to HRI.

As it has been mentioned earlier, the provided evidence for the dependency on the reported participants' preference and their responses to the dimension Warmth can only be considered a first step to understand the ties of the concept of Warmth between social cognition and HRI.

This study design could benefit from a set of changes to provide a deeper understanding of this connection. First and foremost, the assessment of the participant's preferred condition could be achieved in a more fine-grained fashion. In this study, participants had to choose between distinctive conditions or choose that they do not have any preference. In future studies, their responses could instead be collected on a scale letting them decide how much they preferred each condition or, alternatively, participants could answer after the second (and then following) interaction how much they preferred the last interaction to previous ones. This would allow participants to prefer more than one condition and it would broaden the options for analyzing the results. For example, it could be investigated whether participants usually prefer the first interaction more than the second. It would further enable tests which analyze the dependency between the amount of reported preference and the participants' response value to Warmth.

Another possible change would be to design an interview to assess the participants' preference. However, letting participants directly report the variable of interest is prone to participant biases. One popular example was described by Orne (1962) which he coined *demand characteristics*. It describes that participants may try to be particularly *good* in order to satisfy the perceived needs of the researcher. This bias was addressed by keeping the current robot condition hidden from the experimenter. However, since the question asked for a preference and the participants perceived the robots differently, this bias could have caused the participants to choose one interaction as their preference, rather than reporting to have no preference. Therefore, it may be best to avoid or to accompany the direct reporting. For example, instead of evaluating their preference indirectly, one could measure the participants' engagement. The search for engagement measures is a challenging topic in itself and is an active research field in many disciplines. Going along this route was not applicable for this thesis, given the many novel areas this thesis already explored. However, a first step could be to analyze video data and use either the time of or the number of interactions by the participants as a criterion of their engagement. This, however, still needs detailed research, as it is unclear what kind of engagement can be expected when participants perceive a robot as more warm. It is not necessarily clear that a higher amount of interactions is something positive. For example, participants could increase the interactions because they desperately want the robot to change its behavior or, it's the opposite and they really want to interact with it because they perceive it as more warm.

The most promising way to explore the perception of Warmth and participant's preference indirectly could be to let participants quit the interaction by themselves. The length of time they actually interacted with a robot before they quit the experiment can be a measure of their preference. While this has been strongly considered, this path has not been conducted because of, again, the novelty of the research topic. With the knowledge after the third study, focusing on Warmth seems reasonable, but that was not clear by the start of designing the study. Therefore, keeping the condition times comparable to allow for all dimensions to be explored equally seemed crucial. Looking back, however, Fiske et al. (2007) argue that the perception of Warmth can unfold quite quickly. Future studies which focus on the Warmth dimension could therefore benefit from measuring the voluntary interaction time.

Another limitation is that the current evaluation is based on a study where nearly all main effects on the assessed dimensions showed only a small effect or no effect at all. Of course, this makes it even more interesting that Warmth and Competence showed a statistically significant dependency on the participant's reported preference, but it is likely that the indirect responses to the scale dimensions Warmth and Competence differed too little to see a larger, more convincing effect size on predicting people's preference. Future studies could tackle this in two ways: it could be beneficial to take the difference between the highest and the second highest response to a dimension into account when analyzing for dependency, giving more weight to the participant's responses where these values differ the most. On the other hand, it could be beneficial to repeat the collection of participant's preferences with an experiment that shows larger effects from the start. A good candidate would be to repeat study II together with the use of the preference measure used in the current study (or a suggested extension) to underline the current results or gain further insights.

An additional observation is that the results also showed that the participant's preference was not just dependent on their responses to Warmth, but also to Competence. What is very interesting about this is the different main effects of Warmth and Competence. While Warmth showed medium-sized, partially statistically significant main effects, the dimension Competence, in contrast, did not show any effect at all. This was of course intended by the design and therefore expected. However, the results suggested that the participant responses to the Competence dimension also determined their preference. This means that Competence may be a more sensitive discriminator for the participant's preference. Given that Competence ratings take time to unfold (ibid.), this needs special consideration for experiments that run for longer.

The results showed that both central dimensions Warmth and Competence influence human preference to interact with a robot. Future studies can further investigate the interplay between the two central dimensions. For example, it would be helpful to understand whether a specific combination of perceived Warmth and Competence forms the same attitudes in an interacting human participant as it would if the person interacts with another human. The more knowledge that is gained about the similarity in human attitude formation between human-human interaction and HRI, the more the knowledge from one discipline can support research in the other discipline. Untangling this social attitude formation and behaviors for the social interaction between humans and robots of course needs further work, but this study provides evidence that both central dimensions play a role, and it is likely that they do so in relation to each other. This is particularly important for studies with more task-focused, intrinsically motivated robots with behavior beyond the playful, exploratory behavior of this study.

The study showed that the analysis of the relationship between physical HRI and the dimensions Warmth and Competence can accompany future HRI studies, because assessing human participant's preference does not add too much extra time. This also holds true for the evaluation of the data and their suggested changes above. Future studies could collect this data for a variety of contexts and robot morphologies, which then can help to foster the understanding of the concept of Warmth for physical interactions between humans and robots.

There are, of course, an ample amount of further research opportunities, with some outlining the possibilities in the far distance. For example, future studies could investigate the effects of this study for longer interactions and for robots with different morphologies. Most importantly, the formalism as well as the study design has to be advanced so that the resulting behavior goes beyond the playful, exploratory behavior presented in this study.

5.6. Conclusion

This study was the final study of a series investigating the effect of a fully autonomous, intrinsically motivated robot on human perception. This study introduced a proximity sensor and a new motion control for the baseline behavior, to address the two major concerns extracted from the second study: could the large impact on the positive effect for Warmth stem from the robot's motion control, or could it stem from its capability to react to human perturbations directly? The results of this study alleviated these concerns and showed that the main impact on Warmth can be dedicated to the robot's behavior resulting from its IMs. This not only confirms the findings of the previous studies, but it also gives more weight to the effects by using intrinsically motivated autonomy in robots. This study showed that a fully autonomous, intrinsically motivated robot that is closely entrained with a human can elicit a feeling of Warmth in an embodied social cognition scenario. In particular, the experience directed adaptation by the robot's IM plays an important role, because it led to a statistically significant drop in the perception of Warmth when a robot was only not intrinsically motivated.

The dimension Warmth was the measure of choice because it is known that human's perceived as more warm experience more positive social interactions with their peers. This study addressed the question of whether that characteristic carries over to physical interactions between humans and robots. It provided evidence that humans who prefer a specific robot behavior were likely to perceive that robot behavior highest in Warmth. This is the first, but a very important step, to foster the understanding of how to link human-human interaction to physical HRI. It showed that Warmth is a potential candidate to evaluate the robot behaviors to indicate whether they sustain HRI. At the same time, this chapter showed that intrinsically motivated autonomy plays an important role in developing sustained HRI.

Chapter 6.

Conclusion

The thesis was concerned with two research questions: how to assess a human's preference of robotic behavior (RQ1), and whether human participants prefer to interact with an intrinsically motivated robot and, consequently, whether intrinsically motivated autonomy in robots can sustain human-robot interaction (HRI) (RQ2). The questions sparked two intermediate research objectives, which were needed in order to answer them: what is a good computational approach to enable intrinsic motivations (IMs) in robots (O1), and how to design a study to measure this (O2)?

This chapter first summarizes the overall thesis, focusing on the separate chapters and their relevant contributions to the research questions and objectives (6.1). It then takes a closer look at the answers to the two research questions (6.2). And finally, it lists the thesis' contributions to knowledge (6.3) and provides limitations and related ideas for future work (6.4).

6.1. Summary of research

Chapter 2. The focus of the background and development chapter was on *what* exactly constitutes as IM, and *how* it could be applied to the robot. The chosen robot platform to investigate this was a Sphero in its BB8 version and was described in 2.5. A limited, non-humanoid robot was chosen in order to limit the participants' expectations of the robot's capabilities. A computational model of IM needs to fulfill two main criteria in order to be robustly applicable to a robot platform in a real-world HRI scenario: (i) it needs to cover an infinite number of states, i.e., it needs to be able to work on a large range of continuous sensor input and (ii) it needs to be computable.

The chapter proposed to maximize the time-local predictive information (TiPI) as a computational model for IM. Predictive information (PI) quantifies how much of the future states are predictable based on past observations. To maximize this quantity the robot has to generate a rich sensory input, while at the same time keeping itself in a somewhat predictable state. The chapter derived the formulas to compute TiPI, along with a discussion of the needed approximations. Two approximations have a direct impact on designing this experiment: (i) prediction errors need to be very small and Gaussian and (ii) the noise must be independent of controller parameters. The first point (i) determined the choice of sensory input. Most notably, the computational model needed sensor input which underlies a Gaussian error, which essentially restricted the input to sensors that do not have sudden drops, caused by, e.g., incontinuities. A sensible choice, for example, is proprioceptive sensors, such as speed or acceleration sensors. The second point (ii) made it necessary that the noise of the sensor input is independent of the control parameters. Section 2.6 showed that the measured servo speed is not solely dependent on the control for that servo. Any error from using this proprioceptive sensor as an input was therefore dependent on the controller parameters, which violated the assumption. The chapter presented a simple motion model to break this dependency and showed that the use of this model enabled the generation of rich behavior.

The chapter further presented a sensor system which uses the received signal strength of Bluetooth Low Energy (BLE) to derive proximity information between the robot and the human, a sensor that was then used in the final study presented in chapter 5. The sensor system is a contribution of this thesis: it enables a robot to distinguish between humans in its vicinity and to recognize touch gestures. The main contribution of the chapter, however, is the overview of IM and that it addresses the first research objective O1: TiPI maximization is a possible computational model for IM, which can be applied to a real, minimal robot in an HRI scenario.

Chapter 3. This chapter presented the first of three interaction studies. The within-subjects study ($N = 16$) compared an intrinsically motivated robot to a reactive baseline behavior. The systematic development of a baseline behavior is a contribution of this chapter. Before the study started, an intrinsically motivated robot was placed in the same study environment and adapted its parameters. The adaptation was then stopped and the robot used constant, but adapted parameters to serve as a baseline for the study. Both robots were reactive to the same input sensors and both used the built-in balancing controller to locomote. The balancing controller requested heading and speed information and applied the values to the robot in such a way that the robot was kept upright.

The chapter also presented the first approach to design a suitable study. The main paradigm was to *enforce* HRI: the robot moved on a table with one edge open. The participants' task was to keep the robot from falling off the table. The idea was that the task ensured the participants would interact with the robot, so they could see its capability to adapt toward them. Additionally, the table that the robot locomotes on had different areas with varying altitudes and frictions. This was thought to further show the strength of the adaptive robot.

The results of the study showed that the idea of enforcing the interaction needed to be redesigned. Participants were very alert to the robot approaching the edge. However, they did not perceive the robot's behavior as approaching or interaction-seeking, but instead considered it as *faulty* or *suicidal*. Contrary to the hypothesis, the intrinsically motivated robot, which sought to maximize sensory input and was more adaptive to the environmental changes, was considered the least competent. This was due to the task given to the participants', as they implicitly assumed that the robot's goal was to stay on the table. However, the study also showed that the intrinsically motivated robot was perceived as more warm, with a medium effect. Warmth, together with Competence, is a central dimension in social cognition. Together they can be used to describe almost all social attitude formations (e.g. Fiske et al. 2007). Warmth is primary for *positive* social attitude formation, meaning that humans who are perceived as high in Warmth experience more positive social interactions from their peers (e.g. Abele and Wojciszke 2007).

Chapter 4. This chapter presented the second study with design changes following the discussion from the previous study. These changes were (i) concentrate on the perception of Warmth and (ii) enable behavior generation solely based on the robot's IMs.

Firstly, the environment was changed in order to fully focus on the perception of Warmth. The table was now circular, had no friction or altitude variations, and was fully enclosed. Without the open edge, a second change became necessary to motivate interaction. The participants were now presented with a *game-like* task, with instructions "to find out whether the two presented robots are different". The game was thought to prevent the participants from implicitly assuming a robot's goal, which does not match the robot's behavior. This game design also encouraged interaction by exploiting the participant's interest to perform *well* (e.g. Orne 1962). Additional encouragement was provided by handing the participants a wand-shaped tool. The idea was that the tool, of which its presented purpose was to interact, would cause the participants to feel the urge to use the tool and interact with the robot.

Secondly, the intrinsically motivated robot used the motion model (cf. 2.6) to directly change the speed of its two servos, instead of using the balancing controller as in the study of the previous chapter. This way the robot's behavior was only influenced by its IMs, unconstrained from additional software. This allowed to further focus the analysis on the perception of intrinsically motivated autonomy.

The results of the within-subjects study ($N = 24$) showed that both robots were perceived as similarly competent and intelligent. This indicated that changes to the study design were successful, and participants did not project an implicit goal onto the robot which mismatched its behavior. This had the additional benefit that the results on the Warmth dimension were not influenced by the perceived robot's Competence (cf. Fiske et al. 2007). Most importantly, results showed that the intrinsically motivated robot was perceived (statistically significantly) more warm than the baseline behavior. However, results also showed that both behaviors were perceived very differently and the intrinsically motivated robot was even perceived as more animated. It is known that humans perceive robots and even objects as animated if the cause of their movement changes are not obvious to the observer (e.g. Castro-González et al. 2016). This raised the concern that the motion regimes of the baseline robot were too different (maybe even too predictable) in comparison to those of the intrinsically motivated robot.

Chapter 5. This chapter presented the final study of the thesis, which focused on confirming the effect of Warmth while controlling the concerns found in the previous study: was the intrinsically motivated robot perceived statistically significantly more warm because (i) it was more animated and had a different behavioral regime than the baseline, or because (ii) participants could see that the robot was responding more directly to their perturbations, or because of (iii) the exploratory, playful behavior generated by the robot's IMs? This time, the baseline robot in this chapter used the same direct motion model as the intrinsically motivated robot, in order to appear similarly animated. Without the mediating balancing controller, a robot that could directly control its servos based on constant parameters would elicit a very monotone behavior. Therefore, this chapter suggested implementing *fake* adaptivity, which means that the robot received parameter updates, but not based on TiPI maximization. Instead, the updates were recorded during a previous run of an intrinsically motivated robot and then *replayed* for the baseline. This way, the updates were not random, but also not truly adaptive. In order to control the robot's capability to respond directly to participants, the study had two conditions with intrinsically motivated robots: one with a proximity sensor, and one without a proximity sensor. The sensor enabled the robot to perceive human proximity by using BLE signal strength between the wand-shaped tool and the robot (cf. 2.7).

The results of the within-subjects study ($N = 36$) showed that the new baseline behavior was perceived similarly animated and, overall, that participants perceived the baseline as much more similar to the intrinsically motivated robots. This means that the designed baseline behavior, based on a *parameter replaying controller*, was a good candidate for comparison. Similar to the previous study, it was also perceived similarly competent compared to the intrinsically motivated robots. Despite these similarities, the two intrinsically motivated robots were both perceived as more warm. In particular, the intrinsically motivated robot with the proximity sensor was perceived statistically significantly more warm than the fakely adaptive baseline behavior, which also used the proximity sensor. This underlined the evidence found in the previous study: an intrinsically motivated robot in an embodied social cognition scenario elicits a feeling of Warmth, a dimension known to be central to human attitude formation. The changes to the study design provided evidence that this effect was mainly routed in the robot's IMs, and not because of (i) the differently perceived baseline or (ii) the robot's capability to respond directly to human perturbations. The study also showed that the proximity amplified the feeling of Warmth. This means that an intrinsically motivated robot, which can adapt toward the proximity of the human it interacts with, elicits a stronger feeling of Warmth.

Another important contribution of this chapter is its investigation into whether the knowledge of social cognition transfers to HRI: do participants prefer to interact with the robot which they perceived highest in Warmth? This would indicate a parallel to social cognition, where it is known that human's who are perceived as more warm experience more positive social interaction. The study results showed that robots perceived highest in Warmth are the robots that human participants preferred to continue interacting with. In particular, using Warmth as a predictor for participant's intent for future interaction is better than using the dimensions Animacy, Anthropomorphism, Perceived Intelligence, Perceived Safety and Likeabilty from the Godspeed scale. This contribution is important for two reasons. Firstly, it gives more weight to the previous study results: it provides evidence that all intrinsically motivated robots were preferred over the baseline robot. Secondly, it shows that there is a link between human attitude formation toward peers and robots. If further evidence can confirm these results, it would provide future research with a good measure for human preference, by using tools established in social cognition.

6.2. Research questions revisited

With the above overview of the thesis, the two main research questions can now be revisited.

RQ1 **Can dimensions of social cognition be employed to measure human participants' preferences of robot behavior in order to understand what may sustain the interaction between humans and robots?** Prior to conducting any studies, several options for measuring were carefully assessed. Ideally, the sustainability of interaction could be measured by measuring the time a participant voluntarily spends with a robot. This path, however, has not been pursued because it was not clear whether an intrinsically motivated robot would be perceived positively in the first place. This thesis therefore employed secondary measurements to measure sustainability ethically, i.e. to not expose participants to a lengthy study without a solid background and educated hypothesis.

Social cognition motivated the focus on the two central dimensions: Warmth and Competence. Research has argued that these dimensions can express most attitude formations in humans toward other humans, with a high perception of Warmth being primary for a positive attitude formation.

This thesis used the Robotic Social Attribute Scale (RoSAS), which transfers these dimensions to HRI. However, the questionnaire has not been validated for an interaction scenario with a real robot, instead its validation relied on still images. This opened the question of whether the theory from social cognition truly carries over to *physical* HRI. In other words, if we perceive a robot as high in Warmth, does this likewise mean that we prefer to interact with it? This thesis extends the RoSAS evaluation by showing that there is a relationship between a robot behavior scoring high in Warmth and being preferred by human participants for future interaction. This is in line with conclusive research of social cognition: a human perceived high in Warmth experiences more positive social interaction. This is a first step toward understanding whether the knowledge from social cognition transfers to HRI: there is evidence that the central dimensions of Warmth indeed measures our affection toward robots.

This directly answered the research question: the dimension Warmth is a good candidate to measure participants' preference for continuing / repeating an interaction with a robot. Therefore, it is a good indicator of sustainable HRI.

RQ2 Can an autonomously, intrinsically motivated robot, sustain the interaction with humans? In all studies of this thesis the robots' intrinsically motivated autonomy was enabled by TiPI maximization. All the studies found, to a varied extent, that an intrinsically motivated robot is perceived as more warm when compared to an autonomous, reactive baseline robot. This effect persisted even for a baseline behavior that participants perceived to be very similar to the intrinsically motivated robot, like the one used in chapter 5.

Combined with the results of RQ1, namely that the perception of Warmth reflects human preference, it is strong evidence that human participants prefer a robot which is intrinsically motivated over an autonomously, reactive robot, which determines its actions in such a way that it generates similar behavior.

Further research is needed to understand more about the uncovered effects of IMs. First and foremost, sustainability needs to be measured in an everyday context, over a longer period of time, beyond inferring it from the perception of Warmth. However, the results underlined that IMs can help to increase human interest to interact with robots. The thesis adds the argument that this can be achieved by using a computational model of IM, which enables the robot to create intrinsically motivated behavior itself, without the need of an experimenter or scientist to pre-define curiosity.

The answer to this research question was that the TiPI formalism enables intrinsically motivated autonomy in robots, which elicits a feeling of Warmth in human interaction partners. Given that the knowledge of Warmth transfers from human-human interaction over to human-robot interaction (as the answer to RQ1 suggests), this is evidence that such behavior can sustain HRI.

6.3. Original contribution to knowledge

This thesis is, to the best of my knowledge, the first attempt to investigate the impact of intrinsically motivated autonomy on human perception. It used a computational model to implement intrinsically motivated autonomy into a robot. The model, which was based on information theoretical approaches, was then used to analyze the effect of an intrinsically motivated robot on human perception. Furthermore, the thesis links back to social cognition and investigated whether the perceptions of humans and the attitude formation toward humans transfers to human attitude formation toward robots. Therefore, this thesis contributes knowledge to the five disciplines: information theory (6.3.1), robotics (6.3.2), experimental methodology (6.3.3), and HRI (6.3.4).

6.3.1. Information theory

The main contribution to information theory is that this work shows an information-theoretically based computational model of IM which had an impact on human perception. The thesis used PI maximization to implement intrinsically motivated autonomy. It has been argued that PI is “the most natural complexity measure for time series” (e.g., Bialek et al. 2001). The behavior generated by PI maximization was then judged by observations, claiming that it results in “playful and exploratory” behavior (Martius, Der, et al. 2013b). This and all other computational models usually followed theoretically sound approaches, but how they perform in a real-world setting has not yet been analyzed. This thesis therefore adds a quantitative argument to the list. Without a doubt, humans are our best judges of how humans perceive robotic behavior. This thesis shows that human participants perceive the intrinsically motivated robot as more warm, which is the primary concept to understand *positive* human attribute perception. This means that PI maximization, as a candidate for a computational model of IM, creates a feeling of Warmth toward an artificial agent.

On the more practical side, issues with the original TiPI implementation by Martius, Der, et al. (ibid.) were discovered during the process of actually implementing the computational model. In particular, this concerned the parameter computation for more than two steps back. These observations were addressed to the authors and are publicly accessible (see Scheunemann 2018d). In addition, there are contributions to enable compiling the `lpzrobots simulator` on current architectures (see Scheunemann 2018c), a simulator created by the authors and used for their simulation experiments.

6.3.2. Robotics

The thesis further contributes to the field of robotics by providing guidance for *how* to implement PI maximization onto a robot (O1). The thesis emphasizes the approximations of TiPI maximization, and outlines the implications for applying the formalism and its approximations to a real robot. This contributes to further research in this area, especially for research that uses the same approach as this thesis. However, the guidance can also be used as an orientation for related computational models of IM. In general, similar approaches follow up on the same approximations to make the computation of the underlying information theory concepts possible.

On the more technical side of robotics, this work contributed code enhancements to run the Sphero robot. In fact, the off-the-shelf robot was chosen to expedite the start of research. It turned out, however, that the code to run the robot had a variety of issues. The provided official `JavaScript` framework, for example, did not parse the protocol correctly, which yielded a stuck robot behavior that forced a re-start of the controller. Sensor values, such as the roll angle of the robot, were faulty, along with the order of how the servo speed was set. These insights were forwarded to the company, mainly by using public pull requests to the (back then) official framework repository. The four contributions made which are most closely related to this thesis are the following. The first allowed the retrieval of quaternion readings (see Scheunemann 2017d) and another fixed the documentation (see Scheunemann 2017b). Arguably, the most important were the two more technical contributions: one which allowed making a connection to multiple robots simultaneously (see Scheunemann 2017c) and the other enabled parsing the robot's protocol correctly to prevent the robot from crashing (see Scheunemann 2017a).

When the above contributions were proposed to the official framework, it already occurred that the framework development would not continue for long. This is one reason why the knowledge was transferred into an own framework development based on `C++`. The other reason is that the `C++` language allows writing code that is closer to the hardware, which enables controlling the robot from embedded, computationally limited systems. The `C++` library is is publicly available (see Scheunemann 2017e; Scheunemann 2018b). The first attempts of this thesis involved working with autistic children in a nursery. In this context, the library was successfully used on a *Raspberry Pi 2* during the initial play sessions with children.

The public availability of the framework creates a simpler way of reproducing the experiments. More importantly, it allows other researchers to start straight away without the developmental issues present prior to this thesis. The robot platform eventually became discontinued altogether in 2018. This is a very common issue to robotic related contributions: the developed code is bound to a specific robot platform and the contribution is therefore only short-term. However, a technical contribution with a longer *date of expiry* is the derived motion model discussed in this thesis (cf. section 2.6). More specifically, the way the motion model was derived. The ideas can be extended to a variety of other robot platforms.

For the work on this thesis, a proximity sensor based on BLE was developed. It was successfully applied to the wand-shaped tool in chapter 5, which allowed the robot to sense the proximity of the human interaction partner. To the best of my knowledge, BLE has not been used in the context of robotics as a proximity sensor. Related research also shows that it can be used to prototype a touch sensor and that it can help to distinguish between people (e.g. Scheunemann, Dautenhahn, Salem, et al. 2016b). The technology is relatively cheap and easily applicable, and therefore can contribute to faster robot development and a faster design process for robot-related experiments. Instead of developing a whole vision pipeline in order to recognize humans, a researcher can start investigating by using the cheap BLE technology first. The technology can also be applied to service robots already present in human inhabitant environments. For example, one method to increase the functionality of a service robots is to distinguish humans and recognize reoccurring visitors. Instead of relying on images to accomplish this, which consumes modeling time and computational power, the robot can rely on the phone signals or a visitor's badge instead. The code for the sensor system is publicly available (see Scheunemann 2018e; Scheunemann 2018a).

6.3.3. Experimental methodology

This thesis presents novel developments for an experimental methodology in finding a suitable study design to investigate the impact of IM on the interaction between a human and a robot (O2). The traditional HRI angle is to externally impose specific behaviors or behavior patterns in a robot and to study how each of the isolated behaviors impacts the human perception of said robot. In the HRI literature, for example, this has been studied with robots that either asked questions to mimic curiosity, compared to robots that did not ask questions and were seemingly incurious (cf. section 2.1). In contrast, in this thesis, the robot is not provided with external behavior scripts that mimic IM, but instead the robot's behavior is generated by an IM algorithm that makes the whole robot behavior intrinsically motivated. This, however, creates challenges. Any external guidance of said autonomous behavior generation makes the behavior, by definition, not intrinsically motivated anymore. This makes it challenging to isolate specific factors of the generated behavior and understand their impact. This means, to study the effect of IM only, the behavior generation can only be compared between IM-induced behavior and not-IM-induced behavior. This, however, is another challenge as the formalism that creates the IMs does not create any behavior if switched off.

To measure the influence of IM, the methodology needed to ensure the above mentioned behaviors could be compared in a way that the observed effects can be addressed to IM. This resulted in two original contributions: (i) the design of similarly perceived baseline behaviors, and (ii) a game-like study design that encourages interaction with very limited input.

The design of a good baseline behavior is very critical and it needs to fulfill two tasks: it needs to be reproducible by other researchers and it needs to be systematically sound. For example, if I had compared an intrinsically motivated robot to a straight driving robot, the results would have had less quality. In order to be able to reduce any influence on the intrinsic motivation, the behaviors had to look very similar. The baseline proposed in the final study is very promising. It consists of a fakely adaptive robot, i.e., a behavior that is seemingly adaptive because it replays changes from an earlier run. This way it was almost impossible for the participants (and for the author) to tell the behaviors apart. The results show that although the intrinsically motivated robot behavior and the fakely adaptive behavior were similarly perceived, the behaviors appeared differently only because of the robot being intrinsically motivated. This thesis extensively discussed the thought process involved and it is believed that the same process can be applied on different and more complex robot platforms, like the one proposed in subsection 6.4.3.

The other key contribution is the proposed game-like study design in study II and in study IV that encourages participants to interact with the robot. In the first study, the interaction was rather enforced in the sense that participants had to interact with the robot to keep the robot on the table. This let the participants believe that the robot's goal is to stay on the table, making them very alert to exactly this and let them rate the robot's competence in achieving that goal. However, the perceived Competence can have an impact on our social perceptions of others, which makes the comparison between the behaviors less fair. In the last two studies, study II and study IV, the environment was redesigned so that the robots did not need interaction from the participants in order to stay on the table. Furthermore, the participant's task was to understand whether the presented robot behaviors were different. Both reduced the influence of the study on the participant's goal-assumption in the sense that none of the robots were considered more competent than the other. The game-like scenario exploits the participant bias, i.e., the participant's interest to perform particularly well in order to accelerate science (Orne 1962). To perform well, the participants assumed they had to interact in order to perceive any difference. This *proxy* task increases the participant's interest in interacting with the robot, while at the same time shields the true research interest.

Both contributions, the game-like study design and the design of similarly perceived baseline behaviors are very critical to the understanding of the perception of intrinsically motivated autonomy. Both shielded the task of the robot behavior and allowed interpreting the main effects as guided by the IM algorithm (O2).

6.3.4. Human-robot interaction

The contribution to the field of HRI is threefold: (i) the thesis links findings from the field of social cognition to the field of HRI, (ii) the thesis shows that the IM formalism of predictive information maximization can enable robust, and socially perceived behavior and indicates that this can sustain the interaction with a human participant, and (iii) the above experimental methodology provides a systematic framework to investigate behavior that is autonomously generated.

A central contribution of this thesis to the field of HRI is linking the knowledge from social human-human interaction (i.e. social cognition) to HRI. This thesis provides evidence that humans like to sustain the interaction with the agent we perceive highest in Warmth, independently of whether this agent is a human or a robot (RQ1). These findings are supported by recent research that explores the perception of robots by humans that play a game together. Paetzel et al. (2020) found that involvement is positively correlated to the perception of Warmth and Oliveira et al. (2019) found that participants preferred the robot that they perceived highest in Warmth. In this thesis, on the other hand, the findings were derived from studies where the robot's task is not explicitly clear to the human participant. Human-human interactions and relationships have been extensively studied over decades and it remains an active research area. If a connection between human-human interactions and HRI can be further fostered, this would allow HRI research to use a large set of existing methods and tools to evaluate robot behaviors.

Another central contribution is the study of fully autonomously generated intrinsically motivated robot behavior. The difference to existing work on IM is that the robot did not mimic IMs (e.g. Ceha et al. 2019; Law et al. 2017; Gordon et al. 2015) but was in fact truly intrinsically motivated, driven by its *interest* to explore the world through a hysteresis of predictability and change. The thesis shows that the IM formalism of predictive information maximization is perceived as more social compared to an autonomous robot that is seemingly adaptive and overall seemingly similar. The before mentioned link between human-robot interaction and human-human interaction gives more weight to this contribution: intrinsically motivated autonomy in robots is perceived as warm by human interaction partners. This provides evidence that IM could be key to sustain the interaction in HRI (RQ2). Moreover, realizing intrinsically motivated behaviors in robots could be key to have *robust* behavior generation in robots, For many robot systems that are tested or deployed in the real world, human interference can cause a potential risk to the reliability of the system, cause system failures or unwanted results. This thesis, on the other hand, uses a behavior generation that is robust to interactions because the intrinsically motivated autonomy employed is, in contrast, searching for new (i.e. unpredictable) interactions that enriches the robot's perception of the world. For research that aims to bring robots into the human-inhabitant real world, a robot that can handle unforeseen interactions (i.e. is robust to human perturbations) may become essential. The study of such robust, intrinsically motivated behaviors in HRI is another contribution of this thesis.

The robust behavior generation discussed above introduced the challenge of how to investigate human perception. This is because any attempt to constrain that behavior can lead to a system that cannot deal with unforeseen situations or to a behavior that is then, by definition, not intrinsically motivated any longer. The thesis contributes a systematic framework to approach this. The studies presented here let human participants compare two autonomous robot behaviors. This is achieved by encouraging the participants to interact with a robot using only a limited set of instructions: the only task given to the participants is to find out whether the robots' behaviors are different. This game-like approach helps to motivate the participants to interact, without giving away any robot task that can influence the perception of all participants (see subsection 6.3.3). This experimental methodology is an important contribution because it enables experiments that study the perception of fully autonomously generated behavior. In particular, the methodology allows this without the robot having a specific goal or task. This is important because, on the one hand, it is an active research area to look at how to combine behavior generation rules for IM with goal-oriented behavior. On the other hand, any combination will raise the question of whether the observed effects of the IM-induced characteristics and the goal-oriented characteristics can be truly isolated. The presented methodology accompanies the popular approaches of most HRI projects that concentrate on the robots' competence to fulfill a task appropriately. There, human interference could often cause a potential risk to the reliability of a system. Certain interactions would need to be explained to the participants prior to the experiment, in particular the task and the goal of the robot. A participant that deviates too much from the explained interactions could cause system failures or unwanted results, as there would be no adequate responses represented in the scripted behavior. The experimental methodology contributed by this thesis, on the other hand, encourages the participants to explore their interaction fully, while measuring their perception of said robots without biasing their implicit goal assignment of said robots.

So far, the above contributions resulted in three publications, and others are currently being written: (Scheunemann, Salge, and Dautenhahn 2019; Scheunemann, Salge, Polani, et al. 2021; Scheunemann, Cuijpers, et al. 2020). It is hoped that these findings encourage HRI research to explore intrinsically motivated robots further to allow a deeper understanding of our social perception of said robots. The thesis indicates that this could be key to have both sustainable HRI and robots that can autonomously adapt to unforeseen situations such as those in the unpredictable, human-inhabitant environment.

6.4. Limitations and future work

Each interaction study chapter carefully discussed limitations in relation to the particular study and suggested changes to a follow-up study. This section proposes three possible future studies, which can address three more general limitations of the thesis: (i) sustained HRI was indirectly measured based on questionnaires, (ii) the studies were conducted in a laboratory setting, and (iii) a simple robot platform was used.

6.4.1. Measure sustained human-robot interaction directly

This thesis explored different scale dimensions to better understand the human perception of robots. It further provided evidence that the dimension Warmth is a promising candidate to understand whether humans would like to continue interacting with a robot. Current research has also employed the measure of Warmth for repeated interaction with a robot. Paetzel et al. (2020) found that there is a positive correlation between participants' involvement in a game and how warm they perceive a robot. This suggests that the studies of this thesis could be extended for repeated interactions, with the hypothesis that a stable perception of Warmth indicates participants' involvement. However, the results that human participants preferred to interact with an intrinsically motivated robot are promising, and makes the effort of investigating directly whether IM in robots can sustain HRI worthwhile.

An idea that was addressed already in the final study (cf. section 5.5) is to measure directly whether participants interact for longer with the intrinsically motivated robot compared to a baseline robot. The interaction time is a common measure for HRI research that studies sustained interactions. For example, Graaf et al. (2016) measure the self-reported frequency of usage of in-home robots or Iio et al. (2019) measure the time people paid attention to a robot in a public space.

The within-subjects design employed in this thesis is not a good candidate to compare the voluntary interaction time of participants. This is because it is very likely that participants would interact longer with any first robot that they get exposed to, because a new robot may be more interesting in general, irrespective of its behavior or the participants' true preference.

Analyzing the interaction time directly would require a new study design. Two study designs seem most promising: (i) exposing participants to multiple robots at the same time and (ii) a between-subjects study that only presents one robot to one participant.

The two scenarios differ on how to assess the interaction time. The idea of point (i) is that participants could interact with two or more robots simultaneously. For example, they could interact with the intrinsically motivated robot and the baseline robot from study IV. The study could then measure which of the robots the participants preferred to interact with the most. The interaction time could be measured by counting the time between consecutive interactions. An alternative would be to quantify the number of observed interactions that each robot perceives. However, it has been argued in section 5.4 that the amount of interactions does not necessarily imply long-term interest. Generally speaking, a participant might be simply annoyed by one robot and tries to harm it. Instead, the experiment could consist of two different sets. After the participants have interacted with both robots, they could decide to continue interacting with exactly one robot only. This would measure the participants' preference directly. However, this would require the robots to be visually distinguishable, so that participants have a side by side comparison. This in turn increases the risk that the visual differences have an effect on the participant's decision.

The second point (ii) is a more straight-forward approach to measure the time of interaction. Each participant would be randomly assigned to a robot and they would be told that they can stop whenever they like. Then, the interaction times of the experimental group and the control group would be analyzed for differences. While this study design has been strongly considered prior to conducting the studies of this thesis, it needs to be noted that this kind of data assessment would need a considerably larger amount of participants. First of all, this design would require special treatment for participants that are familiar with the movie character that the current robot resembles. In the studies of this thesis, the focus was put on how differently participants perceive the robot in two different conditions. This is similar to the suggestion in (i). However, if a between-subjects design was employed, this could skew results depending on the participants' familiarity. Secondly, between-subjects studies are less powerful than within-subjects studies. This study would likely require access to hundreds of participants in order to see an effect. This is also why an alternative path with questionnaires was taken in this thesis.

Since this thesis provides evidence that an intrinsically motivated robot is preferred by participants, conducting one of the above studies is worthwhile. If one of the two suggested designs can confirm that participants prefer to spend more time with the intrinsically motivated robot, this would further support the results of this thesis.

As outlined above, the need for a large number of participants makes the realization of such a physical HRI study very challenging. One option, however, might be to use robots which are already present in our everyday life. This is something the next future study idea addresses.

6.4.2. Investigate the effect with a robot in the wild

One option to further investigate the effect of intrinsically motivated autonomy is to apply the algorithm to robots that are already present in our everyday lives. An example of a robot present in households is a vacuum cleaning robot. There are already thousands of vacuum cleaning robots placed in households, which makes it potentially possible to roll out a large scale study. An added benefit comes from the fact that the robots are already deployed and the robot behavior could therefore be changed unobtrusively. This kind of research would allow me to step out of the laboratory setting and investigate the effect in the real-world.

The question is: would humans value intrinsically motivated autonomy in their vacuum robot? Or do they rather prefer an *efficient* robot? My hypothesis is that even with robots that have a clear purpose with a quantifiable goal, it would still be beneficial if the robot would be perceived as more warm. This way, we may be more forgiving when the robot does not succeed, which is something a robot will always face in a real-world scenario.

A problem, however, and a limitation of the overall thesis is the applicability of the computational model of IM to a real-world robot. Should the robot be intrinsically motivated *at times*, for example when it reaches an unknown world state? Should the robot's drive to clean co-exist next to a computational model of IM, and how can that be even realized? Or, even more advanced, should the input of sensors be chosen in such a way, that the cleaning process itself is part of exciting its sensors?

A problem is the assessment of interaction time. Measuring the participants' interaction time with the robot may be challenging, since it may not be ethically possible to access data from the participants' homes. For example, facilitating cameras to measure the number of interactions from a participant would require the collection of video data to provide a means for recording the number of interactions. However, this presents obvious privacy issues, which renders the study either not applicable or would reduce the number of participants.

Instead, the knowledge contributed by this thesis could be used and a questionnaire could be designed. The RoSAS could be employed to investigate the development of the perception of Warmth and Competence, while at the same time controlling for participants' perception of the *usability*[1] of the robot.

If the above can confirm a positive development for the perception of the intrinsically motivated robot during a long-term HRI study, a possible next step could involve implementing intrinsically motivated autonomy on robots used for therapy. Wada and Shibata (2007) developed a seal-shaped robot named *Paro*. The robot is designed in such a way that it mimics

[1] A popular questionnaire is the System Usability Scale (Brooke 1986).

the behavior of a real animal. The robot Paro can react to touch, sounds or light, and has an implementation that enables it to adapt its behavior based on user input. For example, it changes states based on whether the robot receives a positive input, such as stroking behavior, or negative input, such as beating. The robot can also proactively generate behavior based on a variety of internal states and parameters.

Studies conducted with the robot in care homes suggest that the robot can help to increase the social network density, as well as reduce stress in participants with dementia or Alzheimer (e.g. Wada and Shibata 2007; Sabanovic et al. 2013; Lane et al. 2016). If future studies show that intrinsically motivated autonomy can yield sustained interaction, it might be worthwhile to implement a computational model of IM on the Paro robot too. That way, the increased interest could in turn increase the usage of Paro and, to that end, also increase its positive effects on patients.

Some of the challenges remain similar to the ones outlined for the vacuum robot. How is it possible to implement intrinsically motivated autonomy in a robot in parallel to, e.g., an implementation of drives. The problem here is that the work involves vulnerable participants. By no means should the robot behavior hinder the stress reduction in patients.

A benefit is the purpose of the robot which is to engage humans in interactions. This allows combining the study with the previous idea to measure the participant's interaction time with the robot easily. Since the patient changes the robot's sensor readings during the interaction, the robot could measure the interaction time autonomously, possibly even anonymized.

This future research avenue needs careful consideration of the ethical implications. Humans are inherently social beings. They assign social attributes to other beings but also to non-living things such as toys or robots. The aforementioned robot Paro, for example, is specifically designed to increase the participant's affection for it. The thesis shows that human participants have an increased social perception of intrinsically motivated robots. In other words, the robot's behavior can further enhance the positive emotions toward the robot. As with all research, there needs to be a balance between possible current harm (i.e., conducting a study that increases emotions toward a robot) with future harm (i.e., the lack of social robots in everyday life). The current COVID-19 pandemic has shown that there is a need for social robots. For example, it is argued that social robots could positively impact mental health issues apparent during times with social distancing measures in place (Ghafurian et al. 2020; Scassellati and Vázquez 2020). A social robot that is intrinsically motivated to interact with humans, without the need for a skilled human to control the robot, could help to elevate the harm arising by social distancing.

6.4.3. Increase complexity of the robot

The presented study was conducted on a non-humanoid robot platform with two degrees of freedom. One motivating idea behind this choice was that more complex robots, such as anthropomorphic robots, “might raise false expectations regarding the cognitive and social abilities that the robot cannot fulfill” (Dautenhahn 2004). The effects of these false expectations can be twofold. Participants could lose interest in the robot because its behavior does not meet their expectations. On the other hand, human participants could get overexcited about a robot that accidentally conducts *gesture-like* motions. For example, they could anticipate a waving gesture, which human participants could perceive as a *will* to communicate. This could highly impact the participants preferred interaction, even if the effect was caused by a random or accidental movement.

The simple robot platform of this study allowed to concentrate on the effects induced by intrinsically motivated autonomy. However, the platform is limited when it comes to studying more complex interactions. The question therefore is whether the computational model of IM used in this study can also be applied to more complex robot platforms. A range of new challenges needs to be carefully considered prior to this investigation. Firstly, how much more complex can the robot be? One challenge will be that, if participants imply robot capabilities, they will be disappointed if the robot cannot enact on them. This means, if participants anticipate that the robot has the capability to make sense of visual input, maybe because they see it has a camera built-in, then the robot should certainly use the visual sensor.

This however leads to other questions. Will the computational model still render computable? How to let the robot explore the input of visual sensors which are prone to occlusions? In theory, TiPI maximization can be implemented onto more complex robots, as it has been shown in simulation. Initial trials on real robots, however, have shown that the more sensors that are connected, the computation becomes increasingly complex.

The question of how to include sensors is a very subtle one and possibly the first step I will take toward understanding IM in more complex robots. In particular, I would like to understand how sensors with different timely implications can be combined.

This issue already appeared in the final study. There, the TiPI maximization used two different accelerations as sensor input: firstly, the self-acceleration of the robot, and secondly the acceleration of the human. These two sensors were the accelerometer of the robot and the proximity of the human. The relatively short experiment time and the human participants’ urge to interact with the robot have made both the accelerations comparable. The question is: what would happen if the human interacts relatively little with the robot? For example during a long-term study with a household robot, the human participant will not constantly interact with the robot. Because TiPI formalism only maximizes PI in a time-local manner, a change of the human proximity has a similar effect on PI compared to the robot’s self-acceleration. This is counter-intuitive, given that the sheer presence of rarely perceived signals should *make* the robot curious, which is the state IMs process too. The first idea that comes to mind is probably best described as a layered approach, where connected controllers maximize TiPI for different time-windows.

6.5. Concluding remarks

It has been shown that robots that appear curious can elicit curiosity in their human interaction partners. Previous research implemented this curiosity with hand-designed scripts or by using a wizard-of-oz design. These approaches require knowledge about the specific setting or they need a qualified scientist operating the robot. In contrast, this thesis used a computational model of IM, which enabled a robot to generate its behavior on its own. This way, the fully autonomous robot was created so that it could adapt to its environment and, in particular, adapted to an interacting participant.

The interaction studies of this thesis used a minimal, particularly non-humanoid, spherical robot platform. This decreased the expectations that participants tend to have of a robot's capabilities, and it allowed for a study of IMs in robots. The studies investigated how humans perceive such an intrinsically motivated robot compared to a reactive, baseline behavior. The results showed that human participants perceive the intrinsically motivated robots as more warm than the baseline robot. This also holds true for a baseline behavior that is perceived similarly competent and animated. Warmth is the primary dimension from social cognition. A human perceived high in Warmth experiences more positive social interaction. This thesis presents the first evidence that this knowledge about Warmth from social cognition transfers to HRI. In particular, this thesis shows that robots perceived highest in Warmth are also the ones that are preferred as interaction partners by the human participants.

The evidence of this thesis suggests that computational models of IM are possible candidates to make robots interesting enough to sustain interactions. This thesis further opens up a variety of possible research directions: from making these models more advanced, to testing them on more advanced robots.

Appendix A.

Dissemination

This chapter presents the publications and dissemination grown out of the work on this thesis. The work has been promoted and disseminated mostly in conferences and scientific talks.

A.1. Publications

The thesis' author was first author of all the following publications and their respective first draft.

- Marcus M. Scheunemann, Raymond H. Cuijpers, et al. (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)*. IEEE. Virtual Conference: IEEE, pp. 1340–1347. DOI: 10.1109/RO-MAN47096.2020.9223478 (62% acceptance rate)

- Marcus M. Scheunemann, Christoph Salge, and Kerstin Dautenhahn (2019). "Intrinsically Motivated Autonomy in Human-Robot Interaction: Human Perception of Predictive Information in Robots". In: *Towards Autonomous Robotic Systems*. Ed. by Kaspar Althoefer et al. Cham: Springer International Publishing, pp. 325–337. ISBN: 978-3-030-23807-0. DOI: 10.1007/978-3-030-23807-0_27 (73% acceptance rate)

- Marcus M. Scheunemann and Kerstin Dautenhahn (Feb. 2017). "Bluetooth Low Energy for Autonomous Human-Robot Interaction". In: *Proceedings of the Companion of the 2017 ACM/IEEE International Conference on Human-Robot Interaction*. HRI '17. Vienna, Austria: ACM, pp. 52–52. ISBN: 978-1-4503-4885-0. DOI: 10.1145/3029798.3036663

- Marcus M. Scheunemann, Kerstin Dautenhahn, Maha Salem, et al. (Aug. 2016b). "Utilizing Bluetooth Low Energy to recognize proximity, touch and humans". In: *2016 25th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*. New York, NY, USA: IEEE, pp. 362–367. DOI: 10.1109/ROMAN.2016.7745156 (47% acceptance rate)

- Marcus M. Scheunemann, Kerstin Dautenhahn, Maha Salem, et al. (Aug. 2016a). "Utilizing Bluetooth Low Energy for human-robot interaction". In: *2016 25th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*. New York, NY, USA: IEEE

The study presented in the second study has been written as a journal article and is currently under review:

- Marcus M. Scheunemann, Christoph Salge, Daniel Polani, et al. (2021). *Human Perception of Intrinsically Motivated Autonomy in Human-Robot Interaction*. arXiv: `2002.05936 [cs.RO]`

Furthermore, an abstract and a poster of preliminary results of the first study was contributed to the conference "Social cognition in humans and robots", which was hosted by the EU project socSMCs and the EUCognition network in *University Medical Center Hamburg-Eppendorf* from 27–28.09.2018.

A.2. Related publications

Two other publications were inspired by this thesis. The author used a simple threshold mechanism to distinguish between touch and close proximity using Bluetooth Low Energy sensors. For more advanced categorization, the idea was to semantically segment between this actions using deep neural network. This idea spawned the publication:

- Sander G. van Dijk and Marcus M. Scheunemann (Aug. 2019). "Deep Learning for Semantic Segmentation on Minimal Hardware". In: *RoboCup 2018: Robot World Cup XXII*. ed. by Dirk Holz et al. Vol. 11374. Lecture Notes in Computer Science. Springer International Publishing, pp. 349–361. ISBN: 978-3-030-27544-0. DOI: `10.1007/978-3-030-27544-0_29`

On a different note, the author was involved in many custom framework developments, mostly because popular framworks, such as the robot operating framework (ROS) were not built for real-time applications. However, the second version of ROS actually supported real-time needs. The author contributed to ROS 2 core development and some custom modules to push the ROS 2 development further, and he applied it to a humanoid robot in the RoboCup context:

- Marcus M. Scheunemann and Sander G. van Dijk (2019). "ROS 2 for RoboCup". In: *RoboCup 2019: Robot World Cup XXIII*. ed. by Stephan Chalup et al. Vol. 11531. Lecture Notes in Artificial Intelligence. Springer International Publishing, pp. 429–438. ISBN: 978-3-030-35699-6. DOI: `10.1007/978-3-030-35699-6_34`

For both publications, the author worked together with SD. SD did the main implementation part of the deep learning paper, but both authors contributed equally in writing. For the ROS 2 publication, both authors contributed equally to the lengthy integration and benchmarking progress.

Appendix B.

Questionnaires

B.1. Study I

ID: ___________________

Pre-test questionnaire

Before we can start, we need a few more information about your person.

1. Gender: ________________
2. Age: ___________________
3. Occupation:____

How familiar are you with interacting with robots:

	1	2	3	4	5	
Not familiar	○	○	○	○	○	Very familiar

How familiar are you with programming robots:

	1	2	3	4	5	
Not familiar	○	○	○	○	○	Very familiar

How familiar are you with the BB8 robot from Sphero (the robot used in this experiment):

	1	2	3	4	5	
Not familiar	○	○	○	○	○	Very familiar

1

ID: ___________________

Post-test questionnaire

In the previous session you have interacted with the spherical robot. We are now interested to find out how you experienced and perceived the interaction and the robot itself.

Please tick only one box (✓) per row. Do not skip any questions and answer them in the given order. There is **no right** or **wrong**.

1. Using the scale provided, how closely are the following attributes associated with the robot?

	1 (not at all)	2	3	4 (a moderate amount)	5	6	7 (very much so)
Dangerous	○	○	○	○	○	○	○
Emotional	○	○	○	○	○	○	○
Compassionate	○	○	○	○	○	○	○
Aggressive	○	○	○	○	○	○	○
Capable	○	○	○	○	○	○	○
Happy	○	○	○	○	○	○	○
Strange	○	○	○	○	○	○	○
Organic	○	○	○	○	○	○	○
Scary	○	○	○	○	○	○	○
Responsive	○	○	○	○	○	○	○
Awkward	○	○	○	○	○	○	○
Knowledgeable	○	○	○	○	○	○	○
Sociable	○	○	○	○	○	○	○
Feeling	○	○	○	○	○	○	○
Competent	○	○	○	○	○	○	○
Reliable	○	○	○	○	○	○	○
Awful	○	○	○	○	○	○	○
Interactive	○	○	○	○	○	○	○
	1 (not at all)	2	3	4 (a moderate amount)	5	6	7 (very much so)

1

2. Please rate your impression of the robot on these scales:

	1	2	3	4	5	
Machinelike	○	○	○	○	○	Humanlike
Ignorant	○	○	○	○	○	Knowledgeable
Awful	○	○	○	○	○	Nice
Unkind	○	○	○	○	○	Kind
Unintelligent	○	○	○	○	○	Intelligent
Dislike	○	○	○	○	○	Like
Dead	○	○	○	○	○	Alive
Foolish	○	○	○	○	○	Sensible
Unconscious	○	○	○	○	○	Conscious
Artificial	○	○	○	○	○	Lifelike
Fake	○	○	○	○	○	Natural
Moving rigidly	○	○	○	○	○	Moving elegantly
Incompetent	○	○	○	○	○	Competent
Mechanical	○	○	○	○	○	Organic
Stagnant	○	○	○	○	○	Lively
Unpleasant	○	○	○	○	○	Pleasant
Apathetic	○	○	○	○	○	Responsive
Unfriendly	○	○	○	○	○	Friendly
Irresponsible	○	○	○	○	○	Responsible
Inert	○	○	○	○	○	Interactive

3. Please rate your emotional state on these scales:

	1	2	3	4	5	
Agitated	○	○	○	○	○	Calm
Anxious	○	○	○	○	○	Relaxed
Quiescent	○	○	○	○	○	Surprised

2

4. Can you describe the different behaviours of the robot? Did the robot have any particular strategy for exploring?

5. What were the best and/or worst aspects of the robot's behaviour?

3

B.2. Study II

ID: ___________________

Pre-test questionnaire

Before we can start, we need a few more information about your person.

1. Gender: ________________
2. Age: ___________________
3. Occupation:____

How familiar are you with interacting with robots:

	1	2	3	4	5	
Not familiar	○	○	○	○	○	Very familiar

How familiar are you with programming robots:

	1	2	3	4	5	
Not familiar	○	○	○	○	○	Very familiar

How familiar are you with the BB8 robot from Sphero (the robot used in this experiment):

	1	2	3	4	5	
Not familiar	○	○	○	○	○	Very familiar

How familiar are you with "Star Wars":

	1	2	3	4	5	
Not familiar	○	○	○	○	○	Very familiar

1

ID: ____________________

Post-test questionnaire

In the previous session you have interacted with the spherical robot. We are now interested to find out how you experienced and perceived the interaction and the robot itself.

Please tick only one box (✓) per row. Do not skip any questions and answer them in the given order. There is **no right** or **wrong**.

1. Using the scale provided, how closely are the following attributes associated with the robot?

	1 (not at all)	2	3	4 (a moderate amount)	5	6	7 (very much so)
Dangerous	○	○	○	○	○	○	○
Emotional	○	○	○	○	○	○	○
Compassionate	○	○	○	○	○	○	○
Aggressive	○	○	○	○	○	○	○
Capable	○	○	○	○	○	○	○
Happy	○	○	○	○	○	○	○
Strange	○	○	○	○	○	○	○
Organic	○	○	○	○	○	○	○
Scary	○	○	○	○	○	○	○
Responsive	○	○	○	○	○	○	○
Awkward	○	○	○	○	○	○	○
Knowledgeable	○	○	○	○	○	○	○
Sociable	○	○	○	○	○	○	○
Feeling	○	○	○	○	○	○	○
Competent	○	○	○	○	○	○	○
Reliable	○	○	○	○	○	○	○
Awful	○	○	○	○	○	○	○
Interactive	○	○	○	○	○	○	○
	1 (not at all)	2	3	4 (a moderate amount)	5	6	7 (very much so)

1

2. Please rate your impression of the robot on these scales:

	1	2	3	4	5	
Machinelike	○	○	○	○	○	Humanlike
Ignorant	○	○	○	○	○	Knowledgeable
Awful	○	○	○	○	○	Nice
Unkind	○	○	○	○	○	Kind
Unintelligent	○	○	○	○	○	Intelligent
Dislike	○	○	○	○	○	Like
Dead	○	○	○	○	○	Alive
Foolish	○	○	○	○	○	Sensible
Unconscious	○	○	○	○	○	Conscious
Artificial	○	○	○	○	○	Lifelike
Fake	○	○	○	○	○	Natural
Moving rigidly	○	○	○	○	○	Moving elegantly
Incompetent	○	○	○	○	○	Competent
Mechanical	○	○	○	○	○	Organic
Stagnant	○	○	○	○	○	Lively
Unpleasant	○	○	○	○	○	Pleasant
Apathetic	○	○	○	○	○	Responsive
Unfriendly	○	○	○	○	○	Friendly
Irresponsible	○	○	○	○	○	Responsible
Inert	○	○	○	○	○	Interactive

3. Please rate your emotional state on these scales:

	1	2	3	4	5	
Agitated	○	○	○	○	○	Calm
Anxious	○	○	○	○	○	Relaxed
Quiescent	○	○	○	○	○	Surprised

2

4. Which words describe the behaviour of the robot best? (up to 5 words only)

3

ID: ___________________

Post-test questionnaire 2

In the previous session you have interacted with the spherical robot. We are now interested to find out how you experienced and perceived the interaction and the robot itself.

Please tick only one box (✓) per row. Do not skip any questions and answer them in the given order. There is **no right** or **wrong**.

1. Using the scale provided, how closely are the following attributes associated with the robot?

	1 (not at all)	2	3	4 (a moderate amount)	5	6	7 (very much so)
Dangerous	○	○	○	○	○	○	○
Emotional	○	○	○	○	○	○	○
Compassionate	○	○	○	○	○	○	○
Aggressive	○	○	○	○	○	○	○
Capable	○	○	○	○	○	○	○
Happy	○	○	○	○	○	○	○
Strange	○	○	○	○	○	○	○
Organic	○	○	○	○	○	○	○
Scary	○	○	○	○	○	○	○
Responsive	○	○	○	○	○	○	○
Awkward	○	○	○	○	○	○	○
Knowledgeable	○	○	○	○	○	○	○
Sociable	○	○	○	○	○	○	○
Feeling	○	○	○	○	○	○	○
Competent	○	○	○	○	○	○	○
Reliable	○	○	○	○	○	○	○
Awful	○	○	○	○	○	○	○
Interactive	○	○	○	○	○	○	○
	1 (not at all)	2	3	4 (a moderate amount)	5	6	7 (very much so)

1

2. Please rate your impression of the robot on these scales:

	1	2	3	4	5	
Machinelike	○	○	○	○	○	Humanlike
Ignorant	○	○	○	○	○	Knowledgeable
Awful	○	○	○	○	○	Nice
Unkind	○	○	○	○	○	Kind
Unintelligent	○	○	○	○	○	Intelligent
Dislike	○	○	○	○	○	Like
Dead	○	○	○	○	○	Alive
Foolish	○	○	○	○	○	Sensible
Unconscious	○	○	○	○	○	Conscious
Artificial	○	○	○	○	○	Lifelike
Fake	○	○	○	○	○	Natural
Moving rigidly	○	○	○	○	○	Moving elegantly
Incompetent	○	○	○	○	○	Competent
Mechanical	○	○	○	○	○	Organic
Stagnant	○	○	○	○	○	Lively
Unpleasant	○	○	○	○	○	Pleasant
Apathetic	○	○	○	○	○	Responsive
Unfriendly	○	○	○	○	○	Friendly
Irresponsible	○	○	○	○	○	Responsible
Inert	○	○	○	○	○	Interactive

3. Please rate your emotional state on these scales:

	1	2	3	4	5	
Agitated	○	○	○	○	○	Calm
Anxious	○	○	○	○	○	Relaxed
Quiescent	○	○	○	○	○	Surprised

2

4. Was the behaviour of the robot different in comparison to the previous interaction:

	1	2	3	4	5	
Not at all	○	○	○	○	○	Very much so

5. Was the behaviour of the robot more adaptive in comparison to the previous interaction:

	1	2	3	4	5	
Not at all	○	○	○	○	○	Very much so

6. Which words describe the behaviour of the robot best? (up to 5 words only)

7. (Optional) Please use the following space if you want to say more about **how** the behaviour of the robot(s) differ:

3

B.3. Study III

ID: ____________________

Pre-test questionnaire

Before we can start, we need a few more information about your person.

1. Gender: _________________
2. Age: ____________________
3. Occupation:____

How familiar are you with interacting with robots:

	1	2	3	4	5	
Not familiar	○	○	○	○	○	Very familiar

How familiar are you with programming robots:

	1	2	3	4	5	
Not familiar	○	○	○	○	○	Very familiar

How familiar are you with the BB8 robot from Sphero (the robot used in this experiment):

	1	2	3	4	5	
Not familiar	○	○	○	○	○	Very familiar

How familiar are you with “Star Wars”:

	1	2	3	4	5	
Not familiar	○	○	○	○	○	Very familiar

1

ID: ____________________

Post-test questionnaire

In the previous session you have interacted with the spherical robot. We are now interested to find out how you experienced and perceived the interaction and the robot itself.

Please tick only one box (✓) per row. Do not skip any questions and answer them in the given order. There is **no right** or **wrong**.

1. Using the scale provided, how closely are the following attributes associated with the robot?

	1 (not at all)	2	3	4 (a moderate amount)	5	6	7 (very much so)
Dangerous	○	○	○	○	○	○	○
Emotional	○	○	○	○	○	○	○
Compassionate	○	○	○	○	○	○	○
Aggressive	○	○	○	○	○	○	○
Capable	○	○	○	○	○	○	○
Happy	○	○	○	○	○	○	○
Persistent	○	○	○	○	○	○	○
Strange	○	○	○	○	○	○	○
Organic	○	○	○	○	○	○	○
Scary	○	○	○	○	○	○	○
Responsive	○	○	○	○	○	○	○
Awkward	○	○	○	○	○	○	○
Knowledgeable	○	○	○	○	○	○	○
Sociable	○	○	○	○	○	○	○
Feeling	○	○	○	○	○	○	○
Competent	○	○	○	○	○	○	○
Predictable	○	○	○	○	○	○	○
Reliable	○	○	○	○	○	○	○
Awful	○	○	○	○	○	○	○
Interactive	○	○	○	○	○	○	○
	1 (not at all)	2	3	4 (a moderate amount)	5	6	7 (very much so)

1

2. Please rate your impression of the robot on these scales:

	1	2	3	4	5	
Machinelike	○	○	○	○	○	Humanlike
Ignorant	○	○	○	○	○	Knowledgeable
Awful	○	○	○	○	○	Nice
Unkind	○	○	○	○	○	Kind
Unintelligent	○	○	○	○	○	Intelligent
Dislike	○	○	○	○	○	Like
Dead	○	○	○	○	○	Alive
Foolish	○	○	○	○	○	Sensible
Unconscious	○	○	○	○	○	Conscious
Artificial	○	○	○	○	○	Lifelike
Fake	○	○	○	○	○	Natural
Moving rigidly	○	○	○	○	○	Moving elegantly
Incompetent	○	○	○	○	○	Competent
Mechanical	○	○	○	○	○	Organic
Stagnant	○	○	○	○	○	Lively
Unpleasant	○	○	○	○	○	Pleasant
Apathetic	○	○	○	○	○	Responsive
Unfriendly	○	○	○	○	○	Friendly
Irresponsible	○	○	○	○	○	Responsible
Inert	○	○	○	○	○	Interactive

3. Please rate your emotional state on these scales:

	1	2	3	4	5	
Agitated	○	○	○	○	○	Calm
Anxious	○	○	○	○	○	Relaxed
Quiescent	○	○	○	○	○	Surprised

2

4. Which words describe the behaviour of the robot best? (up to 5 words only)

ID: ___________________

Post-test questionnaire 2

In the previous session you have interacted with the spherical robot. We are now interested to find out how you experienced and perceived the interaction and the robot itself.

Please tick only one box (✓) per row. Do not skip any questions and answer them in the given order. There is **no right** or **wrong**.

1. Using the scale provided, how closely are the following attributes associated with the robot?

	1 (not at all)	2	3	4 (a moderate amount)	5	6	7 (very much so)
Dangerous	○	○	○	○	○	○	○
Emotional	○	○	○	○	○	○	○
Compassionate	○	○	○	○	○	○	○
Aggressive	○	○	○	○	○	○	○
Capable	○	○	○	○	○	○	○
Happy	○	○	○	○	○	○	○
Persistent	○	○	○	○	○	○	○
Strange	○	○	○	○	○	○	○
Organic	○	○	○	○	○	○	○
Scary	○	○	○	○	○	○	○
Responsive	○	○	○	○	○	○	○
Awkward	○	○	○	○	○	○	○
Knowledgeable	○	○	○	○	○	○	○
Sociable	○	○	○	○	○	○	○
Feeling	○	○	○	○	○	○	○
Competent	○	○	○	○	○	○	○
Predictable	○	○	○	○	○	○	○
Reliable	○	○	○	○	○	○	○
Awful	○	○	○	○	○	○	○
Interactive	○	○	○	○	○	○	○
	1 (not at all)	2	3	4 (a moderate amount)	5	6	7 (very much so)

1

2. Please rate your impression of the robot on these scales:

	1	2	3	4	5	
Machinelike	○	○	○	○	○	Humanlike
Ignorant	○	○	○	○	○	Knowledgeable
Awful	○	○	○	○	○	Nice
Unkind	○	○	○	○	○	Kind
Unintelligent	○	○	○	○	○	Intelligent
Dislike	○	○	○	○	○	Like
Dead	○	○	○	○	○	Alive
Foolish	○	○	○	○	○	Sensible
Unconscious	○	○	○	○	○	Conscious
Artificial	○	○	○	○	○	Lifelike
Fake	○	○	○	○	○	Natural
Moving rigidly	○	○	○	○	○	Moving elegantly
Incompetent	○	○	○	○	○	Competent
Mechanical	○	○	○	○	○	Organic
Stagnant	○	○	○	○	○	Lively
Unpleasant	○	○	○	○	○	Pleasant
Apathetic	○	○	○	○	○	Responsive
Unfriendly	○	○	○	○	○	Friendly
Irresponsible	○	○	○	○	○	Responsible
Inert	○	○	○	○	○	Interactive

3. Please rate your emotional state on these scales:

	1	2	3	4	5	
Agitated	○	○	○	○	○	Calm
Anxious	○	○	○	○	○	Relaxed
Quiescent	○	○	○	○	○	Surprised

2

4. Was the behaviour of the robot different in comparison to the first interaction:

	1	2	3	4	5	
Not at all	○	○	○	○	○	Very much so

5. Which words describe the behaviour of the robot best? (up to 5 words only)

3

ID: ___________________

Post-test questionnaire 3

In the previous session you have interacted with the spherical robot. We are now interested to find out how you experienced and perceived the interaction and the robot itself.

Please tick only one box (✓) per row. Do not skip any questions and answer them in the given order. There is **no right** or **wrong**.

---

1. Using the scale provided, how closely are the following attributes associated with the robot?

	1 (not at all)	2	3	4 (a moderate amount)	5	6	7 (very much so)
Dangerous	○	○	○	○	○	○	○
Emotional	○	○	○	○	○	○	○
Compassionate	○	○	○	○	○	○	○
Aggressive	○	○	○	○	○	○	○
Capable	○	○	○	○	○	○	○
Happy	○	○	○	○	○	○	○
Persistent	○	○	○	○	○	○	○
Strange	○	○	○	○	○	○	○
Organic	○	○	○	○	○	○	○
Scary	○	○	○	○	○	○	○
Responsive	○	○	○	○	○	○	○
Awkward	○	○	○	○	○	○	○
Knowledgeable	○	○	○	○	○	○	○
Sociable	○	○	○	○	○	○	○
Feeling	○	○	○	○	○	○	○
Competent	○	○	○	○	○	○	○
Predictable	○	○	○	○	○	○	○
Reliable	○	○	○	○	○	○	○
Awful	○	○	○	○	○	○	○
Interactive	○	○	○	○	○	○	○
	1 (not at all)	2	3	4 (a moderate amount)	5	6	7 (very much so)

1

2. Please rate your impression of the robot on these scales:

	1	2	3	4	5	
Machinelike	○	○	○	○	○	Humanlike
Ignorant	○	○	○	○	○	Knowledgeable
Awful	○	○	○	○	○	Nice
Unkind	○	○	○	○	○	Kind
Unintelligent	○	○	○	○	○	Intelligent
Dislike	○	○	○	○	○	Like
Dead	○	○	○	○	○	Alive
Foolish	○	○	○	○	○	Sensible
Unconscious	○	○	○	○	○	Conscious
Artificial	○	○	○	○	○	Lifelike
Fake	○	○	○	○	○	Natural
Moving rigidly	○	○	○	○	○	Moving elegantly
Incompetent	○	○	○	○	○	Competent
Mechanical	○	○	○	○	○	Organic
Stagnant	○	○	○	○	○	Lively
Unpleasant	○	○	○	○	○	Pleasant
Apathetic	○	○	○	○	○	Responsive
Unfriendly	○	○	○	○	○	Friendly
Irresponsible	○	○	○	○	○	Responsible
Inert	○	○	○	○	○	Interactive

3. Please rate your emotional state on these scales:

	1	2	3	4	5	
Agitated	○	○	○	○	○	Calm
Anxious	○	○	○	○	○	Relaxed
Quiescent	○	○	○	○	○	Surprised

2

4. If you could interact with one of the robots again, which one would you chose:

1	2	3	no preference
○	○	○	○

5. Was the behaviour of the robot different in comparison to the **first** interaction:

	1	2	3	4	5	
Not at all	○	○	○	○	○	Very much so

6. Was the behaviour of the robot different in comparison to the **second** interaction:

	1	2	3	4	5	
Not at all	○	○	○	○	○	Very much so

7. Which words describe the behaviour of the robot best? (up to 5 words only)

8. (Optional) Please use the following space if you want to say more about **how** the behaviour of the robot(s) differ:

3

Appendix C.

Ethics approvals

C.1. Study I: Ethics Approval of amendments

University of Hertfordshire UH

HEALTH SCIENCE ENGINEERING & TECHNOLOGY ECDA

ETHICS APPROVAL NOTIFICATION

TO	Marcus Scheunemann
CC	Kerstin Dautenhahn
FROM	Dr Simon Trainis, Health, Science, Engineering & Technology ECDA Chairman
DATE	26/04/2018

Protocol number: **aCOM/PGR/UH/03018(1)**

Title of study: Perception of autonomous robots.

Your application to modify and extend the existing protocol as detailed below has been accepted and approved by the ECDA for your School and includes work undertaken for this study by the named additional workers below:

Modification: Amendments to questionnaire as per EC2 Form.

This approval is valid:

From: 26/04/2018

To: 01/11/2018

Additional workers: no additional workers named

Please note:

If your research involves invasive procedures you are required to complete and submit an EC7 Protocol Monitoring Form, and your completed consent paperwork to this ECDA once your study is complete. You are also required to complete and submit an EC7 Protocol Monitoring Form if you are a member of staff. This form is available via the Ethics Approval StudyNet Site via the 'Application Forms' page http://www.studynet1.herts.ac.uk/ptl/common/ethics.nsf/Teaching+Documents?Openview&count=9999&restricttocategory=Application+Forms

Any conditions relating to the original protocol approval remain and must be complied with.

Any necessary permissions for the use of premises/location and accessing participants for your study must be obtained in writing prior to any data collection commencing. Failure to obtain adequate permissions may be considered a breach of this protocol.

Approval applies specifically to the research study/methodology and timings as detailed in your Form EC1/EC1A or as detailed in the EC2 request. Should you amend any further aspect of your research, or wish to apply for an extension to your study,

C.2. Study II: Ethics Approval of amendments

University of Hertfordshire UH

HEALTH SCIENCE ENGINEERING & TECHNOLOGY ECDA

ETHICS APPROVAL NOTIFICATION

TO	Marcus Scheunemann
CC	Kerstin Dautenhahn
FROM	Dr Simon Trainis, Health, Science, Engineering & Technology ECDA Chairman
DATE	06/02/19

Protocol number: aCOM/PGR/UH/03018(3)

Title of study: Perception of autonomous robots

Your application to modify and extend the existing protocol as detailed below has been accepted and approved by the ECDA for your School and includes work undertaken for this study by the named additional workers below:

Modification: Modifications as entered on the EC2 application

This approval is valid:

From: 06/02/19

To: 24/11/19

Additional workers: no additional workers named

Please note:

If your research involves invasive procedures you are required to complete and submit an EC7 Protocol Monitoring Form, and your completed consent paperwork to this ECDA once your study is complete. You are also required to complete and submit an EC7 Protocol Monitoring Form if you are a member of staff. This form is available via the Ethics Approval StudyNet Site via the 'Application Forms' page http://www.studynet1.herts.ac.uk/ptl/common/ethics.nsf/Teaching+Documents?Openview&count=9999&restricttocategory=Application+Forms

Any conditions relating to the original protocol approval remain and must be complied with.

Any necessary permissions for the use of premises/location and accessing participants for your study must be obtained in writing prior to any data collection commencing. Failure to obtain adequate permissions may be considered a breach of this protocol.

Approval applies specifically to the research study/methodology and timings as detailed in your Form EC1/EC1A or as detailed in the EC2 request. Should you amend any further aspect of your research, or wish to apply for an extension to your study,

C.3. Study III: Ethics Approval of amendments

University of Hertfordshire UH

HEALTH SCIENCE ENGINEERING & TECHNOLOGY ECDA

ETHICS APPROVAL NOTIFICATION

TO	Marcus Scheunemann
CC	Kerstin Dautenhahn
FROM	Dr Simon Trainis, Health, Science, Engineering & Technology ECDA Chair.
DATE	04/03/2019

Protocol number: **aCOM/PGR/UH/03018(4)**

Title of study: Perception of autonomous robots

Your application to modify and extend the existing protocol as detailed below has been accepted and approved by the ECDA for your School and includes work undertaken for this study by the named additional workers below:

Modification: Detailed in EC2.

This approval is valid:

From: 04/03/2019

To: 24/11/2019

Additional workers: **no additional workers named**

Please note:

If your research involves invasive procedures you are required to complete and submit an EC7 Protocol Monitoring Form, and your completed consent paperwork to this ECDA once your study is complete. You are also required to complete and submit an EC7 Protocol Monitoring Form if you are a member of staff. This form is available via the Ethics Approval StudyNet Site via the 'Application Forms' page http://www.studynet1.herts.ac.uk/ptl/common/ethics.nsf/Teaching+Documents?Openview&count=9999&restricttocategory=Application+Forms

Any conditions relating to the original protocol approval remain and must be complied with.

Any necessary permissions for the use of premises/location and accessing participants for your study must be obtained in writing prior to any data collection commencing. Failure to obtain adequate permissions may be considered a breach of this protocol.

Approval applies specifically to the research study/methodology and timings as detailed in your Form EC1/EC1A or as detailed in the EC2 request. Should you amend any further aspect of your research, or wish to apply for an extension to your study, you will need your supervisor's approval (if you are a student) and must complete and

Acronyms

BLE Bluetooth Low Energy. 15, 17, 20, 23, 40, 56, 58, 61–65, 67, 68, 113, 119, 120, 145, 152, 154, 158

DOF degrees of freedom. 31, 38, 39

HRI human-robot interaction. 3, 4, 9–18, 20, 21, 23–25, 27, 33–38, 55, 59, 67, 69–71, 74, 77, 80, 90, 93, 95, 98, 99, 102, 103, 110, 113, 117, 118, 120, 123, 143, 146–152, 154–157, 159–162, 164, 167

IM intrinsic motivation. 3, 10–12, 14–19, 21, 23, 24, 26, 37, 38, 42, 44, 55, 69, 73, 90, 93, 110–113, 117, 118, 120, 121, 128, 143–146, 149, 151–154, 156, 157, 159–162, 164–167

IMU inertial measurement unit. 39, 40, 76, 77, 100

PI predictive information. 12, 18, 19, 23, 26–32, 57, 71–76, 87, 92, 96, 110, 116, 121, 151, 157, 166

RFID Radio-frequency identification. 59, 60, 68

RoSAS Robotic Social Attribute Scale. 3, 13, 35, 70, 71, 77, 78, 81, 82, 95, 97, 101, 106, 107, 122, 130, 139, 146, 155, 164

RSS received signal strength. 60–65, 119, 120

SDT Self-Determination Theory. 10, 25

TiPI time-local predictive information. 12–15, 19, 21, 27–34, 36, 38, 40–45, 52–55, 57, 69, 100, 113, 116, 120, 121, 144, 151, 152, 154, 156, 157, 166

References

Abele, Andrea E., Nicole Hauke, et al. (Nov. 2016). "Facets of the Fundamental Content Dimensions: Agency with Competence and Assertiveness—Communion with Warmth and Morality". In: *Frontiers in Psychology* 7. DOI: 10.3389/fpsyg.2016.01810 (pp. 13, 34, 35, 95, 110, 146).

Abele, Andrea E. and Bogdan Wojciszke (2007). "Agency and communion from the perspective of self versus others". In: *Journal of Personality and Social Psychology* 93.5, pp. 751–763. DOI: 10.1037/0022-3514.93.5.751 (pp. 16, 153).

– (2014). "Communal and Agentic Content in Social Cognition". In: *Advances in Experimental Social Psychology*. Elsevier, pp. 195–255. DOI: 10.1016/b978-0-12-800284-1.00004-7 (p. 35).

Ahmad, Bashar I. et al. (Apr. 2019). "Driver and Passenger Identification From Smartphone Data". In: *IEEE Transactions on Intelligent Transportation Systems* 20.4, pp. 1278–1288. DOI: 10.1109/tits.2018.2845113 (pp. 61, 68, 119).

Airenti, Gabriella (Jan. 2015). "The Cognitive Bases of Anthropomorphism: From Relatedness to Empathy". In: *International Journal of Social Robotics* 7.1, pp. 117–127. DOI: 10.1007/s12369-014-0263-x (p. 70).

Antunes, Rodolfo S. et al. (Apr. 2018). "A Survey of Sensors in Healthcare Workflow Monitoring". In: *ACM Computing Surveys* 51.2, pp. 1–37. DOI: 10.1145/3177852 (p. 68).

Apple Inc. (June 2014). *Getting Started with iBeacon*. Version 1.0. Apple (pp. 62, 63).

– (Sept. 2015). *Proximity Beacon Specification*. Release R1. Apple (p. 63).

Ay, Nihat, Holger Bernigau, et al. (Nov. 2012). "Information-driven self-organization: the dynamical system approach to autonomous robot behavior". In: *Theory in Biosciences* 131.3, pp. 161–179. ISSN: 1611-7530. DOI: 10.1007/s12064-011-0137-9 (p. 27).

Ay, Nihat, Nils Bertschinger, et al. (June 2008). "Predictive information and explorative behavior of autonomous robots". In: *The European Physical Journal B* 63.3, pp. 329–339. ISSN: 1434-6036. DOI: 10.1140/epjb/e2008-00175-0 (pp. 26–28, 33).

Bartneck, Christoph et al. (Jan. 2009). "Measurement Instruments for the Anthropomorphism, Animacy, Likeability, Perceived Intelligence, and Perceived Safety of Robots". In: *International Journal of Social Robotics* 1.1, pp. 71–81. ISSN: 1875-4805. DOI: 10.1007/s12369-008-0001-3 (pp. 70, 77).

Belpaeme, Tony et al. (Aug. 2018). "Social robots for education: A review". In: *Science Robotics* 3.21. DOI: 10.1126/scirobotics.aat5954 (p. 11).

Bialek, William et al. (Nov. 2001). "Predictability, Complexity, and Learning". In: *Neural Computation* 13.11, pp. 2409–2463. DOI: 10.1162/089976601753195969 (pp. 18, 27, 28, 157).

Bickmore, Timothy et al. (July 1, 2010). "Maintaining Engagement in Long-term Interventions with Relational Agents". In: *Applied Artificial Intelligence* 24.6, pp. 648–666. DOI: 10.1080/08839514.2010.492259 (pp. 13, 89).

Blow, Mike et al. (Sept. 2006). "Perception of Robot Smiles and Dimensions for Human-Robot Interaction Design". In: *The 15th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*. Hatfield, UK: IEEE, pp. 469–474. ISBN: 1-4244-0564-5. DOI: 10.1109/roman.2006.314372 (p. 72).

Bluetooth SIG (Feb. 2016). *Bluetooth 4.2 core specification*. URL: https://www.bluetooth.org/docman/handlers/downloaddoc.ashx?doc_id=441541 (visited on 12/08/2015) (pp. 60, 61, 64).

Boden, Margaret A. (Feb. 2008). "Autonomy: What is it?" In: *Biosystems* 91.2: *Modelling Autonomy*. Ed. by Xabier Barandiaran and Kepa Ruiz-Mirazo, pp. 305–308. ISSN: 0303-2647. DOI: 10.1016/j.biosystems.2007.07.003 (pp. 10, 12, 25).

Bodenhausen, Galen V et al. (2012). "Social Categorization and the Perception of Social Groups". In: *The SAGE Handbook of Social Cognition*. Ed. by Susan Fiske and C. Neil Macrae. SAGE Publications Sage UK: London, England, pp. 311–329. ISBN: 978-0-857-02481-7 (p. 34).

Braitenberg, Valentino (1984). *Vehicles. Experiments in Synthetic Psychology*. Cambridge: Bradford Books/MIT Press. IX, 152. ISBN: 978-0262521123 (p. 72).

Brooke, John (1986). "SUS: a 'quick and dirty' usability scale". In: *Usability Evaluation In Industry*. Ed. by Patrick W. Jordan et al. London: Taylor and Francis, pp. 189–194 (p. 164).

Calinon, Sylvain et al. (Oct. 2010). "Learning-based control strategy for safe human-robot interaction exploiting task and robot redundancies". In: *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE. DOI: 10.1109/iros.2010.5648931 (p. 58).

Carpinella, Colleen M. et al. (2017). "The Robotic Social Attributes Scale (RoSAS): Development and Validation". In: *Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction*. HRI '17. Vienna, Austria: ACM, pp. 254–262. ISBN: 978-1-4503-4336-7. DOI: 10.1145/2909824.3020208 (pp. 13, 35, 77, 89, 95, 109, 146).

Castro-González, Álvaro et al. (June 2016). "Effects of form and motion on judgments of social robots' animacy, likability, trustworthiness and unpleasantness". In: *International Journal of Human-Computer Studies* 90, pp. 27–38. DOI: 10.1016/j.ijhcs.2016.02.004 (pp. 17, 110, 153).

Ceha, Jessy et al. (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 (pp. 10, 11, 25, 161).

Christensen, Henrik I. et al. (June 29, 2016). “Next Generation Robotics”. In: arXiv: 1606.09205v1 [cs.CY] (p. 11).

Clabaugh, Caitlyn and Maja Matarić (May 2019). “Escaping Oz: Autonomy in Socially Assistive Robotics”. In: *Annual Review of Control, Robotics, and Autonomous Systems* 2.1, pp. 33–61. DOI: 10.1146/annurev-control-060117-104911 (p. 11).

Cohen, Jacob (July 1992). “A power primer”. In: *Psychological Bulletin* 112.1 (1), pp. 155–159. ISSN: 0033-2909 (pp. 84, 107, 131).

Cover, Thomas M. and Joy A. Thomas (2012). *Elements of Information Theory*. John Wiley & Sons (p. 30).

Crutchfield, James P. and Karl Young (July 1989). “Inferring statistical complexity”. In: *Physical Review Letters* 63.2, pp. 105–108. DOI: 10.1103/physrevlett.63.105 (p. 26).

Cuddy, Amy J. C. et al. (Apr. 2007). “The BIAS map: Behaviors from intergroup affect and stereotypes”. In: *Journal of Personality and Social Psychology* 92 (4), pp. 631–648. ISSN: 0022-3514. DOI: 10.1037/0022-3514.92.4.631 (pp. 13, 34).

Czopp, Alexander M. et al. (July 2015). “Positive Stereotypes Are Pervasive and Powerful”. In: *Perspectives on Psychological Science* 10.4, pp. 451–463. DOI: 10.1177/1745691615588091 (p. 34).

Dautenhahn, Kerstin (July 1997). “I could be you: The phenomenological dimension of social understanding”. In: *Cybernetics and Systems* 28.5, pp. 417–453. ISSN: 0196-9722. DOI: 10.1080/019697297126074 (p. 72).

– (2004). “Robots We Like to Live With?! – A Developmental Perspective on a Personalized, Life-Long Robot Companion”. In: *13th IEEE International Workshop on Robot and Human Interactive Communication*. IEEE, pp. 17–22. DOI: 10.1109/roman.2004.1374720 (pp. 38, 166).

Der, Ralf, Frank Güttler, et al. (Jan. 2008). “Predictive information and emergent cooperativity in a chain of mobile robots”. In: *The 11th International Conference on the Synthesis and Simulation of Living Systems (Artificial Life XI)*. Ed. by Seth Bullock et al. MIT Press, pp. 166–172 (pp. 12, 26, 27, 31).

Der, Ralf and Georg Martius (2006). “From Motor Babbling to Purposive Actions: Emerging Self-exploration in a Dynamical Systems Approach to Early Robot Development”. In: *From Animals to Animats 9*. Ed. by Stefano Nolfi et al. Vol. 4095. Lecture Notes in Computer Science. Springer Berlin Heidelberg, pp. 406–421. ISBN: 978-3-540-38615-5. DOI: 10.1007/11840541_34 (p. 28).

– (Sept. 2012). *The Playful Machine. Theoretical Foundation and Practical Realization of Self-Organizing Robots*. Vol. 15. Cognitive Systems Monographs. Springer-Verlag Berlin Heidelberg. 336 pp. ISBN: 978-3-642-20253-7. DOI: 10.1007/978-3-642-20253-7 (pp. 12, 26, 33, 69).

Dijk, Sander G. van and Marcus M. Scheunemann (Aug. 2019). “Deep Learning for Semantic Segmentation on Minimal Hardware”. In: *RoboCup 2018: Robot World Cup XXII*. Ed. by

Dirk Holz et al. Vol. 11374. Lecture Notes in Computer Science. Springer International Publishing, pp. 349–361. ISBN: 978-3-030-27544-0. DOI: 10.1007/978-3-030-27544-0_29 (pp. 59, 170).

Do, Trinh Minh Tri and Daniel Gatica-Perez (June 2011). "GroupUs: Smartphone Proximity Data and Human Interaction Type Mining". In: *2011 15th Annual International Symposium on Wearable Computers*. IEEE. IEEE, pp. 21–28. DOI: 10.1109/ISWC.2011.28 (p. 60).

Eyssel, Friederike and Frank Hegel (July 2012). "(S)he's Got the Look: Gender Stereotyping of Robots1". In: *Journal of Applied Social Psychology* 42.9, pp. 2213–2230. DOI: 10.1111/j.1559-1816.2012.00937.x (p. 35).

Faragher, Ramsey and Robert Harle (Nov. 2015). "Location Fingerprinting With Bluetooth Low Energy Beacons". In: *IEEE Journal on Selected Areas in Communications* 33.11, pp. 2418–2428. ISSN: 0733-8716. DOI: 10.1109/JSAC.2015.2430281 (pp. 61, 63, 65, 68, 120).

Feng, Tao et al. (2018). "A Human-Tracking Robot Using Ultra Wideband Technology". In: *IEEE Access* 6, pp. 42541–42550. DOI: 10.1109/access.2018.2859754 (p. 59).

Fiske, Susan T. (Feb. 2018). "Stereotype Content: Warmth and Competence Endure". In: *Current Directions in Psychological Science* 27.2, pp. 67–73. DOI: 10.1177/0963721417738825 (pp. 34, 35, 89, 95).

Fiske, Susan T. et al. (Feb. 2007). "Universal dimensions of social cognition: warmth and competence". In: *Trends in Cognitive Sciences* 11 (2), pp. 77–83. DOI: 10.1016/j.tics.2006.11.005 (pp. 13, 16, 34, 35, 70, 89, 93, 95, 96, 110, 146, 148, 153).

Ford, Clay (Oct. 17, 2015). *Is R-squared Useless?* URL: https://data.library.virginia.edu/is-r-squared-useless/ (visited on 12/12/2020) (p. 49).

Friston, Karl (Jan. 13, 2010). "The free-energy principle: a unified brain theory?" In: *Nature Reviews Neuroscience* 11 (2), pp. 127–138. ISSN: 1471-0048. DOI: 10.1038/nrn2787 (p. 12).

Froese, Tom and Tom Ziemke (Mar. 2009). "Enactive artificial intelligence: Investigating the systemic organization of life and mind". In: *Artificial Intelligence* 173.3-4, pp. 466–500. DOI: 10.1016/j.artint.2008.12.001 (pp. 10, 25).

Fukuda, Haruaki and Kazuhiro Ueda (Mar. 2, 2010). "Interaction with a Moving Object Affects One's Perception of Its Animacy". In: *International Journal of Social Robotics* 2.2, pp. 187–193. DOI: 10.1007/s12369-010-0045-z (pp. 13, 70, 74, 110).

Ghafurian, Moojan et al. (2020). *Social Companion Robots to Reduce Isolation: A Perception Change Due to COVID-19*. arXiv: 2008.05382 [cs.CY] (p. 165).

Gordon, Goren et al. (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 (pp. 10, 11, 25, 161).

Graaf, Maartje M.A. de et al. (Dec. 2016). "Long-term evaluation of a social robot in real homes". In: *Interaction Studies* 17.3, pp. 461–490. DOI: 10.1075/is.17.3.08deg (p. 162).

Grassberger, Peter (Sept. 1986). "Toward a quantitative theory of self-generated complexity". In: *International Journal of Theoretical Physics* 25.9, pp. 907–938. DOI: 10.1007/bf00668821 (p. 26).

Grillon, Christian et al. (2004). "Anxious Responses to Predictable and Unpredictable Aversive Events". In: *Behavioral Neuroscience* 118.5, pp. 916–924. DOI: 10.1037/0735-7044.118.5.916 (pp. 13, 89).

Guckelsberger, Christian et al. (Sept. 2016). "Intrinsically motivated general companion NPCs via Coupled Empowerment Maximisation". In: *2016 IEEE Conference on Computational Intelligence and Games (CIG)*. IEEE, pp. 150–157. ISBN: 978-1-5090-1883-3. DOI: 10.1109/cig.2016.7860406 (p. 12).

Hoffman, Guy and Wendy Ju (Feb. 2014). "Designing Robots With Movement in Mind". In: *Journal of Human-Robot Interaction* 3.1, pp. 91–122. ISSN: 2163-0364. DOI: 10.5898/jhri.3.1.hoffman (p. 72).

Huang, Hui-Min et al. (2004). "Autonomy Measures for Robots". In: *ASME 2004 International Mechanical Engineering Congress and Exposition*. American Society of Mechanical Engineers Digital Collection, pp. 1241–1247. DOI: 10.1115/IMECE2004-61812 (pp. 10, 25).

Hull, Clark L (1943). *Principles of behavior: an introduction to behavior theory*. Vol. 422. New York: Appleton-Century-Crofts (pp. 10, 23).

Iio, Takamasa et al. (Sept. 2019). "Human-Like Guide Robot that Proactively Explains Exhibits". In: *International Journal of Social Robotics* 12.2, pp. 549–566. DOI: 10.1007/s12369-019-00587-y (p. 162).

Ishiguro, Hiroshi (Jan. 2007). "Scientific Issues Concerning Androids". In: *The International Journal of Robotics Research* 26.1, pp. 105–117. DOI: 10.1177/0278364907074474 (p. 74).

José, Rui et al. (Oct. 2008). "Instant Places: Using Bluetooth for Situated Interaction in Public Displays". In: *IEEE Pervasive Computing* 7.4, pp. 52–57. ISSN: 1536-1268. DOI: 10.1109/mprv.2008.74 (p. 60).

Judd, Charles M. et al. (2005). "Fundamental dimensions of social judgment: Understanding the relations between judgments of competence and warmth". In: *Journal of Personality and Social Psychology* 89.6, pp. 899–913. DOI: 10.1037/0022-3514.89.6.899 (p. 34).

Kanda, Takayuki et al. (Oct. 2010). "A Communication Robot in a Shopping Mall". In: *IEEE Transactions on Robotics* 26.5, pp. 897–913. ISSN: 1941-0468. DOI: 10.1109/tro.2010.2062550 (pp. 10, 11).

Kaplan, Frédéric and Pierre-Yves Oudeyer (2004). "Maximizing Learning Progress: An Internal Reward System for Development". In: *Embodied Artificial Intelligence*. Ed. by F. Iida et al. Vol. 3139. Lecture Notes in Computer Science. Springer Berlin Heidelberg, pp. 259–270. ISBN: 978-3-540-27833-7. DOI: 10.1007/978-3-540-27833-7_19 (p. 12).

Khoramshahi, Mahdi et al. (June 2016). "Role of Gaze Cues in Interpersonal Motor Coordination: Towards Higher Affiliation in Human-Robot Interaction". In: *PLOS ONE* 11.6. Ed. by Eugene V Aidman, e0156874. DOI: 10.1371/journal.pone.0156874 (p. 58).

Klyubin, Alexander S. et al. (2005). "Empowerment: A Universal Agent-Centric Measure of Control". In: *2005 IEEE Congress on Evolutionary Computation*. Vol. 1. IEEE. IEEE, pp. 128–135. DOI: 10.1109/cec.2005.1554676 (p. 12).

Kulms, Philipp and Stefan Kopp (June 2018). "A Social Cognition Perspective on Human–Computer Trust: The Effect of Perceived Warmth and Competence on Trust in Decision-Making With Computers". In: *Frontiers in Digital Humanities* 5. DOI: 10.3389/fdigh.2018.00014 (p. 95).

Lane, Geoffrey W. et al. (Aug. 2016). "Effectiveness of a social robot, "Paro," in a VA long-term care setting." In: *Psychological Services* 13.3, pp. 292–299. DOI: 10.1037/ser0000080 (p. 165).

Law, Edith et al. (Aug. 2017). "A Wizard-of-Oz study of curiosity in human-robot interaction". In: *2017 26th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*. IEEE. DOI: 10.1109/roman.2017.8172365 (p. 161).

Lee, Jae et al. (Oct. 2006). "People Tracking Using a Robot in Motion with Laser Range Finder". In: *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE. DOI: 10.1109/iros.2006.282147 (p. 58).

Leigh, Angus et al. (May 2015). "Person tracking and following with 2D laser scanners". In: *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. DOI: 10.1109/icra.2015.7139259 (p. 58).

Lucasfilm Ltd. (2015). *Star Wars*. URL: https://www.starwars.com/films/ (visited on 11/10/2020) (pp. 36, 74, 125).

Ma, Yunfei et al. (2017). "Minding the Billions: Ultra-wideband Localization for Deployed RFID Tags". In: *Proceedings of the 23rd Annual International Conference on Mobile Computing and Networking - MobiCom '17*. ACM Press. DOI: 10.1145/3117811.3117833 (p. 60).

Martinelli, Francesco (Sept. 2015). "A Robot Localization System Combining RSSI and Phase Shift in UHF-RFID Signals". In: *IEEE Transactions on Control Systems Technology* 23.5, pp. 1782–1796. DOI: 10.1109/tcst.2014.2386777 (p. 60).

Martius, Georg (Nov. 22, 2013). *Implementation of predictive information maximization*. URL: https://github.com/georgmartius/lpzrobots/blob/d2e6bbd164d902cdaa57eef154ed353ee0027236/selforg/controller/pimax.cpp (visited on 12/14/2020) (p. 30).

Martius, Georg, Ralf Der, et al. (2013a). *Appendix Text S1 for Information driven self-organization of complex robotic behaviors with derivations and technical detail.* (P. 30).

– (May 2013b). "Information Driven Self-Organization of Complex Robotic Behaviors". In: *PLOS ONE* 8.5. Ed. by Josh Bongard, pp. 1–14. DOI: 10.1371/journal.pone.0063400 (pp. 12, 15, 18, 19, 26–29, 31, 33, 157, 206).

– (2013c). *Supplementary material for "Information Driven Self-Organization of Complex Robotic Behaviors" Martius, Der, et al. 2013b: program code and videos*. URL: http://robot.informatik.uni-leipzig.de/research/supplementary/TiPI2013 (p. 26).

Martius, Georg, Luisa Jahn, et al. (2014). “Self-exploration of the Stumpy Robot with Predictive Information Maximization”. In: *From Animals to Animats 13*. Springer International Publishing, pp. 32–42. DOI: 10.1007/978-3-319-08864-8_4 (pp. 12, 26, 33).

Maturana, Humberto R. and Francisco J. Varela (1991). *Autopoiesis and Cognition: The Realization of the Living*. Ed. by Robert S. Cohen and Marx W. Wartofsky. Vol. 42. Springer Science & Business Media. ISBN: 978-94-009-8947-4. DOI: 10.1007/978-94-009-8947-4 (pp. 10, 25).

Megalou, S. et al. (Mar. 2019). “Fingerprinting Localization of RFID tags with Real-Time Performance-Assessment, using a Moving Robot”. In: *2019 13th European Conference on Antennas and Propagation (EuCAP)*, pp. 1–5 (p. 60).

Merrick, Kathryn E. and Mary Lou Maher (2009). *Motivated Reinforcement Learning. Curious Characters for Multiuser Games*. 1st ed. Springer Berlin Heidelberg. XIV, 206. ISBN: 9783540891864. DOI: 10.1007/978-3-540-89187-1 (p. 12).

Michel, P. et al. (2006). “Online environment reconstruction for biped navigation”. In: *Proceedings 2006 IEEE International Conference on Robotics and Automation, 2006. ICRA 2006*. IEEE. DOI: 10.1109/robot.2006.1642171 (p. 58).

Mieczkowski, Hannah et al. (Mar. 2019). “Helping Not Hurting: Applying the Stereotype Content Model and BIAS Map to Social Robotics”. In: *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE. DOI: 10.1109/hri.2019.8673307 (pp. 35, 36).

Nehmzow, Ulrich (2006). *Scientific Methods in Mobile Robotics*. Vol. 24. 1-2. Springer-Verlag GmbH, pp. 1–3. ISBN: 9781846280191 (p. 140).

Ng, Pai Chet et al. (Aug. 2019). “A Compressive Sensing Approach to Detect the Proximity Between Smartphones and BLE Beacons”. In: *IEEE Internet of Things Journal* 6.4, pp. 7162–7174. DOI: 10.1109/jiot.2019.2914733 (p. 61).

Oliveira, Raquel et al. (Mar. 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 (pp. 35, 160).

OptiTrack (2020). *OptiTrack for Robotics*. URL: https://optitrack.com/applications/robotics/ (visited on 12/14/2020) (p. 58).

Orne, Martin T. (1962). “On the social psychology of the psychological experiment: With particular reference to demand characteristics and their implications.” In: *American Psychologist* 17.11, pp. 776–783. DOI: 10.1037/h0043424 (pp. 16, 123, 147, 153, 160).

Oudeyer, Pierre-Yves, J. Gottlieb, et al. (2016). “Intrinsic motivation, curiosity, and learning”. In: *Motivation - Theory, Neurobiology and Applications*. Elsevier, pp. 257–284. DOI: 10.1016/bs.pbr.2016.05.005 (pp. 10, 11, 24, 25).

Oudeyer, Pierre-Yves and Frédéric Kaplan (July 2008). “How can we define intrinsic motivation?” In: *Proceedings of the 8th International Conference on Epigenetic Robotics: Modeling

Cognitive Development in Robotic Systems. Vol. 5. Brighton, United Kingdom, pp. 93–101 (pp. 11, 24).

Oudeyer, Pierre-Yves and Frédéric Kaplan (Nov. 2, 2009). "What is Intrinsic Motivation? A Typology of Computational Approaches". In: *Frontiers in Neurorobotics* 1. Ed. by Max Lungarella, p. 6. ISSN: 1662-5218. DOI: 10.3389/neuro.12.006.2007 (p. 10).

Oudeyer, Pierre-Yves, Frédéric Kaplan, and Verena V. Hafner (Apr. 2, 2007). "Intrinsic Motivation Systems for Autonomous Mental Development". In: *IEEE Transactions on Evolutionary Computation* 11.2, pp. 265–286. ISSN: 1941-0026. DOI: 10.1109/tevc.2006.890271 (p. 12).

Paetzel, Maike et al. (Mar. 2020). "The Persistence of First Impressions". In: *Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction*. ACM. DOI: 10.1145/3319502.3374786 (pp. 160, 162).

Paolo, Ezequiel A. Di (June 2004). *Unbinding Biological Autonomy: Francisco Varela's Contributions to Artificial Life*. Vol. 10. 3. MIT Press - Journals, pp. 231–233. DOI: 10.1162/1064546041255566 (pp. 10, 25).

Pinillos, Roberto et al. (May 2016). "Long-term assessment of a service robot in a hotel environment". In: *Robotics and Autonomous Systems* 79, pp. 40–57. DOI: 10.1016/j.robot.2016.01.014 (pp. 9, 11).

Rault, Tifenn et al. (2014). "Energy efficiency in wireless sensor networks: A top-down survey". In: *Computer Networks* 67, pp. 104–122. ISSN: 1389-1286. DOI: 10.1016/j.comnet.2014.03.027 (p. 61).

Rosenthal, Robert et al. (1994). "Parametric measures of effect size". In: *The handbook of research synthesis* 621.2, pp. 231–244 (pp. 84, 107, 131).

Ryan, Richard M. and Edward L. Deci (Jan. 2000). "Intrinsic and Extrinsic Motivations: Classic Definitions and New Directions". In: *Contemporary Educational Psychology* 25.1, pp. 54–67. ISSN: 0361-476X. DOI: 10.1006/ceps.1999.1020 (pp. 10, 23, 24).

Sabanovic, Selma et al. (June 2013). "PARO robot affects diverse interaction modalities in group sensory therapy for older adults with dementia". In: *2013 IEEE 13th International Conference on Rehabilitation Robotics (ICORR)*. IEEE. DOI: 10.1109/icorr.2013.6650427 (p. 165).

Scassellati, Brian and Marynel Vázquez (July 2020). "The potential of socially assistive robots during infectious disease outbreaks". In: *Science Robotics* 5.44, p. 2. DOI: 10.1126/scirobotics.abc9014 (p. 165).

Scheunemann, Marcus M. (June 27, 2017a). *GitHub pull request: Bug/BLE packet parsing*. URL: https://github.com/sphero-inc/sphero.js/pull/91 (visited on 12/14/2020) (pp. 19, 40, 158).

– (June 8, 2017b). *GitHub pull request: Bug/documentation*. URL: https://github.com/sphero-inc/sphero.js/pull/86 (visited on 12/14/2020) (pp. 19, 40, 158).

– (June 16, 2017c). *GitHub pull request: Feature/enhance connection to multiple ble devices.* URL: https://github.com/sphero-inc/sphero.js/pull/89 (visited on 12/14/2020) (pp. 19, 40, 158).

– (June 14, 2017d). *GitHub pull request: Feature/streaming quaternion.* URL: https://github.com/sphero-inc/sphero.js/pull/88 (visited on 12/14/2020) (pp. 19, 40, 158).

– (2017e). *Project page: Intrinsically Motivated Robots.* Videos and code repositories. URL: https://mms.ai/imr/ (visited on 05/18/2020) (pp. 19, 40, 77, 101, 116, 121, 145, 158).

– (Aug. 15, 2018a). *Code repository: BLE4HRI.* URL: https://github.com/scheunemann/BLE4HRI (visited on 12/12/2020) (pp. 20, 64, 159).

– (Aug. 15, 2018b). *Code repository: spherocpp.* URL: https://gitlab.com/scheunemann/spheropp (visited on 12/12/2020) (pp. 19, 40, 158).

– (Feb. 8, 2018c). *Contributions to lpzrobots. GitHub pull request.* URL: https://github.com/georgmartius/lpzrobots/commits?author=scheunemann (visited on 12/14/2020) (pp. 19, 157).

– (Nov. 8, 2018d). *GitHub pull request: TiPI fixes.* URL: https://github.com/georgmartius/lpzrobots/pull/29 (visited on 12/14/2020) (pp. 19, 157).

– (Dec. 15, 2018e). *Project page: Bluetooth Low Energy for Human-Robot Interaction.* Videos and code repositories. URL: https://mms.ai/BLE4HRI (visited on 12/15/2020) (pp. 20, 64, 159).

– (Apr. 26, 2019). *Intrinsically Motivated Autonomy in Human-Robot Interaction: Human Perception of Predictive Information in Robots. Supplementary material.* URL: https://mms.ai/TAROS2019-supplementary/ (visited on 10/10/2019) (pp. 77, 101).

Scheunemann, Marcus M., Raymond H. Cuijpers, et al. (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).* IEEE. Virtual Conference: IEEE, pp. 1340–1347. DOI: 10.1109/RO-MAN47096.2020.9223478 (pp. 36, 162, 169).

Scheunemann, Marcus M. and Kerstin Dautenhahn (Feb. 2017). "Bluetooth Low Energy for Autonomous Human-Robot Interaction". In: *Proceedings of the Companion of the 2017 ACM/IEEE International Conference on Human-Robot Interaction.* HRI '17. Vienna, Austria: ACM, pp. 52–52. ISBN: 978-1-4503-4885-0. DOI: 10.1145/3029798.3036663 (p. 169).

Scheunemann, Marcus M., Kerstin Dautenhahn, Maha Salem, et al. (Aug. 2016a). "Utilizing Bluetooth Low Energy for human-robot interaction". In: *2016 25th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN).* New York, NY, USA: IEEE (p. 170).

– (Aug. 2016b). "Utilizing Bluetooth Low Energy to recognize proximity, touch and humans". In: *2016 25th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN).* New York, NY, USA: IEEE, pp. 362–367. DOI: 10.1109/ROMAN.2016.7745156 (pp. 20, 32, 62, 64, 68, 158, 169).

Scheunemann, Marcus M. and Sander G. van Dijk (2019). "ROS 2 for RoboCup". In: *RoboCup 2019: Robot World Cup XXIII*. Ed. by Stephan Chalup et al. Vol. 11531. Lecture Notes in Artificial Intelligence. Springer International Publishing, pp. 429–438. ISBN: 978-3-030-35699-6. DOI: 10.1007/978-3-030-35699-6_34 (p. 170).

Scheunemann, Marcus M., Christoph Salge, and Kerstin Dautenhahn (2019). "Intrinsically Motivated Autonomy in Human-Robot Interaction: Human Perception of Predictive Information in Robots". In: *Towards Autonomous Robotic Systems*. Ed. by Kaspar Althoefer et al. Cham: Springer International Publishing, pp. 325–337. ISBN: 978-3-030-23807-0. DOI: 10.1007/978-3-030-23807-0_27 (pp. 21, 162, 169).

Scheunemann, Marcus M., Christoph Salge, Daniel Polani, et al. (2021). *Human Perception of Intrinsically Motivated Autonomy in Human-Robot Interaction*. arXiv: 2002.05936 [cs.RO] (pp. 21, 162, 170).

Schmidhuber, Jürgen (1991). "Curious model-building control systems". In: *[Proceedings] 1991 IEEE International Joint Conference on Neural Networks*. Singapore: IEEE, pp. 1458–1463. DOI: 10.1109/ijcnn.1991.170605 (p. 11).

Schulz, Dirk et al. (Feb. 2003). "People Tracking with Mobile Robots Using Sample-Based Joint Probabilistic Data Association Filters". In: *The International Journal of Robotics Research* 22.2, pp. 99–116. DOI: 10.1177/0278364903022002002 (p. 58).

Schwarz, David et al. (2015). "Cosero, Find My Keys! Object Localization and Retrieval Using Bluetooth Low Energy Tags". In: *RoboCup 2014: Robot World Cup XVIII*. Ed. by Reinaldo A. C. Bianchi et al. Springer International Publishing, pp. 195–206. ISBN: 978-3-319-18615-3. DOI: 10.1007/978-3-319-18615-3_16 (pp. 61, 119, 120).

Siegemund, Frank and Christian Florkemeier (Mar. 2003). "Interaction in pervasive computing settings using Bluetooth-enabled active tags and passive RFID technology together with mobile phones". In: *Proceedings of the First IEEE International Conference on Pervasive Computing and Communications, 2003. (PerCom 2003)*. IEEE. IEEE Comput. Soc, pp. 378–387. DOI: 10.1109/PERCOM.2003.1192762 (p. 60).

Silicon Laboratories Inc. (Dec. 2015). *BLE112 DATA SHEET*. Version 1.48. login required. Silicon Laboratories (p. 64).

Sphero, Inc. (2020a). *Licensed products: BB-8*. URL: https://support.sphero.com/category/kxwbdyqeyq-bb-8 (visited on 12/12/2020) (pp. 36, 74).

– (2020b). *Sphero Edu JavaScript: Movement: Stabilization*. URL: https://sphero.docsapp.io/docs/movement#stabilization (visited on 12/12/2020) (p. 41).

Steels, Luc (July 5, 2004). "The Autotelic Principle". In: *Embodied Artificial Intelligence*. Vol. 3139. Lecture Notes in Computer Science. Springer Berlin Heidelberg, pp. 231–242. ISBN: 978-3-540-27833-7. DOI: 10.1007/978-3-540-27833-7_17 (p. 12).

Stroessner, Steven J. (2020). "On the social perception of robots: measurement, moderation, and implications". In: *Living with Robots*. Elsevier, pp. 21–47. DOI: 10.1016/b978-0-12-815367-3.00002-5 (pp. 35, 36).

Stubbs, K. et al. (Mar. 2007). “Autonomy and Common Ground in Human-Robot Interaction: A Field Study”. In: *IEEE Intelligent Systems* 22.2, pp. 42–50. DOI: 10.1109/MIS.2007.21 (pp. 10, 25).

Subhan, Fazli et al. (Apr. 2011). “Indoor positioning in Bluetooth networks using fingerprinting and lateration approach”. In: *2011 International Conference on Information Science and Applications*. IEEE, pp. 1–9. DOI: 10.1109/ICISA.2011.5772436 (p. 60).

Tay, Benedict et al. (Sept. 2014). “When stereotypes meet robots: The double-edge sword of robot gender and personality in human–robot interaction”. In: *Computers in Human Behavior* 38, pp. 75–84. DOI: 10.1016/j.chb.2014.05.014 (p. 35).

Texas Instruments (2013). *CC2540 SimpleLink™ Bluetooth® low energy Wireless MCU with USB datasheet (Rev. F)*. Version SWRS084F. First published 2010. Texas Instruments. 27 pp. (p. 64).

Theil, Henri (July 1970). “On the Estimation of Relationships Involving Qualitative Variables”. In: *American Journal of Sociology* 76.1, pp. 103–154. DOI: 10.1086/224909 (p. 140).

Thrun, Sebastian et al. (2005). *Probabilistic Robotics*. Intelligent Robotics and Autonomous Agents. MIT Press. 672 pp. ISBN: 978-0-262-30380-4 (pp. 29, 32).

Tremoulet, Patrice D and Jacob Feldman (Aug. 2000). “Perception of Animacy from the Motion of a Single Object”. In: *Perception* 29.8, pp. 943–951. DOI: 10.1068/p3101 (pp. 70, 110, 115, 144).

Truijens, M. et al. (2014). “Evaluating the Performance of Absolute RSSI Positioning Algorithm-Based Microzoning and RFID in Construction Materials Tracking”. In: *Mathematical Problems in Engineering* 2014, pp. 1–8. DOI: 10.1155/2014/784395 (p. 60).

Van Rensburg, E. J. Janse et al. (2001). “Self-averaging sequences in the statistical mechanics of random copolymers”. In: *Journal of Physics A: Mathematical and General* 34.33, p. 6381 (p. 33).

Wada, K.. and T.. Shibata (Oct. 2007). “Living With Seal Robots—Its Sociopsychological and Physiological Influences on the Elderly at a Care House”. In: *IEEE Transactions on Robotics* 23.5, pp. 972–980. DOI: 10.1109/tro.2007.906261 (pp. 164, 165).

Wainer, Joshua et al. (Sept. 2014). “Using the Humanoid Robot KASPAR to Autonomously Play Triadic Games and Facilitate Collaborative Play Among Children With Autism”. In: *IEEE Transactions on Autonomous Mental Development* 6.3, pp. 183–199. DOI: 10.1109/tamd.2014.2303116 (p. 9).

Walters, Michael L. et al. (Nov. 2007). “Avoiding the uncanny valley: robot appearance, personality and consistency of behavior in an attention-seeking home scenario for a robot companion”. In: *Autonomous Robots* 24.2, pp. 159–178. DOI: 10.1007/s10514-007-9058-3 (p. 70).

White, Robert W (1959). “Motivation reconsidered: The concept of competence.” In: *Psychological review* 66.5, p. 297 (pp. 10, 23).

Wojciszke, Bogdan, Andrea E. Abele, et al. (Feb. 2009). "Two dimensions of interpersonal attitudes: Liking depends on communion, respect depends on agency". In: *European Journal of Social Psychology* 39.6, pp. 973–990. DOI: 10.1002/ejsp.595 (p. 35).

Wojciszke, Bogdan, Roza Bazinska, et al. (Dec. 1998). "On the Dominance of Moral Categories in Impression Formation". In: *Personality and Social Psychology Bulletin* 24.12, pp. 1251–1263. DOI: 10.1177/01461672982412001 (pp. 34, 35).

Wood, Luke J. et al. (2017). "The Iterative Development of the Humanoid Robot Kaspar: An Assistive Robot for Children with Autism". In: *Social Robotics*. Springer International Publishing, pp. 53–63. DOI: 10.1007/978-3-319-70022-9_6 (p. 59).

Yatani, Koji (Mar. 23, 2016). "Effect Sizes and Power Analysis in HCI". In: *Modern Statistical Methods for HCI*. Ed. by Judy Robertson and Maurits Kaptein. Human–Computer Interaction Series. Springer. Chap. 5, pp. 87–110. ISBN: 978-3-319-26631-2. DOI: 10.1007/978-3-319-26633-6_5 (pp. 84, 107, 131).

Zahedi, Keyan et al. (2013). "Linear combination of one-step predictive information with an external reward in an episodic policy gradient setting: a critical analysis". In: *Frontiers in Psychology* 4. DOI: 10.3389/fpsyg.2013.00801 (p. 26).

Ziemke, Tom (Feb. 2008). "On the role of emotion in biological and robotic autonomy". In: *BioSystems* 91.2: *Modelling Autonomy*. Ed. by Xabier Barandiaran and Kepa Ruiz-Mirazo, pp. 401–408. ISSN: 0303-2647. DOI: 10.1016/j.biosystems.2007.05.015 (p. 10).