

Title	Prediction Error Minimization for the Emergence of Prosocial Helping Behavior				
Author(s) Baraglia, Jimmy Camille					
Citation 大阪大学, 2016, 博士論文					
Version Type VoR					
URL	https://doi.org/10.18910/59600				
rights					
Note					

The University of Osaka Institutional Knowledge Archive : OUKA

https://ir.library.osaka-u.ac.jp/

The University of Osaka

Doctoral Dissertation

Prediction Error Minimization for the Emergence of Prosocial Helping Behavior

Baraglia Jimmy Camille July 2016

A dissertation submitted to the department of Adaptive Machine Systems in partial fulfillment of the requirements for the degree of doctor of Philosophy in Engineering at the

GRADUATE SCHOOL OF ENGINEERING OSAKA UNIVERSITY

Abstract

The aims of the work presented in this dissertation are to understand the emergence and the development of prosocial helping abilities in infants and to reproduce it into robots. The ensuing main objectives of our studies are:

- To propose a general motivation mechanism that explains how infants may acquire prosocial helping abilities via the design of a computational model to reproduce such behavior in robot.
- To construct an assistive robot capable to fluently and efficiently interact with others in the context of joint tasks collaboration.

We propose a robotic model that build representations of its environment based on predictive learning and perform actions to help others motivated by the minimization of prediction error. By doing so, our systems become capable to develop cognitive and prosocial abilities such as helping behavior or collaboration, similarly to what infants are capable to do. We performed two studies that describe major steps from the development of prosocial helping abilities to efficient human-robot collaboration.

1. The first study presents a general mechanism explaining the emergence of prosocial helping behavior based on goal understanding. Traditional theories suggest that empathy and emotion contagion play capital roles in the motivation for early prosocial helping behavior. However, recent studies have shown that infants can similarly help simple animated shapes to achieve their goal, implying that another mechanism, more general, may be involve in the emergence of prosocial helping behavior. We suggest that when observing others actions, infants are capable to predict their goals based on self-experience and to perform actions if the goal is not achieved. In other words, infants help others to minimize prediction error. To evaluate our hypothesis, we design a computational model based on psychological studies and implement it in real and simulated robots. Our experimental results demonstrate that our robots could spontaneously generate prosocial helping behavior by being motivated by the minimization of prediction error.

2. The second study focuses on how to use previously described abilities to generate efficient and fluent human-robot interactions. We address the question of whether and when a robot should help to minimize prediction error during collaborative human-robot task execution. Based on our previous model, we design a robot capable to autonomously perform table-top manipulation tasks while monitoring the environmental state and human activity. To evaluate our system, we implement three different initiative conditions to trigger the robot's actions. Robot-initiated reactive help triggers robot assistance when prediction error is high; robot-initiated proactive help makes the robot help whenever it can, even with low prediction error; human-initiated help gives control of the robot action timing to the user; Our user study results (N=18) give us significant proofs that proactive robots perform best, while users prefer to be in control rather than interacting with the reactive robots.

Together, these two studies describe mechanisms explaining how a robot can develop prosocial helping behavior based on prediction error minimization and perform efficient and fluent collaboration. Our results contribute to the field of developmental science and robotics by proposing a general and likely mechanism for the development of prosocial helping abilities that can be used to create efficient assistive robots.

Acknowledgements

Foremost, I would like to express my sincere gratitude to my supervisors Prof. Nagai and Prof Asada for the continuous and never failing support during the course of my Ph.D. study and research. I greatly valued their patience, guidance, motivation and knowledge. Along my 4 years at the Osaka University, they allowed me to flourish in my field and gave me plenty of opportunity to widen my research horizons. I could not have imagined having better advisors.

Besides I would like to thank the rest of my thesis committee: Prof. Ishiguro, Prof. Hosoda, and Prof Rao for their encouragement, insightful comments, and hard questions that forced me to think more and deepen my thinking.

My sincere thanks also goes to Maya Cakmak and her students who welcomed me in their laboratory at the University of Washington and allowed me to have one of the best time ever. Together we achieved amazing results, boosted by their strong motivation and wisdom.

I thank my fellow lab mates in Asada Lab: Ryo, Jihoon, Takato, Hiroki Yokoyama, Jorge and JJ, but also the previous exchange students: Edouard, Paul, Francois, Manoracth, Benedict, Hagen and many more for the stimulating talks and all the fun. Also I thank all my friends, especially Roy, Shao, Mahdi, Amr and Omid, for allowing me to appreciate Japan even more and to give me opportunities to discover new cultures.

I would like to also thank my mother for her moral support; my father for his never failing positiveness; my brothers for always being there for me; my grandparents for guiding me, and the rest of my family for being simply the best and giving me the motivation to go ahead.

Last but not least, I would like to thank my wife Sookyeun for always being with me and for giving me happiness.

Contents

1	Introduction		
	1.1	Prosocial Behavior in Robots	16
	1.2	Targets and Key Ideas of our Study	17
	1.3	Overview of the thesis	18
2	\mathbf{Rel}	ated Work	22
	2.1	Prosocial Helping Behavior in Infants	22
		2.1.1 When does Prosocial Helping Start?	23
		2.1.2 Why do Infants Help Others?	24
	2.2	Helping Robots and Human-Robot Collaboration	26
3	A t	heory of Predictive Learning to Understand Infant Devel-	
	opn	nent	30
	3.1	Predictive Learning	30
	3.2	Overview of Our General Model	33
4	\mathbf{Em}	ergence of Prosocial Helping Behavior in Robots	36
	4.1	Hypothesis for the Emergence of Helping Behavior	37
	4.2	Model for the Minimization of Prediction Error	37
		4.2.1 Perception Module	39
		4.2.2 Predictive Module	39
		4.2.3 Minimization of Prediction Error Module	43
	4.3	Experiment 1- Setup	44
		4.3.1 Robots Actions	45
		4.3.2 Training and Observation Phase	45
	4.4	Experiment 1- Results	47

		4.4.1 Evaluation of Helping Performances
	4.5	Experiment 1 - Discussion
	4.6	Experiment 2 - Setup
		4.6.1 System Implementation
		4.6.2 Training and Observation Phase
	4.7	Experiment 2 - Results
		4.7.1 Evaluation of the Training Phase $\ldots \ldots \ldots \ldots \ldots$
		4.7.2 Evaluation of the Observation Phase $\ldots \ldots \ldots \ldots$
	4.8	Experiment 2 - Discussion
	4.9	General Discussion and Future Work
5	From	m Helping to Efficient Human-Robot Collaboration
	5.1	Task Representation and Scenario
	5.2	Joint Task Execution Model
		5.2.1 Perception Module $\ldots \ldots \ldots$
		5.2.2 Predictive Module
		5.2.3 Action Selection Module
	5.3	Experimental Setup
		5.3.1 Platform
		5.3.2 Robot Actions $\ldots \ldots \ldots \ldots \ldots \ldots \ldots \ldots \ldots \ldots \ldots$
	5.4	User Study
		5.4.1 Study Design \ldots \ldots \ldots \ldots \ldots \ldots \ldots \ldots \ldots
		5.4.2 Study Setup
		5.4.3 Procedure
		5.4.4 Measurement \ldots \ldots \ldots \ldots \ldots \ldots \ldots \ldots
	5.5	Experimental Results
		5.5.1 Objective Metric
		5.5.2 Subjective Metric
	5.6	Discussion and Future Work
6	Cor	clusion
	6.1	Summary of our Study
	6.2	Contributions
	6.3	Limitations and Future Work
		6.3.1 Extension of our Model

7	App	endix 10	2
	.1	Appendix A	2
	.2	Appendix B	4

List of Tables

4.1	Experiment 1: Six actions the system experienced	45
4.2	Experiment 1: List of others' failed actions $\ldots \ldots \ldots \ldots$	47
4.3	Experiment 2: List of action primitive	54
4.4	Experiment 2: Experimental results	55

List of Figures

1.1	Example of helping behavior shown by 18-month-old infants	15
2.1	Emergence of prosocial helping behavior and development of cog- nitive abilities	23
		-0
3.1	Overview for the predictive learning architecture	31
3.2	Example of robot using predictive learning	32
3.3	Overview for the proposed model	33
4.1	Model for the minimization of $PE \dots \dots \dots \dots \dots \dots \dots$	38
4.2	Example of graph construction	40
4.3	Example of PE estimation	44
4.4	Experiment 1: Example of action graph	46
4.5	Experiment 1: system performances	48
4.6	Experiment 2: experimental setting	50
4.7	Experiment 2: visual processing	52
4.8	Experiment 2: learning example	54
4.9	Experiment 2: graph after learning	55
4.10	Experiment 2: successful prosocial helping behavior	56
5.1	Different initiative models for robot assistance	61
5.2	Goal states for the two task categories used in our evaluation	62
5.3	Particular instances of the tasks used in the user study	63
5.4	Model for collaborative robots based on predictive learning \ldots	64
5.5	Examples of task knowledge represented as states transition	66
5.6	Examples of object detection likelihood estimation	67
5.7	Examples of interactions	75

5.8	Objective results: system performance	76
5.9	Objective results: breakdown of task completion times	77
5.10	Analyzed gazing zones	79
5.11	Objective results: gaze to robot	80
5.12	Subjective results: mean Likert-scale ratings	82

Chapter 1

Introduction

The ability to act prosocially and to help our conspecifics is viewed by many researchers as the roots for the development of higher social capabilities such as concern for others, collaboration and altruism [27, 82]. Human infants start to prosocially help others from their second year of life. From 14 months of age, they have been shown to help in simple scenario, for example giving an experimenter out-of-reach objects [90]. Later in their second year of life, as their cognitive and motor maturity increases, infants gradually start to help others in achieving more complex goals. They become able to, for instance, open a closet to allow an experimenter to store books or stack boxes correctly on a pile [88, 89, 66, 50, 73, 91] (see Fig. 1.1).

Why being social? Why is it important for humans (and robots) to behave prosocially toward our conspecifics? One obvious advantage is that it allows us to be part of a society in which we can count on each other's; help and being helped. If humans were not social, we would have no reasons not to be secluded, and we would fail to create the bonds that allow us to live, work and stay together. Then why humans, contrary to most other animal species, have abandoned selfishness and individual thinking when in a group and started being social? According to Darwin [18], being prosocial is part of an evolutionary process to deal with threats or increase reproductive capabilities. This theory of evolution can explain the epigenetic changes leading to more social humans, but not why infants are capable to help other from as early as 14 months of age.

The emergence of prosocial behavior observed in the early development of



Figure 1.1: Example of helping behavior shown by 18-month-old infants [88]. (a) The infant opens the cabinet door after seeing the adult trying to put a stack of book inside. (b) The infant hands a clothespin to the adult trying to reach for it. (c) The infant stacks boxes that the adult "accidentally" misplaced.

infants has, for many years, intrigued and still fascinates scientists in the fields of psychology, cognitive and developmental sciences, and neuroscience. To find out why and how infants become prosocial (or if they are from birth), scientists have performed a large number of psychological experiments. Some concluded that prosocial acts of help are based on the ability to differentiate self needs and emotions to others' and that infants help to alleviate perceived distress [4]. Some claimed that infants are biologically predisposed to empathy and altruism [42]. Others argued that immature self-other differentiation and goal contagion are responsible for the emergence of prosocial behavior [50].

This interest and hypotheses resulted in a number of theories, or models, proposed to explain why and how infants help others. For instance, the emotionsharing model suggests that an early form of empathy, in the form of emotional contagion, could be the primary behavioral motivation for infants to act prosocially [90, 20, 21]. In other words, infants help others to alleviate shared distress [4]. On the other hand, the goal-alignment model proposes that infants do not need to feel another person's distress, but in fact align their goal with others through a contagion process and are prompted to achieve them [49, 50] (see Chapter 2 for more details).

These theories are based on solid and thorough evidences, but limitations unable us to propose concrete and clear mechanisms for the emergence of prosocial helping behavior. The emotion sharing model is based on the ability to differentiate self from others, which means that being able to recognize someone else's mental state is required to help him [69]. However, some studies showed that helping emerges from 14 months of age, which is before the age at which infants acquire mental state differentiation [90]. The goal alignment model does not depend on mental state recognition, but fails to propose any motivational mechanism to explain why infants perform actions after aligning their goals with others. While these theories provide useful insights on prosocial behavior, the underlying mechanisms and the motivation to help are still largely debated and are not clearly defined.

1.1 Prosocial Behavior in Robots

Let us now come back to the question asked at the beginning of this introduction on why being social, but this time from the point of view of robotics. In the middle of the 20th century, robots became common tools in the factories, providing incredibly accurate and efficient help for hard labor and brain-less tasks. In the last two decades, robots started to become more and more sophisticated and allowed for even faster and more precise industrial production. At this point, robots also started to leave industry and enter into our household in the shape of vaguely autonomous toys like Aibo or tools like Roomba [31, 83]. However, we are still very far from having prosocial robots that help and assist humans such as those who inspired many robotics like Astro-boy, R2D2 or Wall-E.

To design artificial agents capable to interact with us socially and help efficiently, it appears necessary that they acquire human-like social capability. To make this possible, scientists have recently proposed computational models and algorithms capable to learn similarly with humans. An example is machine learning, which is a part of computer science in which agents are no longer pre-programmed, but learn models of the world by interacting with it. In particular, a subfield called cognitive developmental robotics (CDR) [2] aims to design robots that can learn and develop cognitive abilities based on human development. Progress in this area has allowed in the last years to develop more prosocial robots that are becoming a step closer to their science-fiction counterpart. Some examples of these new types of robots can learn to predict the effect of their actions on objects [86, 85] or acquire the ability to imitate others' intended goal [14].

In addition to create robots that are becoming more and more helpful and social, CDR may help psychologists and cognitive scientists to better understand human development. Indeed, studies in CDR are strongly inspired by infant development and propose bio-plausible robotics models for the emergence of cognitive and prosocial abilities. Therefore, we may be able to better understand humans by proposing infant-inspired model that replicate the different phases of our development.

1.2 Targets and Key Ideas of our Study

The development of prosocial helping behavior in infants is a fundamental step of their learning and integration into our highly social environment. Even if the evidences of its early emergence are plenty, only few theories have tried to explain infants' motivations to act prosocially and the mechanisms or the cognitive functions that foster its development. The main target of our work is to understand why and how infants become capable to perform prosocial helping behavior from early in their second year of life. To that end, we designed a cognitive developmental robotics system that replicates infants' development. Furthermore, we want to know when and how such system should generate prosocial helping behavior in order to be efficient during human-robot interaction and collaboration.

To achieve these research targets, we propose the three following issues to be solved in our study:

- 1. What is the motivation for infants to help other?
- 2. What is the role and effect of cognitive maturity on the development of prosocial helping behavior?
- 3. How to design an efficient autonomous helping robots inspired by infants?

Our key idea and hypothesis is that infants are motivated to help others by the minimization of the prediction error between predicted and observed goal. We assume that infants first learn their sensorimotor model through experience with the world. Due to immature cognitive abilities and the absence self-other differentiation, others' action goals are assimilated by the infant. Finally, infants would be prompted to perform an action in order to minimize the prediction error and achieve the assimilated goal if others failed it.

1.3 Overview of the thesis

In this dissertation, we proposed an infant-inspired computation model that replicates the emergence of prosocial helping behavior observed in infants. We used a model based on predictive learning (see Chapter 3) that learns from experience and acquires infant-like cognitive abilities. We then performed a series of experiments with robots to evaluate our ideas and hypothesis. Our first results demonstrated that our infant-inspired systems could acquire prosocial helping behavior based on the prediction error minimization (see Chapter 4). We also showed how the robot's performances changed in human-robot collaboration when varying the initiative variable during helping, and how these different behaviors are perceived by users (see Chapter 5). We finally summarized our contribution and introduced our future work the conclusion (see Chapter 6).

This thesis is composed of six chapters summarized below:

Chapter 1 - Introduction

The aim of this study is to understand and reproduce the emergence of prosocial helping behavior in infants. In this chapter, we showed why being able to help others is important for the development of infants and how it scaffolds their insertion into our society. We then suggested the advantages of replicating such behavior into robotic systems, both to create more efficient and social robots and to better understand human development. Next, we introduce the main targets of our work and our key ideas based on the minimization prediction error. Finally, we gave an overview of our contributions and the expected implications of our results.

• Chapter 2 - Related work

We first review the most significant psychological and developmental studies related to the emergence of helping behavior in infants. We then detail the achievement in the field of infant-inspired assistive and helping robotics agents. To finish this chapter, we show how other scientists tackled the problem of social human-robot collaboration.

• Chapter 3 - Predictive learning to understand infant development

Predictive learning describes how agents can learn and build a model of its environment through experience. By interacting repeatedly on the world in various context and with different behavior, the agent becomes able to learn a sensorimotor predictor that estimates the effect of its actions on the environment. In this chapter, we introduce predictive learning and present our general model to explain the underlying mechanism of prosocial helping behavior.

• Chapter 4 - Prediction Error Minimization for the Emergence of Helping Behavior

In this chapter, we suggest that infants are capable to predict others' goals based on self-experience and to perform actions if there are not achieved. In other words, infants help others to minimize prediction error. We evaluate our claim with two experiments showing how a robot can acquire helping behavior similar to what is observed in infants.

- Chapter 5 From the emergence of helping to efficient collaboration
 - Based on our model for the minimization of prediction error, we show how and when a robot should help during collaboration with human to be perceived efficient. We present three initiative conditions to trigger the robot's help: Robot-initiated reactive help triggers robot assistance when it detects high prediction error; robot-initiated proactive help makes the robot help whenever it can, even with low prediction error; humaninitiated help gives control of the robot action timing to the user; The result of our user study showed that proactive robots perform best and are perceived more efficient and natural that the robot reacting to high prediction error. However, reactive and prosocial robots elicited fewer face gazes than the human initiated ones, which may suggest that the interactions were less natural.
- Chapter 6 Conclusion

In this final chapter, the main contributions are presented. We suggest that prediction error minimization can account for the emergence of prosocial helping behavior based on our experimental results. We show that cognitive maturity, and in particular the ability to understand actions, directly affected the robot's helping performances. Next, we prove that users interact better with a robot if it helps proactively rather than when it reacts to high prediction error values. The chapter is concluded by the enumeration of our model's limitations and the presentation of our future work.

Chapter 2

Related Work

In the introduction, we showed the importance of infants' prosocial development for their integration into our society. We highlighted the remaining issues that prevent us from fully understanding why and how infants help others early in infancy. In this chapter, we review the literature related to these points. We start by looking at evidences and theories about when and why infants start to perform prosocial helping behavior. We then present studies that replicate infants' behavior and endow robots with the ability to perform helping and collaborative tasks.

In Fig. 2.1, the main findings related to the emergence of prosocial helping behavior and the development of related cognitive abilities are summarized. We show when infants become able to help in different scenario and when the corresponding required cognitive capabilities develop.

2.1 Prosocial Helping Behavior in Infants

In this section, we review the work orbiting around the emergence of prosocial helping behavior, especially when and why it develops. Additionally, we look at the development of action understanding, which is described by some as the minimal cognitive requite for helping.

Helping ability	Birth	6 months	12 months	18 months	24 months
Out-of-reach [90] and [24]					
Physical obstacle [89, 90]					
Wrong result* [89, 90]					
Wrong means** [89, 90]					
Cognitive capabilities	Birth	6 months	12 months	18 months	24 months
Action understanding [94]					
S/O assimilation [9, 28]					
S/O differentiation [9, 28, 69]					
N	o evidence in	favor Few ev	idences in favor	Strong evidences in favor	

Figure 2.1: Time-line for the emergence of prosocial helping behavior and the parallel development of cognitive capabilities. S/O: "self-other". *Wrong result: represents scenario in which an action has a different outcome from what can be expected. **Wrong means: represents scenario in which the action to achieve a predictable goal is not adapted. The gradient represent the stage of acquisition of the helping ability or cognitive capabilities. Light gray is equivalent to no evidence in infants. Dark gray represents small evidences or low maturity. Black represent strong evidence of high maturity of the ability or capability. This time-line is based on the work of Paulus *et al.* [69], Dunfiel *et al.* [24], Woodward [94], Warneken *et al.* [89, 90], Fein *et al.* [28] and Brownell *et al.* [9].

2.1.1 When does Prosocial Helping Start?

In the past decades, experimental studies have provided a large amount of evidences showing infants performing prosocial helping behavior at the end of the sensorimotor development stage (24-month-old according to Piaget [71]). Progressively, psychologists tested younger and younger infants to find out the earliest age at which they can help others (see Fig. 2.1 for a summary of our evidences).

At the end of their second year of life, children are known to help others in

fairly complex situations. Dunfield *et al.* for instance showed that at this age, infants had no difficulties helping adult that try to grab out for reach objects [24].

Even earlier, Warneken *et al.* [89, 88] showed that 18-month-old children could help others spontaneously in various situation. For instance, their studies showed children opening the door of a cabinet to allow an experimenter to put books inside, or correctly stacking boxes on a pile when an adult accidentally misplaced them. Over *et al.* [66] also showed that same age infants would more often and spontaneously help others in need when shown photographs evoking affiliation, showing the importance of the social factor.

Finally, Warneken and Tomassello recently performed a new set of experiments to show if even infants at 14 months of age were capable to be prosocial [90]. Despite the fact that cognitive maturation is very low at the beginning of their second year of life, infants surprisingly were still capable to help others. When the experimenter reached for a paper ball and failed to grab it, the toddler spontaneously gave him the targeted objects without hesitation.

2.1.2 Why do Infants Help Others?

While we now know that infants can help from 14-month-old, their cognitive maturity as such a young age is very limited and therefore the mechanism that triggers adult help cannot be applied. To solve this puzzle, we introduced two models theorizing about the emergence of infants' prosocial behavior in the previous chapter. In this section, we review more in details these theories and highlight their limitations and issues yet to be solved (see Paulus *et al.* [69] for more complete review). To further clarify our argumentation, the cognitive development evidences presented in this section are reported in Fig. 2.1, alongside the emergence of prosocial helping behavior.

Emotion-Sharing and Shared Distress

Emotion-sharing models suggest that an early form of empathy, in the form of emotional contagion, could be the primary behavioral motivation for infants to help [90, 20, 21]. Studies related to emotion-sharing models indeed posit that infants are primed to generate prosocial behavior in order to alleviate others' distress [95, 25]. This requires the ability to actually "feel" another person's emotional state, which is often called emotional contagion and represents "an automatic response resulting in a similar emotion being aroused in the observer as a direct result of perceiving the expressed emotion of another" (definition by Decety [21]). This ability is accepted as one of the lowest forms of empathy [20] and the cognitive requisite to prosocial helping behavior. Some scientists claim that infants experience an empathy-based feeling toward individuals in need of help, and that it serves as the primary motive for prosocial helping behavior [76, 44]. In practice, Warneken and Tomasello [88, 90] showed that infants helped others in achieving their goals and postulated that it substantiated the existence of a prosocial motivation in early infancy, closely related to empathy. It has been argued that empathetic concern is also independent from self-reflective abilities [19] and that empathy may be an innate capacity [43]. Studies have shown that very young children, before the age at which they develop self-other discrimination [9, 28], attempted to alleviate the distress of others and showed empathetic concern [95], and that even 12-month-old infants were concerned for others in distress and sometimes intervened by comforting them [25].

However, the cognitive abilities required by infants to feel empathetic concern for others, and thus to develop prosocial helping behavior on the basis of the alleviation of the shared distress remain very controversial. Some scientists have argued that self-other differentiation is required to acquire empathetic concern for others and to help, which implies that only infants that passed the self-recognition task would help others prosocially [6, 7, 69, 49]. Nevertheless, undeniable proofs of prosocial helping behavior have been shown during the first half of the second year of life, even though self-other differentiation and self-concept are immature as shown in Fig. 2.1. On the basis of these findings, we can assume that another source of motivation, more general in nature, may provide behavioral motivations to help others.

Goal-alignment and Action Understanding

Unlike emotion-sharing models, which are based on emotional contagion and empathetic concern, the goal-alignment models propose that more general mechanisms, based on the understanding of others' goals, serve as behavioral motivations for infants to help others. In other words, inferring or feeling others' mental or emotional state is not required for acting prosocially, but the ability to understand others' goals is a sufficient prerequisite. Indeed, Kenward [50] reported that 18-month-old infants could help spherical objects with no humanlike body to reach their goals, which may imply that empathy elicited by direct body matching is not applicable. They postulated that infants may be "primed by an unfulfilled goal", which supports the possibility of a general mechanism different from empathy and concern for others that could motivate infants to exhibit prosocial helping behavior.

In fact, infants very early in their infancy undergo a developmental process that allows them to perceive others' actions as goal-directed. Several studies have been carried out to reveal when and how infants start understanding goaldirected actions. A remarkable work conducted by Woodward [93] shows that infants (5, 6 and 9 months old) with goal-directed action experience react differently to actors reaching for and grasping objects. Later, Sommerville *et al.* [81] studied the impact of 3-month-old infants' action production on the perception of others' actions. The participants of the experiment were 3-month-old infants. Their findings reflected infants' ability to detect the goal structure of action after experiencing object-directed behavior, and to apply this knowledge when perceiving others' actions.

In addition, some experimental results tend to show that infants can assimilate others' action goal as their own through a goal contagion process. This mechanism is closely related to the mirror neuron systems, which are groups of neurons that fire similarly when performing or observing a same goal-directed action [74]. Therefore, infants would be affected by this contagion process and perform actions to achieve the shared goals when others fail to do so, leading to what would seem to be prosocial helping behavior.

To explain why infants help others, the goal alignment models suggest that young children are prompted by unachieved goal. However, even if it has been shown that infants can assimilate others' goals, the behavioral motivation explaining why infants would perform actions in response is not known.

2.2 Helping Robots and Human-Robot Collaboration

Given our emphasis on the importance of understanding infant prosocial helping behavior using a computational approach in Chapter 1, previous work on assistive robots and their associated methodology are highly relevant. Akin of our interest, Chung *et al.* [14] proposed a developmental robotic system based on statistical inference of others' action goal to generate the imitation of intended goal. Using a Bayesian network, their system learns probabilistic models of actions through experience with its environment. Then it uses this learned models to infer humans' action goals, and finally achieve the predicted goal. While their work focuses more on the imitation of intended goal than helping, their approach is highly consistent with ours and provides useful hint and direction for our study.

Other works by Verma *et al.* [87] and Grimes *et al.* [35] also used probabilistic inference of others' action and goal similar to ours. They showed that probabilistic model could be used to generate goal-based and action-based imitation. Similarly, Gray *et al.* [34] presented an architecture for action parsing and goal inference using self as simulator. Based on the well-known simulation theory, their model uses the same cognitive structure to generate behavior and simulate others' mental states when performing actions. The robot presented in their work is thus capable to anticipate others needs and to offer meaningful assistance.

Liu *et al.* [52] proposed a human intention recognition algorithm to allow a robot to collaborate with a human. They used a finite state machine in order to allow their system to estimate one's goal and to guide the robot's actions. This method allowed for more efficient interaction, but more importantly cooperation robust to uncertainty and noise. Similarly, Najmaei and Kermani's prediction-based reactive control model for collaboration [62] used neural network to predict an upcoming event based on the current sensory observation. Based on this prediction, the robot is capable to react and minimize the potential danger of others' actions.

In recent years, collaborative robots designed to work side-by-side with humans have gained momentum in real-world settings. This has fueled a large body of research on human-robot collaboration, tackling for instance the problem of *task planning* for joint human-robot tasks. Shah generates a robot action plan so as to minimize human-idle time [78]. Another vein of research focuses on lowlevel motion planning for the robot within a collaborative context [54, 53, 79], with an eye towards improving team fluency and the user's sense of safety.

Researchers have studied other low-level behaviors, besides robot motion, that impact collaboration and enable coordination of actions during task execution [59]. For example, St. Clair and Mataric demonstrated that robot verbal feedback improves team performance [15]. Awais proposed mechanisms to mitigate breakdowns in the joint tasks [3]. Others focus on the coordination of micro-interactions that occur during collaboration, such as object hand-overs, using gaze [58] or adapting timing of motions to the human's state [46]. Chao and Thomaz developed mechanisms to coordinate sharing of common resources during collaboration, such as the speaking floor or part of the workspace that both the human and the robot need to access [13].

Besides generation of robot behaviors, another key problem in human-robot collaboration is perception of the human. Preliminary work by Hoffman and Breazeal suggests that anticipatory perceptual simulation improves efficiency and fluency in teamwork [41, 40]. With the help of new sensing and human tracking technologies, many others followed with models of action or motion anticipation in the context of human-robot collaboration [63, 3, 48, 37].

In the context of human-robot collaboration, one study by Gombolay *et al.* is particularly relevant. They investigate decision-making authority in the planning process and find that people are willing to give control to the robot for the efficiency benefits [32]. While our assumptions are consistent with theirs, our study differs in its focus on authority over assistance timing *during task execution*, as opposed to authority over assistance allocation *during task planning*. Groten *et al.* looked at shared decision making in the context of haptic collaborations [36]. Cakmak *et al.* investigated initiative in robot question asking [10]. In addition, the large body of work on mixed-initiative control in the context of robot teleoperation [29] has some relevance to our work.

Finally, several of the challenges related to fluent and efficient human-robot collaborations have already been tackled. Some of the main research threads have investigated ways to compute robot action plans that improve joint task performance while reducing the workload on the human [38], tracking and anticipating human motion to enable execution of such task plans [70, 64], and designing robot behaviors to improve team effectiveness and fluency [23, 13]. Others have given guidelines on the behavior the robot should perform in order to be perceived as social. For instance, Li *et al.* [51] suggested that a social robot should be able to recognize the presence of human, engage in physical acknowledgment, use physical motions, express/perceive emotions and engage in some sort of communication.

Chapter 3

A theory of Predictive Learning to Understand Infant Development

To explain the emergence of prosocial helping behavior in infants, we decided to emulate their development using a computational model. Based on the related work, we identified the main components required for infant to start helping others, such as goal understanding or the assimilation of others' action goal. To reproduce these cognitive and motor capabilities into a robotic system, we decided to use a model based on predictive learning.

In the next sections of this chapter, we first introduce the idea of predictive learning and how it could allow artificial agents to acquire causal knowledge of the world (*e.g.*, the effect of an action on the environment). We then present the general model used throughout our experiments that reproduce the cognitive and motor capabilities of infants that are involved in the emergence of prosocial helping behavior.

3.1 Predictive Learning

In industry, robotics methods are mostly based on a set of pre-programmed models and heuristics describing the expected effects of a robot on its environment



Figure 3.1: Basic architecture for sensorimotor predictive learning proposed in [61] that consists of two parts: a sensorimotor system (the lower box) and a predictor (the upper box). The sensorimotor module performs actions under a certain sensory states and receives a new state in return. The predictor module learns to estimate sensory and motor signals, through interaction with the world. The target of the system is then to minimize the difference between the observed new state and its predicted value.

and vice-versa. In step planning, a robot would need to compute all the possible solutions and select the shortest path based on a human-designed cost function. While these approaches are efficient for repetitive tasks in known situations, they do not scale to complex and changing environment such as human-robot joint task collaboration.

To solve these issues, machine learning allows agents to construct models of the world by interacting with it and to make prediction about future events. Systems using machine learning are trained through different processes, such as trial and error or supervised learning, during which value functions are optimized. This enables agents to make decisions based on new data with a minimum of pre-programmed behavior.

Predictive learning describes such a technique in which an agent can build a model of its environment through experience. By interacting repeatedly on the world in various contexts and with different behaviors, the agent learns to predict action-state sequences using a sensorimotor model. Predictive learning allows an agent to learn to make decision by using the knowledge of the effects its actions to create planning operator. This approach is popular in robotics or



Figure 3.2: Training: A robot trains to associate an observed current state (stain), its action (wipe) and the future state (no stain). After training: this training allows the robot to predict the effect of its actions on the current state.

statistics where it is used for decision making (*e.g.*, gaze shift) [17], classification problems [30] or in general to identify the interaction between a set of variables.

Fig. 3.1 describes a basic architecture for predictive learning. The architecture consists of two main elements that describe a sensorimotor system and the predictor. The sensorimotor module is the part of the system that interacts with its environment, *i.e.*, the body. It performs an action $a_i(t) \in \mathbf{a}(t)$ under a certain sensory state $s_i(t) \in \mathbf{s}(t)$, and perceives a new state $s_i(t+1)$ as a consequence. The predictor module simulates the sensorimotor system and estimates the next state of the environment and future actions. It learns to estimate sensory and motor signals, $\hat{s}_i(t+1)$ and $\hat{a}_i(t+1)$ through interaction with the world. The target of the system is to minimize the prediction error noted $e_i(t+1)$, which is the difference between the observed new state and its predicted value.

To further illustrate our predictive learning model, let us imagine a robot learning to clean a room (see Fig. 3.2). The robot is trained, either through demonstration or guided learning, to associate observed states (stain $s_1(t)$), possible future actions (wipe $a_1(t+1)$) and future state (no stain $s_2(t+1)$). As the learning progresses, the sensorimotor model (cf. predictor) is optimized and the prediction error is minimized. This process allows the robot to then predict the effect of actions on the current state, for instance, predicting that wiping



Figure 3.3: Overview for the proposed model. The perception module: allows the model to perceive and recognizes the scene and to extract any meaningful information. The predictive learning module: learns sensorimotor model through learning and predicts future actions and states. The prediction error minimization module: estimates prediction error and performs actions that minimize it.

a stain would remove it $(s_1(t) \rightarrow a_1(t+1) \rightarrow s_2(t+1))$. As the system learn sensorimotor association though learning, the robot can learn various behaviors if training with different experience.

3.2 Overview of Our General Model

In this study, we want to create agents capable to acquire infant-like prosocial helping behavior using a predictive learning architecture. The general model structure used throughout this dissertation is described in Fig. 3.3 and is adapted the predictive learning model presented in Fig. 3.1. The predictor remains identical from the general architecture of predictive learning described in Section 3.1. The sensorimotor system is separated into two part that respectively perceive states and actions, and act upon the environment. These parts are distributed into 3 modules presented below:

1. The perception module: recognizes and categorizes the different elements of the scene, such as observed objects or others actions. All the perceived elements are factorized into the sensorimotor vectors s(t) and a(t).

- 2. The predictive learning module: purpose is twofold. First it memorizes causal relationship between action-state sequences by building a sensorimotor model through the interaction with its environment. Second, it generates prediction of future state (s(t+1)) or action (a(t+1)) based on the observation made by the perception module (c.f., s(t) and a(t)). This correspond to the predictor part of the predictive learning architecture in Fig. 3.1.
- 3. Prediction error minimization module: estimates and minimizes the prediction error based on the perceived and predicted environmental state. If a predicted state is not achieved, the module will minimize the error by generating a predicted action.

Together, these modules respectively perform an action predicted under a certain sensorimotor state by the predictive learning module and receives a new state in return.

This model is based on evidences related to the goal alignment model for the emergence of prosocial helping behavior introduced in Chapter 2. According to this theory, the abilities to build a sensorimotor model of others or to recognize their intention are not required to help. In other words, this system does not need to differentiate self from others. In fact, prosocial helping behavior emerges as a by-product of the minimization of prediction error estimated over the goals assimilated by the system.
Chapter 4

Emergence of Prosocial Helping Behavior in Robots

In the introduction and related work chapters, we reviewed evidences about the development of prosocial helping abilities in infants. Two theories for the emergence of helping were presented: the goal-alignment and the emotion sharing models. We found out that the emotion sharing model required the abilities to differentiate self from others and to understand others' mental state, which develop later than onset of prosocial helping behavior. On the other hand, the goal-alignment model did not require such ability, and seemed to be primed by a more general mechanism. This ability is known as action understanding and develops in infants around the age of 5 to 9-month-old [72, 5, 75, 93] and allow them to comprehend action and predict its goal. However, the motivation for infants to perform prosocial helping behavior based on the prediction of action's goal was not clearly defined.

In this chapter we attempt to answer the question of the motivation for the emergence of prosocial helping behavior based on the ability to understand action goals. We begin by proposing a hypothesis based on the prediction error minimization as motivation to help. We propose a system based on predictive learning introduced in Chapter 3 to emulate how infant might acquire the ability to help. Next, we present two experiments to test our hypothesis along with our model, and study the effects of cognitive maturity on the emergence of prosocial helping behavior. Additionally, we evaluate our claim by endowing a humanoid robot with the ability to help other. The chapter is then concluded by an analysis of our results and a summary of our contributions.

4.1 Hypothesis for the Emergence of Helping Behavior

To shed light on a possible mechanism for the emergence of prosocial helping behavior in infants, we hypothesize that helping emerges as a by-product of the minimization prediction error (hereafter called PE) estimated from others' action goals. We suggest that infants start by learning their sensorimotor model though interaction with their environment. This training allows them to predict the effect of their actions on the environment, namely their action goals. Due to immature self-other differentiation in early infancy, they may assimilate or take over others' actions as their own and predict their goals [50]. Prediction error then arises when the predicted action goals and the observation mismatch, regardless of who is executing the action. The prediction error estimated during this process then triggers infants' actions to minimize it, which results in prosocial helping behavior. An important point of this hypothesis is that infants do not have any explicit intention to help others by design. Instead, goal alignment suggests that infants always try to achieve predicted action plans, leading to the emergence of prosocial helping behavior as a by-product.

4.2 Model for the Minimization of Prediction Error

Our computational model is based on behavioral evidences of infant development and attempts to understand and reproduce the mechanisms of the emergence of prosocial helping behavior in 14-month-old infants. We then assume that our system has the cognitive and motor capabilities of a 14-month-old infant:

- 1. Assumption 1: It is capable to learn to predict sensorimotor sequence through the interaction with their environment.
- 2. Assumption 2: It can assimilate others' action. This cognitive capability is related to the finding on mirror neuron systems, which are a group of



Figure 4.1: Model for the minimization of PE. The three modules are used to recognize action primitives and the associated objects (conditions), to predict the next action primitives, and to estimate PE and to generate the primitive to minimize PE.

neurons that fire both when performing and observing a same goal directed action [74].

3. Assumption 3: It is capable to perform object/goal directed actions.

These cognitive and motor capabilities represent the requirements for our model and are known to be available in infants before the onset of prosocial helping behavior at 14 months of age.

Fig. 4.1 shows the overview of our model for the minimization of PE, which consists of three interdependent modules: the perception, predictive and minimization of PE modules. This model is trained by performing various action with its environment and tested during an observation phase in which the robot observes others' uncompleted actions. During training, the predictive module trains a directed graph, hereafter called an action graph, which represents the robot's sensorimotor model (assumption 1). During observation, the minimization of PE module estimates the prediction error and executes object directed actions to minimize it if needed (assumption 3 and main hypothesis). The perception module recognizes objects using color filters and others' action based on own sensorimotor model during both training and testing (assumption 2). More details about the three modules are given in the following sections.

4.2.1 Perception Module

This module recognizes action primitives (noted A in the next sections) and objects contained in the scene (noted C in the next sections). Action primitives are simple motions that the robot can execute like reaching for a ball, grasping a mug, covering a marker, etc. Hereafter, a sequence of action primitives is called "action". For instance, a "pushing" action contains the action primitives "reach for" an object and "move" it. Objects, hereafter called "conditions", are elements of the scene that can be interacted with like a ball, a mug, etc.

4.2.2 Predictive Module

The predictive module estimates the goal of a motion as the future action primitive based on the current observation. This process is presented below in the action prediction paragraph. The prediction is performed using the action graph, which memorizes the robot's past experience and represents its sensorimotor model. In the following parts, we describe how the action graph is generated and how it is used for the action prediction.

Statistical Action Graph

The action graph uses the robot's experience with environment to make prediction based on ongoing action. An action graph (G) is made of two kind of nodes, representing the system's sensorimotor representation, namely the previously experienced action primitives and their associated conditions. The two types of node are:

- The action nodes **A** that represent action primitives performed by the robot. The number of times an action node has been performed by the robot is noted NB_A .
- The condition nodes **C** that represent the conditions for an action primitive to be executed, namely the object the robot interacted with while performing action primitives.

Action nodes are connected by directed edges $\mathbf{E}_{\mathbf{A}}$ that encode the number of times a transition between two action nodes was experienced. The number of times a transition $\mathbf{E}_{\mathbf{A}}$ has been activated is noted $NB_{A_i \to A_j}$, where A_i and A_j are two different action nodes. The conditional relation of condition nodes to



Figure 4.2: Example of the different steps in the creation of an action graph after executing two different actions. Action I: "*Reach for a Ball, Grasp* the *Ball*, and then *Put* the *Ball* in an *Opened Box*"; Action II: "*Reach for a Ball, Grasp* the *Ball, Open* a *Closed Box*, and then *Put* the *Ball* in the *Opened Box*". Action I was experienced first, then Action II and finally Action I again. The small numerals inside the action nodes represent the number of times the action primitive corresponding to the node was successfully executed, namely NB_A . The small numerals by the directed edge represent the number of times the connected child and parent nodes were performed successively, noted $NB_{A_i \to A_j}$.

action nodes is represented by another type of edges noted $\mathbf{E}_{\mathbf{C}}$. The graph is then represented by:

$$\mathbf{G} = (\mathbf{A}, \mathbf{C}, \mathbf{E}_{\mathbf{A}}, \mathbf{E}_{\mathbf{C}}), \tag{4.1}$$

where all nodes are Boolean variables and can take a value of 1 (active) or 0 (inactive).

Fig. 4.2 shows an example of how an action graph is generated while experiencing three actions (Action I twice and Action II once). Action I: "Reach for a Ball, Grasp the Ball, and then Put the Ball in an Opened Box"; Action II: "Reach for a Ball, Grasp the Ball, Open a Closed Box, and then Put the Ball in the Opened Box". Action I was experienced first, then Action II and finally Action I again. In this example, $\mathbf{A} = \{A_1, A_2, A_3, A_4\}$ and $\mathbf{C} = \{C_1, C_2, C_3\}$. A_2 is the child node of A_1 , while A_1 is the parent node of A_2 . The numerals inside the action nodes in Fig. 4.2 represent the number of times the primitives were successfully executed, noted NB_A . For instance, *Reach for* was executed during all actions, $NB_{A_1} = 3$, while *Open* was executed once during Action II, $NB_{A_3} = 1$. The numerals by the directed edge in Fig. 4.2 represent the number of times the connected child and parent nodes were performed successively, noted for instance $NB_{A_1 \to A_2}$. For instance, *Open* was performed 1 time after executing *Grasp*, thus $NB_{A_2 \to A_3} = 1$.

The action node corresponding to the currently recognized action primitive is denoted as $A_{i(n)} \in \mathbf{A}$, and the condition nodes representing its conditions are contained in the subset $\mathbf{C}_{A_i(n)} \subset \mathbf{C}$. *n* represents the current discrete time step. In Fig. 4.2, for instance, the action primitive "Put a Ball inside an Opened Box" is described by the action node "Put" $(A_{4(n)})$ and the condition nodes "Ball" and "Opened Box", which are contained in $\mathbf{C}_{A4(n)} = \{C_1, C_3\}$.

In practice, the action graph is constructed when the system executes actions with its environment. This process is ruled by the following mechanisms:

- (i) The system performs a primitive from its action repertoire involving objects in the scene.
- (*ii*) For the executed primitive, the corresponding action node A_i and condition node(s) \mathbf{C}_{A_i} are added to the action graph. The condition nodes are connected to the action node by directed edge E_{CA_i} . If an action primitive is performed several times with different objects, multiple instances of the action nodes are created and connected to the corresponding subset of conditions. The delay between the onset and the completion of the action primitive is measured as T_{A_i} . The value NB_{A_i} , representing the number of times this primitive has been executed, is initialized at 1.
- (*iii*) If the node corresponding to the performed primitive with the same subset of conditions is already contained in the graph (i.e., if the system has already experienced the primitive before), the delay T_{A_i} is averaged and the value NB_{A_i} is incremented.
- (iv) If two action primitives are performed consecutively within a delay shorter than a value T_{max} (fixed at five seconds in the current implementation), the corresponding action nodes $A_{h(n-1)}$ and $A_{i(n)}$ are connected by a directed edge E_{Ah-i} . The value $NB_{A_i \to A_j}$, representing the number of times A_i was executed after A_i , is initialized at 1 and incremented each time the

same transition occurs. If the two action primitives are performed consecutively with a delay higher than T_{max} , the newly performed primitive is considered as part of another action. Therefore, the two action node are not connected by any edge.

By performing these learning operations multiple times with different objects and for all action primitives in the system's repertoire, the system becomes able to perform action prediction, which is explained in more details below.

Action Prediction - Goal Understanding

Based on the experience represented in the action graph, the system calculates the probability of observing a primitive $A_{j(n+1)}$ when a node $A_{i(n)}$ is activated. $A_{i(n)}$ can either be activated when the system is executing the action primitive or when it is observing another individual performing the same primitive. This probability is represented by the conditional probabilities $P(A_{j(n+1)} = 1 | A_{i(n)})$, which is calculated as follows:

$$P(A_{j(n+1)} = 1 | A_{i(n)}) = \frac{NB_{A_i \to A_j}}{NB_{A_i}},$$
(4.2)

where NB_{A_i} represents the number of times the primitives A_i was previously executed by the system, and $NB_{A_i \to A_j}$ represents the number of times A_j was performed after A_i . The sum of the probabilities for a given current state $A_{i(n)}$ respects:

$$\sum_{j} P(A_{j(n+1)} = 1 | A_{i(n)}) = 1;$$
(4.3)

The system then tries to find the most likely future action node $\hat{A}_{(n+1)}$. To that end, the system detects the node $A_{j(n+1)}$ with the highest probability $P(A_{j(n+1)} = 1 | A_{i(n)})$ and that can be activated. Indeed, if the value of at least one of its conditions $C_k \in \mathbf{C}_{A_j(n+1)}$ is 0, the corresponding primitive $A_{j(n+1)}$ cannot be activated. Therefore, the future primitive $\hat{A}_{(n+1)}$ is:

$$\hat{A}_{(n+1)} = \underset{A_{j(n+1)}}{\arg \max(\min(\mathbf{C}_{A_{j}(n+1)}))} (4.4)$$
$$(P(A_{j(n+1)} = 1 | A_{i(n)})); \forall j.$$

If two or more nodes have the same conditional probability and if all their corresponding condition nodes are activated, $\hat{A}_{(n+1)}$ is randomly selected among these nodes. If $\hat{A}_{(n+1)} = 0$, the system remains idle and no future action is selected.

4.2.3 Minimization of Prediction Error Module

The minimization of PE module estimates PE signal from observation and prediction and generates actions to minimize it. It is separated in two main parts, the estimation of prediction error and the action execution.

Estimation of Prediction Error

To estimate PE when primitive $\hat{A}_{(n+1)}$ is predicted, two main components are taken into account:

- (i) The conditional probability of $\hat{A}_{(n+1)}$, which is hereafter noted P_{Max} .
- (*ii*) The difference between the delay T_{A_i} and the elapsed time (called t_e) since the current node $A_{i(n)}$ was activated.

PE is then measured as P_{Max} discounted by a time dependent function as follows:

$$PE = P_{Max} \cdot \beta \cdot (1 - e^{(T_{A_i} - t_e)}); \tag{4.5}$$

where $\beta = 0$ when $T_{A_i} \ge t_e$; else $\beta = 1$. β fixes PE = 0 when the elapsed time is shorter than the average delay T_{A_i} of the observed action $A_{i(n)}$. Therefore, PE starts to increase only as t_e becomes greater than $T_{A_i(n)}$. An example of PE estimation is depicted in Fig. 4.3 where a primitive is observed but not completed, leading to an increase of PE. PE is defined such as its value increases if a prediction is not achieved within a certain amount of time. This definition is based on psychological and neuroscience observations, and is simplified to fit our experimental conditions.

Action Execution

We hypothesize that observing others' failure in action execution would lead to the robot performing the predicted action primitive. If PE is greater than a threshold (empirically fixed at 60% of P_{Max} in our current experiments), the PE minimization module executes the predicted primitive $\hat{A}_{(n+1)}$ as an output of the system (see Fig. 4.3). For example, when the system observes another individual trying and subsequently failing to achieve an action (e.g., opening a Closed Closet), the minimization of PE will lead to the robot executing the predicted action (e.g., the robot opening the Closed Closet). From the point of



Figure 4.3: An example of PE estimation. The system observes a primitive $A_{i(n)}$ and can predict the next action primitive $\hat{A}_{(n+1)}$. When the elapsed time becomes greater than T_{A_i} , PE starts increasing. Finally, when PE passes the threshold, the robot performs the predicted primitive to minimize PE.

view of the other individual, this process looks as if the robot helped the person even though it does not have such an intention.

4.3 Experiment 1- Setup

In this experiment, we tested our hypothesis and studied the effect of cognitive maturity on prosocial helping behavior. We used a fully simulated environment to remove any noise coming from the perception module and focus on studying the relevance of the predictive and PE minimization modules. The perception module was replaced by a symbolic representation of actions and conditions instead.

The experiment was separated in two phases: the training phase, during which the robot trained its action graph (or sensorimotor representation) by performing series of actions; The observation phase (or testing), during which the robot observed others' actions and tried to minimize PE. The actions used during our experiment were inspired by the experiments performed by Tomassello and Warneken [90, 89], in which they showed that infants could help others trying to reach out-of-reach objects or overcome obstacles (e.g. out-of-reached cloth pins, and closed cabinet doors).

During the training phase, the robot was trained with several actions in a randomized order and built an action graph, as presented in Section 4.2.2. During the observation phase, seven non-accomplished actions were presented to the system, during which PE was estimated. In order to study the effects cognitive maturity on the ability to estimate PE and help, the amount of experience given to the system during the training was varied, from the execution of one action to that of all six possible actions. This incremental learning had the effect of modulating the complexity of the sensorimotor model, and consequently the ability to predict action goals and emulated the cognitive maturity.

4.3.1 Robots Actions

For this experiment, the system could experience six actions (Act1 to Act6) that are combinations of eight different objects (Ball, Mug, Car, Switch, Opened Closet, Closed Closet, Opened Box, and Closed Box) and six different action primitives (Reach for, Grasp, Open, Put, Move, and Flip). The actions the system experienced are described in Table 4.1. Act1 to Act4 contain "Reach for" a Ball **and** a Mug because both objects are present in the environment. As we assume that our system cannot identify which of the "Ball" or the "Mug" is the target due to perception ambiguity caused by the objects being too close to each other, both objects are conditions for the "Reach for" primitive.

Actions					
Act1	Reach for a Ball and a Mug, Grasp the Ball, Open a				
	Closed Box, Put the Ball in the Opened Box				
Act2	2 Reach for a Ball and a Mug, Grasp the Ball, Put the Ball				
	in an Opened Box				
Act3	Reach for a Ball and a Mug, Grasp the Mug, Open the Closed				
	Closet, Put the Mug in the Opened Closet				
Act4	Reach for a Ball and a Mug, Grasp the Mug, Put the Mug				
	in an Opened Closet				
Act5	Reach for a Car, Move the Car				
Act6	Reach for a Switch, Flip the Switch				

Table 4.1: Experiment 1: Six actions the system experienced.

4.3.2 Training and Observation Phase

During the training phase, the robot experienced up to six different actions. All action primitives were correctly performed. The actions were designed so that the number of child and parents for the different action nodes varies. In some



Figure 4.4: Experiment 1: Example of action graph for all possible actions Act1 to Act6 executed once. The red nodes denote conditions, and the black nodes represent actions. The numbers inside the action nodes denote the number of times the primitives were observed.

cases, action nodes only had one child and parent node; For instance in Act5, "Reach for the Car" could only be followed by "Move the Car". In contract, some action nodes had several parent or child nodes; For instance in Act1 to Act4, "Reach for the Ball and a Mug" could be followed by "Grasp the Ball" or "Grasp the Mug". Fig. 4.4 shows an example of an action graph built after performing all the actions presented in Table 4.1.

During the observation phase, other individuals performed seven uncompleted actions (F1 to F7) listed in Table 4.2. The action primitives and objects used in actions F1 to F7 were the same as those used during the training. These actions could be uncompleted for two reasons:

• Out-of-reach: Other individuals may fail to reach for an object if it is too

ID	Performed primitives	expected primitive	
F1	Reach for Ball	Grasp Ball	
F2	Reach for Mug	Grasp Mug	
F3	Reach for Ball and Mug	Grasp Ball	
F4	Reach for Mug, Grasp Mug	Open Closed Closet	
F5	Reach for Ball, Grasp Ball	Open Closed Box	
F6	Reach for Car	Move Car	
$\mathbf{F7}$	Reach for Switch	Flip Switch	

Table 4.2: Experiment 1: List of others' failed actions. Expected primitive are not achieved by others.

far from them. In this case, the next primitives predicted after "Reach for" (e.g., "Grasp") cannot be observed (activated).

• Physical obstacle: Other individuals may fail to use or interact with an object because of a physical constraint (e.g., cannot open a box if the hands are occupied with balls).

When observing others' actions, our system tried to predict the most likely next action primitives. Because some action nodes had several child nodes, the prediction could be ambiguous. For instance, if F3 was observed and if our system had previously experienced Act1 and Act3, both the primitive "Grasp the Ball" and "Grasp the Mug" could be predicted. This is later called prediction ambiguity.

4.4 Experiment 1- Results

We trained our system for six different conditions, each with a different number of actions performed during the training. The number of actions performed was incremented from one in the first condition, to six in the last. During the training the order of action execution was randomized. We then tested our system for seven different tasks in which another individual performed uncompleted actions.



Figure 4.5: Experiment 1: Column plot representing our system's *acted*, *helped*, and *failed* performances. The error bars represent the standard deviations.

4.4.1 Evaluation of Helping Performances

For each trial, we observed whether the system could successfully produce an action to minimize PE (hereafter denoted as *acted*). If the *acted* primitive could help others in achieving their goals it was denoted as *helped*. If it did not help achieving the goal, the action was categorized as *failed*. In other words, a *failed* primitive is a behavior that successfully minimized PE estimated by the system, but was not helpful from the other's point of view.

Fig. 4.5 shows the *acted*, *helped*, and *failed* performances of our system as a function of the number of actions experienced. The sum of *helped* and *failed* values represent 100% of the *acted* value (*acted* = *helped* + *failed*). The results show that the performance of our system improved as the number of actions experienced increases. The *helped* value got higher than the chance level of (16.67%) after experiencing three different actions. Some actions could be generalized better than others as shown by the *acted* values and the standard deviations. Indeed, if only Act6 (see Table 4.1) is experienced, only actions involving the Switch can be recognized, but if only Act1 is experienced, our system can make predictions for all action involving the Mug or the Ball. As the *helped* value increased, the *failed* decreased proportionally. However, the proportion of primitive that unsuccessfully helped others remained high for two main reasons presented in the following paragraphs.

Effects of Recognition and Prediction Ambiguity

Recognition ambiguity: If multiple objects are located close to each other and are associated with a same action primitive, which is currently activated, the system cannot identify the target object of the ongoing action.

Prediction ambiguity: If multiple action primitives are experienced after a same action primitive (i.e., single parent node connected to multiple child nodes), the system cannot predict accurately which action primitive should be executed next.

Effects of Perspective Difference

The action primitive performed by the robot cannot always help others in accomplishing their intended behavior due to the perspective difference between the robot and others. For instance, when others intend to "Grasp a Mug", the robot performs the action primitive "Grasp the Mug" after observing "Reach for the Mug" to minimize PE. This resulted in the Mug in the robot's hand, but not in others' hand.

4.5 Experiment 1 - Discussion

Our first experiment showed that the minimization of PE could explain the motivation for infants to help others achieve unsuccessful actions. Furthermore, we demonstrated that cognitive maturity greatly influenced the system ability to estimate PE and therefore directly affect helping performances. These results are consistent with evidences presented in Chapter 1.

In most cases, executing the predicted future action primitives could help others achieve their actions. However, it happened that even though PE was minimized by the robot, its actions failed to help others. As mentioned in Section 4.4, we observed different scenarios that could explain why our system failed to help others: recognition and prediction ambiguity, and perspective difference. The ambiguity errors can be explained by the lack of training and



Figure 4.6: Experiment 2: Setting. The blue Car is shown on the right of the robot, and the red Marker is presented on the left of the robot.

generalization of our model. On the other hand, the perspective taking issues are more challenging to address and bring additional questions. Indeed, infants seem to rarely get affected by any sort of perspective between them and the individual they are helping [88]. The mechanism allowing infants to cope with perspective differences is not clear and several possible solutions to this problem will be presented in the discussion.

4.6 Experiment 2 - Setup

The second experiment was designed to demonstrate whether our system could exhibit similar prosocial helping behavior in a more complex and noisy environment. For this experiment, we used a humanoid iCub robot (see Fig. 4.6). We used 19 of the 53 degrees of freedom: 7 in each arm and 5 in the head. The head, the right arm, and the left arm were used during our experiment.

The robotic system is presented in Fig. 4.6. The robot was placed 0.1 m away from a 1-m-high table on which a black mattress was placed. Two objects (a toy Car and a Marker) were positioned on the black mat at a reachable distance from the robot's arms. The object positioned on the left was manipulated by the left arm, and conversely for the object on the right. The objects had spe-

cific affordances: the Car was move-able but not hide-able; the Marker was not move-able but hide-able. The robot was able to perform four action primitives: "reach from the side", "reach straight", "move", and "hide"; the action primitives were executed using the YARP Cartesian interface [68]. The primitives were combined into two actions: push ("reach from the side" and "move") and cover ("reach straight" and "hide"). Below, we detail the experiment specific definition of the perception module and the action graph.

4.6.1 System Implementation

In this experiment, we introduced and tested the perception module using camera images, which was not implemented during the first experiment. Additionally, the predictive module was slightly modified to prevent some experiment related artifacts. These two changes are presented in the following paragraphs.

Perception Module Modification

The perception uses the RGB camera $(640 \times 480 \text{ pixels})$ placed in the robot left eye to detect the objects and action primitives. Objects are detected by combining pixels with similar color and (x, y) position (see Fig. 4.7 (a)-(b)). We use a set of predefined colors (e.g., blue or red) for the detection. The objects are then tracked based on their position and average hue (see Fig. 4.7 (c)-(d)) unless they are not visible for longer than two seconds. The objects are categorized into three states depending on their position history:

- (i) Stationary: the object is stable in position;
- (ii) Moving: the distance traveled by the object during the ongoing action (no time limit) reached 50 pixels;
- (*iii*) Occluded: the object is not detected for more than 500ms and less than 2s.

Action primitives are recognized by looking at the relative position of the hand to the objects. The x and y coordinates of the hand in the image are detected using the predefined skin color like for the object detection (see Fig. 4.7 (e)). Our system can recognize two types of reaching, either "reaching for the side" if the hand is positioned on the side of the object in the x axis or "reaching straight" if the hand is aligned with the object.



Figure 4.7: Visual processing. (a). Raw image. (b): Extraction of all colors. (c): Color extraction without skin color. (d): Object tracking. (e): Hand recognition.

Predictive Module Modification

For this experiment, the actions performed by the robot could result in no effect on the targeted objects due to their specific affordances. To cope with this issue, series of action primitives performed by the robot and the corresponding condition nodes (objects) are memorized in the graph if and only if the performed action modifies the state of at least one object in the scene. For instance, "reach from the side" for the Car and "move" the Car would lead to the Car's movement, and therefore the action is memorized. In contrast, "reach from the side" for the Marker and "move" the Marker would have no effect on the Marker's state, and thus the action is not memorized.

4.6.2 Training and Observation Phase

The experiment was divided into 10 trials with five subjects, each composed of two phases: a training phase and an observation phase. The subjects were chosen from our laboratory and were not familiar with our study.

During the training phase, the robot interacted with the objects presented in front of it. The robot was instructed to either push or cover the objects on the left or right side on the table. During each trial, the robot performed all the four possible actions twice in a random order.

During the observation phase, the robot was placed in front of a subject and observed his behavior. When the subject performed an action primitive with an object, the node corresponding to the primitive in the action graph was activated. The action prediction module then predicted the next primitive to be executed. If the subject failed in achieving the predicted action primitive within a certain time, PE started to increase. If PE exceeded a fixed threshold, a trigger signal was sent to the minimization of PE module, which executed the predicted action primitive in order to minimize PE.

4.7 Experiment 2 - Results

The results gathered during the training and the observation phases for the 10 trials are presented below.

4.7.1 Evaluation of the Training Phase

The Car and the Marker were randomly placed either on the left or the right side of the mat on the table. During each trial, the robot performed all the action presented in Table 4.3 twice in a random order. Fig. 4.8 shows the robot performing 2 actions learned by our system: (a): "reach from side for" the Car and "move" the Car and (b): "reach straight for" the Marker and "hide" the Marker. When moving the Car, the state of the Car switched from "stationary" to moving, and when hiding the Marker, the Marker's state switched from "stationary" to "occluded". The action graph after performing all four actions is presented in Fig. 4.9.

Action primitives	Objects	Status
"Reach from side for" & "move"	Car	Memorized
"Reach from side for" & "move"	Marker	Not memorized
"Reach straight for" & "hide"	Marker	Memorized
"Reach straight for" & "hide"	Car	Not memorized

Table 4.3: Experiment 2: List of action primitives, objects, and status of their memorization in the action graph



Figure 4.8: Experiment 2: A scene from the robot's training: (a) Push the Car and (b) Cover the Marker.

4.7.2 Evaluation of the Observation Phase

During the observation phase, the robot observed participants trying to push or cover either the Car or the Marker. All actions were performed once for each trial. Fig. 4.10 shows the robot's camera image capturing participants' actions and successfully estimating and minimizing PE by executing the predicted action primitives. This figure shows the followings:

- (i) (a1, b1): The robot observes the subject and recognizes the action primitives: "reach from side for" (a1) or "reach straight for" (b1). After observing these primitives, the robot predicts the future action primitives "move" (a2) and "hide" (b2).
- (*ii*) (a2, b2): PE increases after our system predicts the future action primitives and the elapsed time is greater than the estimated delay.



Figure 4.9: Experiment 2: Action graph after experiencing x times "reach from the side for" and "move" the Car, and y times "reach straight for" and "hide" the Marker.

Performed primitive	Success rate	$\rm PE$	Delay (seconds)
"Reach from side for"	80%	0.265173	5.12149
the Car		(SD: 0.00063)	(SD: 0.25)
"Reach straight for"	100%	0.266002	5.19768
the Marker		(SD: 0.00127)	(SD: 0.40)
"Reach from side for"	0%	0.0	0.0
the Marker		(SD: 0.0)	(SD: 0.0)
"Reach straight for" the Car	0%	0.0	0.0
the Car		(SD: 0.0)	(SD: 0.0)

Table 4.4: Experiment 2: Experimental results. Performed primitive: primitive performed by the participants. Success rate: percentage of times the robot successfully helped achieving an action. PE: average maximum prediction error measured before PE minimization. Delay: time between the recognition of user's primitive and the onset of the robot's action. The standard deviation is calculated for the 10 trials.

(*iii*) (a2, b3): The robot performs the predicted action primitives to minimize PE, namely "move" the Car (a3) and "hide" (b3) the Marker.

After 10 trials (training and observation), we measured:



Figure 4.10: Experiment 2: This figure depicts the successful cases during which the robot minimized PE after observing an unachieved action. The black and gray lines represent the distance between the human and the robot's hand to the targeted object, respectively. The red filled line denotes the estimated PE, and the dashed line indicates PE threshold above which the robot performs an action to minimize PE. Here, the robot successfully estimates and minimizes PE. (a1, b1): The subject reaches for the red Marker and the blue Car, respectively. (a2, b2): The estimated PE reaches the threshold, and the robot starts its action to try minimizing PE. (a3, b3): The robot's action successfully minimizes PE.

- Performed primitive: the action that the participant was doing.
- Success rate: the amount of time the robot successfully achieved the participants' goal.
- PE: the average PE at the moment of the robot's primitive onset (PE fixed threshold is 60% of the probability of the next primitive).
- Delay: the elapsed time between the first detection of the subject's primitive and the onset of the robot's action.

These results are summarized in Table 4.4.

Results show that the robot could reliably achieve the predicted goals of the participants (success rate: 80% and 100%) within a relatively short five

seconds delay (SD = 0.25 and SD = 0.40). This was only true if the observed actions were previously experienced and had visible effects on the associated objects during the training. It shows that the system could cope with the noisy perception and generate action to minimize PE. In fact, the robot failed once because the participant removed his hand while PE was getting greater than the threshold and tried to perform another action, leading to a robot performing the previously predicted primitive instead.

4.8 Experiment 2 - Discussion

The second experiment intended to show if our system could also exhibit prosocial helping behavior in more complex and noisy environment while interacting with real participants. These new conditions led to variable interaction patterns with the robot. For instance, when asked to try reaching for an object, some participants repeated several times the same primitives to try enacting the robot's action. In contrast, other participants maintained their hand in the same position. These different behaviors generated multiple PE estimation dynamics throughout the experiment. Even with these new challenges, the robot succeeded in helping others achieve their actions by minimizing PE. The results support our hypothesis that the minimization of PE can be used as a behavioral motivation to help others.

4.9 General Discussion and Future Work

The emergence of prosocial helping behavior in infants from 14 months of age is one of the key milestones of their prosocial development. In past decades, several theories, such as the emotional-sharing models and the goal-alignment models, have been proposed to explain the evolution of prosocial tendencies, but few of them clearly described the motivations and mechanisms allowing infants to help others. In this study, we attempted to explain the emergence of prosocial helping behavior in infants based on goal alignment by proposing PE minimization as a behavioral motivation. It can be argued that PE minimization is not the only possible motivation for early prosocial helping behavior, but because of the generality and central role of PE in the brain (see [61, 22]), we chose to mainly focus on this mechanism. To demonstrate our hypothesis, we conducted two experiments to examine to what extent PE minimization could provide a possible motivation for prosocial helping behavior. In addition, we analyzed the effect of action experience on helping.

Our first experiment analyzed the effect of the system's own action experiences on the emergence of prosocial helping behavior. We first showed that prediction error minimization could indeed lead to the emergence of prosocial helping behavior in our robot. We also showed that as the system got more experience, it became able to more efficiently predict action goals and help. The behavior generated by our system was, in some aspects, similar to the comportment observed in infants in Tomassello' and Warneken's experiments [90, 89]. Indeed, their experimental results showed that 14-month-old infants are good in helping "out-of-reach" actions, where the others' goals are easy to predict, whereas older infants could help in more complex and non-transparent situations. Based on these evidence it is evident that the ability to help others is strongly dependent on the robot's (or infant) experience with the involved actions. Therefore, as the robot (or infant) acquires more experience through the interaction with its environment, its ability to understand actions will improve and more extensive prosocial helping behavior will emerge.

In the second experiment, we integrated our model into a humanoid iCub robot and showed that the robot could also perform prosocial helping behavior. This result was not evident as the second experiment with human participants brought a whole new spectrum of challenges. Indeed, due to variable interaction patterns between the human and the robot, the estimated PE was not always stable and could have led to lower success rate. In addition, using the robot's camera images added noise to the detection of objects and others' action. Even with these new challenges, the robot succeeded in generating action to help others in achieving their actions. Results once more showed that the minimization of PE could explain the emergence of prosocial helping behavior.

Based on these two experiments, we confirmed our hypothesis and proved that minimizing PE is a possible behavior motivation to account for the emergence of prosocial helping behavior in robot. Our result cannot prove that our hypothesis also valid for infant development, but can serve as a guide for psychologist to refine their theories and perform new meaningful experiments. We do believe that such results can greatly contribute to the understanding of the development of prosocial tendencies in infants, but also help the creation of more social robots that can be used in our household and in industry.

Despite these promising results, our experiments also showed that the actions performed by our system to minimize PE were not always able to efficiently help others in accomplishing their actions. Indeed, the prediction of the future action primitive was sometimes incorrect, leading to non-appropriate robot response, or the prediction was correct but the robot's action failed to help others achieving their goal. An issue is that, due to the lack of self-other differentiation in our system, the robot does not take others' perspective and executes the predicted action primitive to minimize PE and achieve its own goal, regardless of whether it helped the other achieving his goal. Some literatures show that infants at 14- or 18-month-old are actually able to help others even when the perspective difference should affect their behavior [88, 90] (i.e. handing over an out-ofreach object instead of keeping it). In fact, infants may change their visual perspective while observing others performing actions. This cognitive ability is noted by Tomassello [84] as a socio-cognitive need for infants' prosocial helping behavior.

Moll [57] showed that 24-month-old infants required the perspective-taking ability in order to help others achieve unsuccessful goal-directed actions. However, self-other differentiation, which is needed to perform such perspectivetaking, is not yet acquired by 14-month-old infants [57]. Another possible solution, which does not need change in perspective, is to estimate PE in terms of states and not in terms of actions. Instead of predicting the future action primitive, our system will predict the impact of the observed action on the environment, and minimizing PE would mean achieving the predicted state. Some researches indeed showed that infants first perform actions that help in achieving the goal rather than imitating the means of an action with no predictable goal [12, 67]. Furthermore, it is strongly suggested that infants, from the age of 3 to 5 months, can represent actions in terms of goals, independent of the spatio-temporal properties of the target [80], which supports the idea of employing state prediction over action prediction.

Chapter 5

From Helping to Efficient Human-Robot Collaboration

In the previous chapters, we showed how psychology and neuroscience can give us the theoretical tools to build system capable of acquiring social-cognitive abilities. In Chapter 4, we demonstrated that a robot could develop helping behavior by assimilating others actions as its own and being motivated by the minimization of prediction error.

Here, we use the previously developed mechanism to create a robotic system capable to perform efficient task-oriented collaboration with humans. Instead of trying to unravel issues on the nature of cognitive development, this chapter focuses on demonstrating when and if a robot should minimize prediction error to help others, or if it should wait for being requested to help. Indeed, robots able to efficiently help others will improve productivity and the quality of everyday tasks while reducing the workload. Although such interactions come naturally to human-human teams, achieving similar fluency and comfort in human-robot teams poses many challenges.

Past work introduced in Chapter 2 provided useful insights into *how* a robot should help as part of joint human-robot interaction and what behavior it should display to be perceived social [38, 41, 70, 64, 23, 13, 51]. However, the question



Figure 5.1: Different initiative models for robot assistance during collaborative task executions: human-initiated help (left), robot-initiated reactive help (middle), and robot-initiated proactive help (right).

of *when* a robot should help and how it impacts the participant's perception of the robot was not clearly answered. To shed light on this point, we investigate the factor of **initiative** in robot assistance during **joint task execution**. We ask two questions:

- Should the robot take initiative by spontaneously minimize prediction error, or let the human control the robot's participation in the task?
- When should the robot take initiative?

In the following sections, we first introduce a task representation for joint table-top manipulation and the interaction scenario. We then propose an improved predictive learning model for collaborative robot using a joint task execution system capable of autonomously performing a number of object manipulation tasks as well as monitoring end-to-end human task executions. The different mechanisms for triggering robot assistance in the context of joint table-top manipulation tasks are explained. Next, we perform a user study and evaluate the subjective and objective performances of our system. We complete this chapter with a discussion and analyze our results, contributions and limitations.

5.1 Task Representation and Scenario

In this study, we focus on joint preparation tasks. This category of tasks shares many properties of tasks previously studied in the context of human-robot collaboration (*e.g.*, circuit building [38], Lego model assembly [77], food preparation [23], industrial assembly [64]), including partially ordered action sequencing and shared physical space. More specifically, we consider food tray preparation with n objects, m tray locations and three non-overlapping table regions.



Figure 5.2: (a) Goal states for the two task categories used in our evaluation.(b) Pictorial description of a sample task instance (category Task B), used for explaining the task to participants in the user study.

Objects can be uniquely recognized and their location is represented as a 2D coordinate on the table. For each object, we also represent its relation to other objects and targets with the three predicates is-on(object), is-at(position), and is-in(region). Note that is-on(object) is inferred based on the task knowledge, while the two other predicates are detected directly through the perception module. The table is split into three regions based on who is allowed to manipulate in them. These zones, depicted in Fig. 5.2, are: robot-only (near robot), human-only (near human), and both-allowed (middle). Task goals are represented as a conjunction of instantiated predicates; *i.e.*, the set of relations that need to be true.

Our experiments involve six specific tasks from two task categories (Tasks A and B) in slightly different domains. All tasks in the same task category have the same set of predicates in their initial state and goal descriptions; however specific tasks differ in the particular objects and locations with which the task is instantiated. Task A involves four objects to be placed in four target locations on the tray. Task B involves six objects to be arranged on two locations on the tray. The two task categories are described in Fig. 5.2a and individual task instances are shown in Fig. 5.3.

Both the human and the robot are assumed to perform one task-relevant action: pick-and-place(object, x, y). The x and y coordinates can be anywhere on the table, including particular tray locations or on other objects. The action is applicable for an agent (human or robot) only on objects whose



Figure 5.3: Particular instances of the tasks used in the user study: (a) practice task, (b-d) three instances of Task A, and (e-g) three instances of Task B performed by participants in the three different conditions.

current location is within the regions allowed to the agent. For all our tasks, one object is initially placed in the robot-only region for both tasks; two objects are placed in the human-only region for Task B.

The scenario for each task is then the same: pick all objects individually and place them to their given final position as shown in Fig. 5.2b. For instance, in Fig. 5.3a, the task can be achieved by picking the yellow ball and placing it to the "L" position; and picking the blue ball and placing it to the "R" position.

5.2 Joint Task Execution Model

To study different helping trigger mechanisms, we develop an end-to-end system for joint task execution that allows a robot to perform object manipulation actions as well as monitor the execution of the same actions by a human. This model is based on the idea of predictive learning and the general model introduced in Chapter 3. In this section we present the details of our system.

The joint task execution model is improved based on insights from previous



Figure 5.4: Model for collaborative robots based on predictive learning. It recognizes the current environmental state, predicts the possible future states using a dynamic Bayesian network and generates actions to achieve the desired end-states.

research presented in Chapter 4, which showed that instrumental helping could be generated using prediction error minimization. We suggested that due to self-other correspondence, referred as the like-me hypothesis [56], the robot can project its own task state onto others performing similar acts due. This mechanism allows our system to assist users in achieving their tasks based on the minimization of prediction errors.

The overall system for joint task execution is illustrated in Fig. 5.4. At the core of this system are three modules for (i) perceiving the state of the environment, (ii) tracking the state of the task and anticipating future actions, and (iii) selecting a robot action based on the observed and anticipated states accreting to different help strategies. A more detailed description of these modules is given in the following sections.

5.2.1 Perception Module

The robot can segment and recognize tabletop objects using the point cloud obtained from the robot's RGBD sensor. It uses the Point Cloud Library implementation of tabletop segmentation, which detects the table plane with the RANSAC algorithm. It then extracts a point cloud segment corresponding to each object on the table. If an object is inside or in contact with another object, they are segmented as one object with possibly multiple colors. The robot represents and recognizes objects based on their color, location on the table and size extracted from the segmented point cloud. Color is discretized into six values (red, blue, yellow, green, pink and orange) and size in three values (small, medium and large).

The robot then estimates the current environmental state as the combination of all object states in the scene. The state corresponding to each recognized object is represented by the 3-tuple (Color, Size, Location). The "Location" variable contains one or several of the predicates is-on(object), is-at(position) and is-in(region) presented in Section 5.1. For instance, if a "small red cup" at the location l_1 and a "medium blue plate" in the "human-only" region are recognized in the scene, the object states are noted $s_1 = (Red, Small, is-at(l_1))$ and $s_2 = (Blue, Medium, is-in(human-only))$ In the case the "small red cup" is on the "medium blue plate" at the location l_1 , the corresponding state is noted $s_3 = (Red, Small, is-on(Blue, Medium, is-in(l_1)).$

5.2.2 Predictive Module

Our system uses Dynamic Bayesian Networks (DBN) to predict future states and robot actions that lead to those states. DBNs are multi-time-slice Bayesian networks where variables are connected to one another over adjacent time steps as well as within the same time step. They are computationally efficient generalization of hidden Markov models and have been used to model multi-modal robot behavior in uncertain environments (*e.g.*, work by Huang *et al.* [45, 14]).

For this study, each time-slice of the DBN contains an object state and an action node, corresponding to two multinomial discrete random variables S and A. S can be one of all possible states $\{s_0, s_1, ..., s_N\}$ that are distinct according to the defined predicates for a finite set of objects and named locations. Two states in which an object's position is different but both positions are not at a named location are considered the same discrete state. The variable A is one of all possible action instances $\{a_0, a_1, ..., a_M\}$ that involve the combination of all objects and named locations in the environment, regardless of whether they are available to the human or the robot. S(t) represents the current observed object states in the scene, and S(t+1) the predicted states at time t+1. Within a single time-slice, the state influences the action. Between consecutive time-slices, the



Figure 5.5: Examples of task knowledge represented as states transition. Object o_1 can be one of two states: s_{11} or s_{12} if positioned in l_1 or l_2 , respectively. Object o_2 can be one of two states: s_{21} or s_{22} if positioned in l_1 or l_2 , respectively. The initial environmental state contains two object states: s_{11} and s_{21} .

state and action from the previous time-slice influence the next state.

The DBN encodes the task knowledge in the conditional probabilities noted P(A(t)|S(t)) and P(S(t+1)|S(t), A(t)), which represent the action policies the robot could use if it were to execute the task on its own. Since the tasks are known a priori in our scenario, these conditional probabilities were computed based on the known task structure (Section 5.1), assuming each path for completing the task is equally likely. The conditional probabilities P(S(t+1)|S(t), A(t)) encode the environment and action dynamics and were determined empirically. Future states and actions are predicted by computing the marginal probabilities P(S(t+1)) using Bayesian inference. The action A(t) to perform in order to transit from S(t) to S(t+1) is inferred by maximizing the conditional probability P(S(t+1)|S(t), A(t)). The predictions are then sent to the action selection module.

To illustrate this mechanisms, let us imagine a simple task in which a table containing two objects should be cleaned (see Fig. 5.5). The robot's task knowledge, known a priori, contains the necessary information to represent task. The two objects are "small red cup", noted o_1 , and a "medium blue plate", noted o_2 . The table is separated in two discrete locations: $l_1 = "dirtyZone"$ and $l_2 = "cleanZone"$. Therefore, the possible object states in this example are:



Figure 5.6: Examples of object detection likelihood estimation. The value increases if the object is recognized and decreases when it is not. If the object detection likelihood decreases below the detection threshold, the object is lost.

 $s_{11} = (red, small, is - in(l_1)),$

 $s_{12} = (red, small, is - in(l_2)),$

- $s_{21} = (blue, medium, is-in(l_1)),$
- $s_{22} = (blue, medium, is-in(l_2)).$

The initial environmental state contains two object states s_{11} and s_{21} . The robot can perform pick and place actions, noted a_1 and a_2 , to move the "small red cup" or the "medium blue plate", respectively, from l_1 to l_2 .

The transitions between the different object states as described in the task knowledge are illustrated in Fig. 5.5. When s_{11} and s_{21} are initially recognized by the robot, marginal probabilities P(S(t+1)) are calculated individually for each object state. As we assume all paths for completing the task are equally likely for this task, we obtain: $P(S(t+1) = s_{11}) = 0$, $P(S(t+1) = s_{12}) = 0.5$, $P(S(t+1) = s_{21}) = 0$ and $P(S(t+1) = s_{22}) = 0.5$. The system then infers what action to perform in order to achieve the predicted environmental states by maximizing the conditional probabilities P(S(t+1)|S(t), A(t)), which are in this case equal to:

$$\begin{split} &P(S(t+1)=s_{12}|S(t)=s_{11},A(t)=a_1)=1,\\ &P(S(t+1)=s_{12}|S(t)=s_{11},A(t)=a_2)=0,\\ &P(S(t+1)=s_{22}|S(t)=s_{21},A(t)=a_1)=0,\\ &P(S(t+1)=s_{22}|S(t)=s_{21},A(t)=a_2)=1. \end{split}$$

Here also, conditional probabilities are calculated individually for each object state.

In this example, the robot estimates that it can perform a_1 or a_2 in order to reach s_{12} or s_{22} , respectively, from the currently observed object states s_{11} and s_{21} . When a new state is reached, the robot reiterates the same inference process until it can no longer predict new states.

When to perform an action and which action to execute is decided by the action selection module presented in the next section.

5.2.3 Action Selection Module

The action selection module implements a policy that specifies what the robot should do at each time step. If the robot were to execute the task completely on its own, this module would directly return one of the possible actions predicted by the DBN immediately after every action. During joint task execution, on the other hand, the robot's policy needs to account for the human's direct input or their actions that result in changes in the world state. We implement three policies that differ in terms of *when* a robot action is triggered.

- Robot-initiated Reactive Help (R): Robot actions are initiated by the robot when it estimated prediction error higher than a fixed threshold. The robot monitors the human's task execution and detects when one of the next states predicted by the DBN is not reached within an expected time window, indicating a delay or difficulty in the task progress.
- Robot-initiated Proactive Help (P): This policy involves performing actions whenever they are possible, even when low prediction error is estimated. However, different from a robot-only task execution, the robot needs to take into account human actions that might be *in progress* before a stable environmental state is reached. If at least one executable action exists that does not conflict with the human actions, the trigger is initiated.
- Human-initiated Help (H): This policy gives complete control of robot actions to the user. The robot performs an action only when the user explicitly says "Robot, can you help me?' and ignores prediction error.

Mechanisms behind the robot's behavior for each policy are similar, but differ in some fundamental aspects. When the system observes object states S(t), it predicts future object states S(t+1) that have non-null marginal probabilities. Object states are noted s_{ij} , where *i* is the object number and *j* represents the number of states the object can be on. Actions on the objects *i* are noted a_i . To decide which action should be performed by the robot, several values are estimated.

- Firstly, for each possible future object state, an object detection likelihood, noted $\mathcal{L}s_{ij}(t+1)$, is calculated. The value is initialized at 0.6 and increases linearly while the object is recognized (max 1). This value represents how well an object corresponding to a predicted state is recognized by the perception module. If the object is momentarily not perceived, $\mathcal{L}s_{ij}(t+1)$ decreases linearly. If $\mathcal{L}s_{ij}(t+1)$ becomes lower than 0.4, the object is considered lost. For instance, if a user repeatedly touches an object in the scene, the corresponding $\mathcal{L}s_{ij}(t+1)$ will be low because the object recognition will be noisy. Examples of object detection likelihoods for good and bad object recognitions are represented in Fig. 5.6.
- Secondly, a prediction error, noted $PEs_{ij}(t+1)$, is estimated for each possible future object state. The value of $PEs_{ij}(t+1)$ is a function of $\mathcal{L}s_{ij}(t+1)$ and of the elapsed time, noted t_e , since the current environmental state S(t) has been first recognized.

A prediction error $PEs_{ij}(t+1)$ triggers an action when it gets higher than a threshold (θ) fixed at 0.8. The robot then executes the corresponding action a_i . The prediction error is calculated as follows:

• In condition R, when possible future states are predicted, the prediction error values are calculated as function of the elapsed time t_e and the corresponding object detection likelihood as follows:

$$PEs_{ij}(t+1) = 0.3 \times \mathcal{L}s_{ij}(t+1) \times (1 - \frac{2T - t_e}{T});$$
(5.1)

where T is the estimated action duration, fixed at 4 seconds empirically.

If one prediction error gets higher than the threshold θ , it triggers a robot action. If two or more prediction errors are higher than θ at the same time, the one with the action on the closest object to the robot is activated. The distance between the robot and an object is noted d_i .

• In condition P, when possible future states are predicted, the prediction error values are calculated as function of the corresponding object detection likelihood only:

$$PEs_{ij}(t+1) = \mathcal{L}s_{ij}(t+1);$$
(5.2)

The prediction error is therefore maximum as soon as objects are detected with enough certainty.

If one prediction error gets higher than the threshold θ , it triggers a robot action. If two or more prediction errors are higher than θ at the same time, the prediction error with the lowest d_{ok} is activated.

• In condition H, when a user asks the robot for help, the trigger signal with the highest object detection likelihood $\mathcal{L}s_{ij}(t+1)$ value is activated. If two or more trigger signals have the same object detection likelihoods, the trigger signal with the action on the closest object to the robot is activated. The distance between the robot and an object is noted d_i .

The robot always uses the gripper closest to the object of the executed action.

Let us now consider the example where our system and a user jointly collaborate during a task, as presented in Section 5.1 (see Fig. 5.5). At first, the robot detects the current environmental state, noted S(t). The tasks state prediction module then predicts possible future object states $S(t+1) = s_2$ or $S(t+1) = s_4$. The action to reach s_2 is estimated to be a_1 . The action to reach s_4 is estimated to be a_2 . When the states are predicted, the robot estimates for each of them an object detection likelihood value $\mathcal{L}s_{2(t+1)}(o_1)$ and $\mathcal{L}s_{4(t+1)}(o_2)$.

In all conditions, if the next state is correctly reached by the robot or by the user, the tasks state prediction module predicts new future states. In this example, if s_2 is reached, the new predicted states will be $S(t+1) = s_4$. Conversely, if s_4 is reached, the new predicted states will be $S(t+1) = s_2$.

In the R condition, if the predicted states are not achieved within few seconds, prediction errors will increase. If the user performs an action before any of the prediction errors reach the threshold, the robot does nothing. Else, if one of the prediction error reaches the threshold, the robot performs the corresponding action. In the P condition, the robot perform an action as soon as possible, namely when one of the object detection likelihood corresponding to predicted states is higher than the threshold. It is similar to the R condition with a much lower threshold. In the H condition, the robot will not perform any of the action until it receives a command.
5.3 Experimental Setup

In this section we present the details of our experiment and user study conducted to analyze the effects of initiative on human-robot interaction.

5.3.1 Platform

Our system is built around the PR2 robot platform (see Fig. 5.1). PR2 has two 7 degrees-of-freedom arms giving it a large workable space for tabletop manipulation tasks. Each arm has 1 degree-of-freedom parallel-finger gripper that can grasp objects up to a width of 8cm. PR2's arms are passively balanced and actuated with low-power motors, making it safe to work around humans. For perception, it has a Kinect sensor attached to the head that has a high-speed pan and tilt motion. Note that most of the system was designed independently of the platform while the action execution part was designed for and with the PR2.

5.3.2 Robot Actions

The robot's pick-and-place actions are parametrized with an *object* to be picked and a *location* at which the object is to be placed. The actions are defined as a sequence of poses relative to the object (pre-grasp, grasp, and lift poses) followed by poses relative to the target location (transfer, lower, and drop poses). While the overall action templates remain the same, some of the poses in the actions are tuned to the particular object being manipulated. The actions were trained using a learning by demonstration approach developed by Alexandrova *et al.* [1].

5.4 User Study

The help trigger mechanisms described in Section 5.2.3 are expected to yield different joint task execution dynamics. Furthermore, each mechanism on its own can result in a wide variety of behaviors depending on the particular user. For example, when interacting with the *human-initiated* policy, users may request help at every step or only when they need it. When interacting with the *robot-initiated proactive* policy, they might select their own actions such that the robot has many opportunities to help or they might (unintentionally or intentionally) block the robot's actions. The differences across and within each policy

can reflect on objective task execution measures, as well as the user's subjective attitude towards the robot. To investigate these differences, we performed a user study that allows us to (i) characterize people's behaviors while interacting with each policy, and (ii) compare the alternative policies for triggering robot help.

5.4.1 Study Design

We performed a within participants study with one independent variable (robot helping behavior) with three conditions: H, R, P (Section 5.2.3). In each condition, participants performed two tasks with the robot, one from each category (Task A and B). The order of the three conditions were counterbalanced.

5.4.2 Study Setup

The robot was placed in front of a 68cm high table. Participants sat across the table. The table top was separated into three zones as shown in Fig. 5.3. Participants were asked not to touch objects that are in the red zone (near the robot). Similarly, the robot could not enter the blue zone (near the human). Both were allowed to manipulate objects in the middle zone. In the middle of the table there was a tray with four target positions.

Tasks were explained to participants with a one-page pictorial description involving (i) the set of objects and targets involved in the task and (ii) the final state of the tray when the task is complete. An example task description is shown in Fig. 5.2b. An additional small table was placed to the right of the participant. Printed task descriptions were placed on this table, together with a tablet for logging task steps and a laptop for responding to our questionnaire. The complete setup can be seen in Fig. 5.1.

5.4.3 Procedure

Participants were recruited from a campus and nearby neighborhoods through mailing lists. Interested individuals signed up for a 45 minutes time slot in advance. When participants arrived at their scheduled study time, we first explained the purpose of the study and asked them to sign a consent form. Then they were taken to the participant seat, introduced to the robot and the workspace, and given an overview of the procedure. Next, the robot was activated and participants performed a practice task (Fig. 5.3a). They were explained what the task is using the corresponding pictorial description. The robot made a specific sound to indicate that it was ready. Participants were told that they can start the task when they hear this sound. They were told to perform one step of the task and then log the step on the tablet. The logging was done throughout the study as a mechanism to space human actions apart and give the robot an opportunity to detect intermediate states of the task. Each log required indicating who performed the step (human or robot), the two letter identifier for the object involved (as indicated in the task description), and the one letter identifier for the target position where the object was placed. The second step of the task was performed by the robot to familiarize participants with the robot's motion. The robot made another sounds when it detected the task completion. Participants were told that they will perform similar tasks together with the robot in three conditions where the robot's behavior will be different.

Next we moved on to the actual study. For each condition, the experimenter first gave condition specific instructions. In the human-initiated help (H) condition, participants were told that they can request the robot's help by saying "Robot, can you help me?" This was done in a wizard of Oz fashion and without using a microphone. As soon as users asked for help, the experimenter discretely pressed a button. In the other conditions (R and P), they were told that the robot will decide when and how to help out with the task. Then the experimenter set up the initial state of the first task, told participants to start when they hear the robot sound, and left them alone with the robot. The experimenter came back to set up the next task after the robot detected that the task was complete. After completing both tasks in the same condition, participants were asked to respond to the condition-specific questionnaire. After the three conditions were complete, participants responded to additional questions drawing comparisons between the three conditions. At the end participants were thanked for participating and given the promised compensation of 10 USD equivalent gift card.

5.4.4 Measurement

The study was recorded with two cameras; one mounted on the robot's head and another overseeing the workspace together with the robot and the participant. In addition, we logged the progression of tasks and robot actions with timestamps throughout the study (see Appendix .1 for complete questionnaire). The extracted data was used to evaluate three main components: the quality of interaction and the system performance.

From the study logs we extracted the task completion time and the number of actions performed by each agent. From the videos we extracted quantitative measure that characterized each participant and the robot behaviors. These measures included times when the robot and the human were moving alone or in concurrence, their idle times and the number of s the participants performed to the face of the arms of the robot during the joint task execution. The coding was performed by two coders (IRR¹= 0.72), including one without prior knowledge of the study.

To compare the three conditions subjectively from the user's perspective, we administered several questions after each condition as well as at the end. First we asked an open ended question to elicit the participants own description of the robot's assistance behavior. Another question asked them to describe their strategy. Then we asked a set of Likert scale questions, similar to those commonly used in human-robot collaboration research [39]. These questions addressed the user's perception of: the robot's helpfulness, its awareness of the human and task progress, its contribution to the task, team fluency and efficiency, and naturalness of the interaction (see questions in Fig. 5.12). Additional questions at the end asked a forced ranking of the three conditions and open ended questions about perceived distinction between the two robot-initiated conditions and how different behaviors would be combined in an ideal interaction. The complete questionnaire and answers are provided in the Appendix .2.

5.5 Experimental Results

Our study was completed by 18 participants (9 females, 9 males, and ages 18 to 35). This section presents our findings based on data collected from these participants. A one-way ANOVA was conducted to compare the effect of conditions H, R and P on the different objective metrics. We performed post-hoc tests (two-tailed paired-t-test) to explore differences between pairs of conditions.

¹Cohen's kappa.



Figure 5.7: Examples of interactions.

To analyze the experimental data, we segmented each interaction between a participant and the robot into temporal action sequences. These actions could be of 3 types for the participants: acting (upon the table), logging or idling; and 2 types for the robot: acting or idling. Based on these temporal action sequences, we could also extract concurrent actions between the participants and the robot. Additionally, we segmented the different gazing patterns of the participants to the robot's face and arms. Two examples of interactions for each of the three condition during the Task B are shown in Fig. 5.7(a-f).

5.5.1 Objective Metric

We first examine common task and collaboration metrics. Fig. 5.8a shows the average number of task actions performed by the robot in each condition (Task A: F(2,51) = 15.35, p < .001; Task B: F(2,51) = 13.24, p < .001) and Fig. 5.8b shows the overall task completion times by the human-robot team (Task A:

F(2,51) = 2.20, p = .12; Task B: F(2,51) = 6.15, p < .005).

Fig. 5.9(a-d) show the breakdown of task completion times into robot-only, human-only, concurrent, and no motion segments and Fig. 5.9(e-f) separately show the human idle time and robot idle time. The results of the ANOVA for results in Fig. 5.9 are as follows: (a): (Task A:F(2,51) = 7.64, p < .005; Task B: F(2,51) = 12.62, p < .001); (b): (Task A: F(2,51) = 10.51, p < .001; Task B: F(2,51) = 0.61, p = .55); (c): (Task A: F(2,51) = 4.41, p < .05; Task B: F(2,51) = 2.79, p = .07); (d): (Task A: F(2,51) = 4.78, p < .05; Task B: F(2,51) = 6.97, p < .005); (e): (Task A: F(2,51) = 2.73, p = .070; Task B: F(2,51) = 9.22, p < .001); (f): (Task A: F(2,51) = 5.32, p < .01; Task B: F(2,51) = 5.61, p < .01).

Finally, Fig. 5.11(a-d) shows the number of times the participant looked and the robot's face and arms and the average duration of the gazes. The results of the ANOVA for results in Fig. 5.11 are as follows: (a): (Task A: F(2,51) = 1.99, p = .14; Task B: F(2,51) = 8.14, p < .001); (b): (Task A: F(2,51) = 5.19, p < .01; Task B: F(2,51) = 4.75, p < .05); (c): (Task A: F(2,51) = 5.68, p < .01; Task B: F(2,51) = 2.61, p = .084); (d): (Task A: F(2,51) = 1.31, p = .28; Task B: F(2,51) = 3.32, p < .05).



Figure 5.8: (a) Number of actions performed by the robot for each task category in each condition. (b) Task completion time for each task category in each condition. Error bars represent standard deviation.

Proactive Versus Reactive

First we focus on the comparison of robot-initiated help strategies. Proactive help results in the robot having a greater contribution to the task, as indicated



Figure 5.9: Breakdown of task completion times into (a) robot-only, (b) humanonly, (b) concurrent, and (d) no motion time segments. These include only motion related to the joint task. (e) Human idle time. This excludes the time during which the human is performing their secondary task of *logging* task actions. (f) Robot idle time.

by the significantly higher number of actions performed by the robot (Task A: p < .001, Task B: p < .001) (Fig. 5.8a). This is also reflected in the significantly lower robot idle times for the proactive robot (P) as compared to the reactive robot (R) (Task A: p < .001, Task B: p < .05) (Fig. 5.9f). The average number of actions performed by the reactive robot was around 1 (Task A: M = 1.17, SD = .38, Task B: M = 1.56, SD = .76), which is the minimum number of actions required by the robot. Whereas, the proactive robot performed around 2 (Task A) and 3 (Task B) actions (Task A: M = 2.17, SD = .48, Task B: M = 1.56, SD = .76)

3.00, SD = .82), which are about half of the actions needed to complete the task. This finding is expected and confirms that our model produced the intended behavior.

Despite the difference in the number of robot actions, there was no significant difference in the total task durations in Task A (Task A: p = .12) and little difference in Task B. A potential reason for this could be lack of parallelization between human and robot actions. However, the significant increase in the *concurrent* human-robot motion (Fig. 5.9c) in the proactive condition indicates that parallelization did indeed happen at least in Task A (Task A: p < .005). In addition, the total task durations appeared to be greatly influenced by the difference in human and robot action speeds as humans are several orders of magnitude faster at pick-and-place actions. Hence they were not slower in completing the overall task in the reactive condition. Despite this difference, human idle times were not significantly higher in the proactive robot condition (P) (Fig. 5.9e).

Human-initiated versus Robot-initiated

Next, we look at comparisons between the human-initiated help (H) condition and robot-initiated help conditions to characterize how people chose to get help from the robot when they had control. From Fig. 5.8a, we see that the number of actions performed by the robot in the H condition was about half of all task actions, as in the P condition. The number of actions performed by the robot was significantly higher than in the R condition (H-R - Task A: p < .001, Task B: p < .001). It resulted in significantly higher concurrent motions in Task A for the H condition compared to the R conditions (H-R - Task A: p < .05) (Fig. 5.9c). We believe that it is because participants asked for help and then started doing their own actions as soon as they understood the robot's intention. This is similar to the P condition, where participants briefly waited until they recognized what the robot was doing and then acted. The added waiting time in the H condition was reflected in overall task completion times (Fig. 5.8b), which was significantly higher than in the R condition for Task B (H-R - Task B: p < .005) and in the P condition for both tasks (H-P - Task A: p < .05, Task B: p < .05). This was also reflected in the human idle times (Fig. 5.9e) which was highest for the H condition in both tasks (H-R - Task A: p = .27, Task B: p < .001; H-P - Task A: p < .05, Task B: p < .01). We noticed that one participant made the robot do all actions for Task 1; two participants made the robot do all possible actions for Task 2 in the H condition. This contributed to the high human idle time and task completion time, while making the variance in this condition high.

Face gaze zone

Gaze toward the Robot

Figure 5.10: Gazing zones: The face gaze zone (purple line) is situated on the robot's "head" part. The arms gaze zone (red line) is situated on the lower body part of the robot.

We then look at the participants' gazing patterns toward the robot during the different tasks. Two gazing targets were analyzed: the robot's face and arms zones, which are shown in Fig. 5.10. The number of times the participants looked at each zone and the duration of each gaze were extracted from the video recording of the experiment. The gazes to the robot's arms were only counted when the robot was moving its arms.

The number of gazes to the face and arms of the robot are described in Fig. 5.11a and Fig. 5.11b, respectively. The average duration of each gaze to the face and arms of the robot are described in Fig. 5.11c and Fig. 5.11d, respectively.

The participants gazed significantly more to the face of the robot in the H condition compared to the R and P conditions during Task B (H-R - Task B: p < .05; H-P - Task B: p < .001) (see Fig. 5.11a). The average gaze duration



Figure 5.11: (a) average number of gaze to the face of the robot; (b) average number of gaze to the robot's arms; (c) average duration of each gaze to the face of the robot; (d) average duration of each gaze to the robot's arms.

was also significantly longer in the H condition during Task A (H-R - Task A: p < .005; H-P - Task A: p < .005) (see Fig. 5.11c). This can be explained by the participants having to vocally command the robot when needing the robot's help. In fact, the participants almost always gazed to the robot's face when asking for help and kept gazing until the robot would start its action.

The number of gazes to the arms of the robot is significantly greater for the H and P conditions compared to the R condition in Task A (H-R - Task A: p < .05; R-P - Task A: p < .001) and greater in condition P compared to the R condition in Task B (R-P - Task B: p < .001) (see Fig. 5.11b). This result is strongly correlated with the number of actions performed by the robot (see Fig. 5.8a). Indeed, gazing at the robot's arms allows the participants to identify its actions and to log them as part of the task. Therefore, if the robot executes more actions, the number of gazes should increase proportionally.

Finally, the average gaze duration to the face and arms is significantly longer in the H condition compared to the two others in Task B (H-R - Task B: p < .05; H-P - Task B: p < .05) (see Fig. 5.11d). The average gaze durations appear to be correlated to the tasks average completion times (see Fig. 5.8b). In case of face gazes, this can be explained by the fact that users looked at the robot's face after asking for help and kept looking until the robot started moving. On the other hand, the reason for the larger duration of gazes to the arms in the H condition can only be hypothesized to be due to a higher willingness of the participants to verify the effect of their commands on the robot.

5.5.2 Subjective Metric

Participant responses to the Likert-scale questions are summarized in Fig. 5.12. The inter-condition differences were analyzed using the Wilcoxon signed rank $test^2$, which is a standardly used non-parametric test. As suggested in [11] and to avoid family-wise errors, we grouped the seven scales into two sub-scales representing the quality of interaction (Fig. 5.12a) and the system performance (Fig. 5.12b). There were no statistically significant differences between the human-initiated help (H) and proactive robot (P) conditions in any of the subscales, despite the differences observed in objective metrics (e.q., the task completion time shown in Fig. 5.8b) between these two conditions. Subjective ratings of the quality of interaction appeared to be correlated with the number of actions performed by the robot (Fig. 5.8a), rather than the overall task efficiency (Fig. 5.8b). The reactive robot (R) condition was rated significantly lower than the other two (H and P) conditions, indicating that participants agreed significantly more that the quality was better in the H and P conditions (see Fig. 5.12a). Whereas the significant differences were observed in the quality of interaction, participants did not rate differently the system performance. It seems they did not attribute the robot's behavior in the R condition to its inability to perceive the human or keep track of task progress. In the forced ranking question administered at the very end of the study, 72% of participants (13/18) indicated P as their most preferred behavior, while 22% (4/18) indicated H and only 6% (1/18) indicated R. 78% of participants (14/18) indicated R as their *least* preferred behavior, with 17% (3/18) for H and 6% (1/18) for P. The question yielded a clear ranking of the three conditions as P > H > R from most preferred to least preferred.

Furthermore, in a separate two-choice questions, 67% of participants (12/18) indicated they prefer letting the robot take initiative, while the remaining 33% said they preferred having control over the robot's actions. These results demonstrates the take initiative in the take the take the take the take takes a statement of the take takes a statement of the take takes a statement of the takes a statement of the take takes a statement of the takes a statement of takes a statement of

 $^{^{2}}$ We also conducted parametric tests and obtained similar results.



Figure 5.12: Mean Likert-scale ratings in questionnaire responses. Significant differences according to Wilcoxon signed rank tests are indicated with p-value ranges.

strate that although there were no significant differences between the H and P conditions in the Likert-scale ratings, people are more likely to prefer P over H in favor of the improved objective metrics (Section 5.5.1).

Perceived differences of robot strategies

An open-ended question asked participants to describe the differences between the two conditions R and P in which the robot decided when to act, if they noticed any difference at all. All participants reported that they noticed a difference. The reactive robot was perceived as "slow" and characterized as "lazy" and "hesitant" by some of the participants. The proactive robot, on the other hand, was perceived "fast" and "pro-active". Descriptions of the perceived robot behaviors were accurate; for example:

• M, 35: "... [P] felt more natural to have unprompted collaboration while I was performing the task, rather than the robot waiting for me to finish as it did during [R]".

- M, 20: "[P] was more **proactive** in its help ... [R], by contrast, would only complete actions that I was unable to complete".
- M, 22: "[In P] the robot took the *initiative* a lot more than [R]".

These answers clearly highlight that the participants understood how the robot was behaving in the P condition, but did not felt that the robot was very motivated to help in the R condition. This understanding of the proactive robot's "mind" may be the reason why participants rated this condition as their favorite.

Collaboration enhancing human behaviors

The differences in the objective and subjective task metrics can be further dissected by examining the occurrence of certain events. Firstly, we saw that concurrent motion was significantly higher in the P and H conditions for Task A (Fig. 5.9c), which shows better team work took place in these conditions. Secondly, in Task B, two objects (a container and a ball) were placed in the human-only zone. We observed that most people intuitively encouraged collaboration by starting tasks with objects that were in the human-only region of the table. Indeed, in the H and P condition, only 3 participants in average did not start with the ball in the human-only zone. In the R condition, 7 users started with one of the balls placed in the both-allowed zone, showing lower collaboration. Participant descriptions of their strategies, in a free form question in the questionnaire, reflected their intent to enhance the collaboration; for example:

- F, 22: "[In R] I chose objects closest to me or that were obscuring the place of the objects needed to be. I also moved slower than I would without the robot to give it time to help".
- M, 19: "[In H] I first wanted to set up the two bowls on the table before putting any of the balls in. This was to ensure that [the robot] would not attempt to put a ball in a space without a bowl".
- F. 24: "[In P] I Moved the bowls and objects from the blue zone first and then help fill them [the bowls] one at a time".

One participant placed the objects from this zone into the common central zone to make the robot perform the task actions while he would perform the logging task. He described his strategy as: • M, 19: "[In P] I moved objects from the blue zone into the collaboration zone, and placed objects in-between logging and [the robot's] actions".

5.6 Discussion and Future Work

We implemented three initiative conditions to evaluate the robot's performances during joint task collaboration with humans. The first condition, robot-reactive help enables the robot to help when the prediction error of a predicted state is higher than a large threshold, which is similar to our hypothesis for the emergence of prosocial helping behavior in Chapter 4. The robot-proactive help condition initiates actions as soon as possible, which is similar to reducing the prediction error threshold to 0. Finally, the human-initiated condition gives the control of the robot's actions to the human.

In all conditions, the robot successfully helped the human to accomplish the different tasks it was given. In addition to the results presented in Chapter 4, it shows that prediction error minimization can allow a robot to help others. However, this study did not focus on explaining the emergence prosocial helping behavior, but on *whether* and *when* a robot should help to be perceived efficient and helpful.

On the question of *whether* a robot should take initiative, our results demonstrate that the answer depends on the robot's behavior. People were happy to give away control if the robot is proactive, but they would rather have control if it is reactive. Given its other benefits in terms of objective task and team metrics, this suggests that collaborative robots should be designed to always be proactive. In practice, this might not always be possible. Challenges such as partial task knowledge and uncertain perception might reduce the robot's ability to help the user when it is actually possible for it to help. While the simplistic help request used in our experiments would not be sufficient, enabling users to ask for particular types of help by commanding actions could result in more effective collaboration in such circumstances.

Our study also demonstrate that the behavior of the proactive robot was similar to the behavior people asserted when they had control over the robot's actions. In turn, the similar high subjective rating of the proactive and the human-controlled robots could be partially ascribed to this similarity. Furthermore, we believe that the behavior that was common in these two conditions is similar to how a human would collaborate in the same role. Participants indeed described the robot's behavior during the proactive condition much better than for the reactive one, attesting of a better understanding of the robot's "mind". They then indicated that the collaboration was most natural.

On the natural aspect of the interaction, we show that the proactive or the reactive robots receives less gazes to the face than the human-initiated robot. It is argued by psychologists that face-gaze (or eye-gaze) is an important component of social interactions [26, 16] and mutual understanding [84, 60], which contribute to fluent and natural interactions. Therefore, in scenarios where a robot is expected to interact socially and naturally with humans, the proactive and reactive robots may be less appropriate than the human-initiated one.

The overall implication of our study is that mixed-initiative help triggers seems to be ideal for collaborations in realistic settings. We suggest that the robot should be first help in response to human request or reactively and slowly become more proactive. We believe that increasing pro-activity over time, after observing the user's collaboration preferences (*e.g.*, [64]) might improve the collaboration, while benefit from the social aspect of the early human-initiated behavior. This mix-initiative process seems similar to human development, in which infants start helping to minimize prediction error and slowly help more pro-actively as they acquire higher social motivations.

Finally, the model used in this chapter answered the question asked in Chapter 4 on whether or not state prediction would solve the issue of perspective for helping scenario. The results of this experiment indeed showed that the perspective taking is no longer affecting the robot's behavior if only the state of the environment is predicted. However, the system is no longer to directly take into account the actions of the user, which in definitive can be disadvantageous for the interaction. We would therefore argue that predicting both the environmental state and others' action would solve the issue of perspective taking while adding the user in the prediction loop.

Chapter 6

Conclusion

6.1 Summary of our Study

This dissertation aim was twofold. First, we attempted to explain how human infants are capable to help others prosocially from as early as 14-month-old. Second, we wanted to replicate the development of prosocial helping behavior into an artificial agent and design an efficient assistive robot. To reach these objectives, we performed a series of experiment to answer the following issues:

- 1. What is the motivation for infants to start prosocially helping other?
- 2. What is the role and effect of cognitive maturity on the development of prosocial helping behavior?
- 3. How to design an efficient autonomous assistive robots inspired by infants' development?

To address these points, we first reviewed the literature in psychology, cognitive developmental science, but also robotics in chapter 2. The work related to infant development gave us useful insight on when and how infants start to help. The theories we introduced highlighted the possible mechanisms thought to play roles in the emergence of prosocial behavior, such as the emotion-sharing and goal alignment models. In addition, similar robotics studies guided our choice of tools to use in order to design a robotic system that could closely replicate infants' development and perform efficient interactions. In chapter 3, we introduced our general model based on predictive learning, which describes a learning paradigm for artificial systems. It suggests that by interacting repeatedly with the world in various contexts and with different behaviors, an agent can learn to predict action-state sequences by building its sensorimotor predictor. The system then becomes able to foresee the effect of its actions on the environment and to perform meaningful actions that minimize prediction error.

Based on predictive learning, we proposed a general motivational mechanism to explain the emergence of prosocial helping behavior in infants in chapter 4. We suggest that due to low self-other differentiation in early infancy, infants assimilate others' action as their own and used they sensorimotor model to predict the outcome of observed actions. We hypothesized that infants may be motivated to help others to minimize the prediction error that is then estimated over the assimilated others' goal. Our experimental results validated our hypothesis by showing that robot endowed with the same cognitive and motor capabilities as a 14-month-old infants could spontaneously help others, motivated by the minimization of prediction error. In addition, we showed the importance of cognitive maturity, and especially the positive effect of action experience on action understanding for the emergence of prosocial helping behavior. This answered the first and second aforementioned question.

In chapter 5, we showed when and if a robot should take initiative help others during joint task human-robot collaboration. We implemented three initiative conditions: robot-initiated reactive help, robot-initiated pro-active help and human-initiated help. The reactive robot triggered help when the estimated state prediction error was higher than a large threshold (reactive to error in prediction). On the other hand, the proactive robot initiated actions as soon as possible, even when prediction error was low. Finally, the human-initiated condition waited a user request to trigger helping and ignored prediction error. We performed a user study and showed that the robot taking initiative proactively, without waiting for high prediction error, had better performance and was preferred by users. However, it appeared that the human-initiated condition resulted in larger number of gaze toward the robot's face than during both the reactive and proactive robot. According to some studies [51], it may represent an important component of social interactions. The number of gaze to the robot's face seems to be related to the fact that users had to verbally address the robot, making the social hypothesis more plausible. These results addressed the third question stated above.

6.2 Contributions

The work presented in this dissertation and summarized above, contributes to a field known as cognitive developmental robotics, which is at the intersection of engineering and science. It attempts to study and widen the understanding of infant development though the conception of artificial systems, themselves designed based on scientific evidences. Our experiments involved solving original and novel technical challenges along with suggesting new hypothesis and mechanisms to explain the emergence of prosocial abilities in infants. Our specific contributions are separated into two categories and described in the following paragraph.

Scientific Contributions

To help understanding infant development, we suggest a model based on predictive learning that emulated the development of prosocial abilities. In particular, we suggested a motivation mechanism based on the minimization of prediction error to explain the emergence of prosocial helping behavior. Through our experiments, we showed that an artificial agent could spontaneously help others by assimilating their action goal and to minimize prediction error. Additionally, we showed that cognitive maturity, and in particular the ability to understand action based on oneself sensorimotor model, directly affected the variety and efficiency of helping behavior. Our results were similar to what was observed in infant studies by Tomassello and Warneken [90, 89]. Namely, that 14-month-old infants are good in helping others when goals are easy to predict, whereas older infants can help in more complex and non-transparent situations. By taking advantage of such result, we hope that our work can help or guide psychologist toward revising and/or improving the current theories on the emergence of prosocial helping behavior.

Engineering Contributions

The conception of an infant inspired robotic agent able to prosocially help other required to overcome a certain number of challenges and resulted in several engineering contributions. We proved that predictive learning, and in particular the minimization fo prediction error, could endow a robot with the ability to help and collaborate with others. In addition, our model did not use a "model of others" as done in other studies [34]. We then showed that proactive robots help others more efficiently, naturally and fluently in joint task collaboration, completing insight from other human-robot interaction studies [3, 63, 47]. Finally, we showed that even if proactive robots are preferred by users, the humaninitiated robots led to more natural interactions with more face gazes.

6.3 Limitations and Future Work

As summarized above, our studies provide great insights on how infants may acquire early prosocial helping behavior through the minimization of prediction error. It also suggests tools and guidelines to design efficient and fluent collaborative robots, fit to interact with humans. However, our work only scratched the surface of infant development and robotics design, and is limited in some aspects. Firstly, while our proposed mechanism can provide hint to support the goal alignment model presented in Chapter 2, we cannot provide any proof strictly validating our claims. In addition, the proposed mechanism only accounts for the emergence of prosocial helping behavior and does not provide insight on other prosocial abilities such as imitation, sharing and caring. Finally, the task and interaction performed by the robot in our experiments were, to some extent, limited and have little concrete application in real world.

6.3.1 Extension of our Model

Our model focused so far on explaining the emergence of instrumental helping behavior. However, not only infants can help others in need, but they are known to display a wide range of prosocial behavior. They can comfort distressed persons, imitates their motions or intended goals (see [55]), share food with their caregivers, etc. While these behaviors are also observed very early in infancy, alike prosocial helping behavior, the motivation to perform them is unclear.

We suggest that prediction error minimization might also be involved, to some extent, for their development. Indeed, imitation might be explained as an attempt to minimize the prediction error by executing a perceived motion and improve the sensorimotor predictor. To prove our theory right, we will perform new experiments in which the system will interact with other, not only after learning its sensorimotor model, but also while learning it. We expect to see the emergence of imitating behavior as well as helping behavior as a byproduct of the minimization of prediction error along with other motivational mechanism such as curiosity [33, 65].

Additionally, we shortly highlighted in Chapter 5 the important of having social assistive robots for companies or household. However, while our system was capable to efficiently and fluently interact with users, the complexity and range of its capacity was limited. By allowing our system to acquire a wider range of prosocial abilities, but also to train longer with more actions in more complex environments, we hope to observe the emergence of robust behavior usable in real situations.

Finally, we will attempt to collaborate with psychologists and cognitive science teams to perform a new experiment to prove our hypothesis. By carefully designing the conditions in which we expect to see infants to help others, we expect to be able to prove that the prediction error minimization can explain the emergence of prosocial helping behavior.

Bibliography

- S. Alexandrova, M. Cakmak, K. Hsiao, and L. Takayama. Robot programming by demonstration with interactive action visualizations. *Proceedings* of Robotics: Science and Systems, Berkeley, USA, 2014.
- [2] M. Asada, K. F. MacDorman, H. Ishiguro, and Y. Kuniyoshi. Cognitive developmental robotics as a new paradigm for the design of humanoid robots. *Robotics and Autonomous Systems*, 37(2):185–193, 2001.
- [3] M. Awais and D. Henrich. Proactive premature intention estimation for intuitive human-robot collaboration. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, pages 4098–4103. IEEE, 2012.
- [4] C. D. Batson. The altruism question: Toward a social-psychological answer. Psychology Press, 2014.
- [5] B. Bertenthal and C. Von Hofsten. Eye, head and trunk control: the foundation for manual development. Neuroscience & Biobehavioral Reviews, 22(4):515–520, 1998.
- [6] D. Bischof-Köhler. Self object and interpersonal emotions. identification of own mirror image, empathy and prosocial behavior in the 2nd year of life. Zeitschrift fur psychologie mit zeitschrift fur angewandte psychologie, 202(4):349–377, 1993.
- [7] D. Bischof-Köhler. Empathy and self-recognition in phylogenetic and ontogenetic perspective. *Emotion Review*, 4(1):40–48, 2012.
- [8] S.-J. Blakemore, C. D. Frith, and D. M. Wolpert. Spatio-temporal prediction modulates the perception of self-produced stimuli. *Journal of cognitive neuroscience*, 11(5):551–559, 1999.

- C. A. Brownell and M. S. Carriger. Changes in cooperation and self-other differentiation during the second year. *Child development*, pages 1164–1174, 1990.
- [10] M. Cakmak, C. Chao, and A. Thomaz. Designing interactions for robot active learners. *IEEE Trans. Autonomous Mental Development*, 2(2), 2010.
- [11] J. Carifio and R. J. Perla. Ten common misunderstandings, misconceptions, persistent myths and urban legends about likert scales and likert response formats and their antidotes. *Journal of Social Sciences*, 3(3):106–116, 2007.
- [12] M. Carpenter, J. Call, and M. Tomasello. Twelve- and 18-month-olds copy actions in terms of goals. *Dev. Sci.*, 8(1):13–20, 2005.
- [13] C. Chao and A. Thomaz. Controlling social dynamics with a parametrized model of floor regulation. J. of Human Robot Interaction, 2013.
- [14] M. J.-Y. Chung, A. L. Friesen, D. Fox, A. N. Meltzoff, and R. P. Rao. A bayesian developmental approach to robotic goal-based imitation learning. *PloS one*, 10(11):e0141965, 2015.
- [15] A. S. Clair and M. Mataric. How robot verbal feedback can improve team performance in human-robot task collaborations. In *Proceedings of the Tenth Annual ACM/IEEE Intl. conf. on Human-Robot Interaction*, pages 213–220, 2015.
- [16] M. Cook. Gaze and mutual gaze in social encounters. Nonverbal communication: Readings with commentary, 2:77–86, 1979.
- [17] J. L. Copete, Y. Nagai, and M. Asada. Development of goal-directed gaze shift based on predictive learning. In 4th International Conference on Development and Learning and on Epigenetic Robotics, pages 351–356. IEEE, 2014.
- [18] C. Darwin. The descent of man. J. Chem. Inf. Model., 1971.
- [19] M. Davidov, C. Zahn-Waxler, R. Roth-Hanania, and A. Knafo. Concern for Others in the First Year of Life: Theory, Evidence, and Avenues for Research. *Child Dev. Perspect.*, 7(2):126–131, 2013.

- [20] F. B. M. de Waal. Putting the altruism back into altruism: the evolution of empathy. Annu. Rev. Psychol., 59:279–300, 2008.
- [21] J. Decety and M. Svetlova. Putting together phylogenetic and ontogenetic perspectives on empathy. Dev. Cogn. Neurosci., 2(1):1–24, 2012.
- [22] H. E. Den Ouden, P. Kok, and F. P. De Lange. How prediction errors shape perception, attention, and motivation. *Frontiers in psychology*, 3, 2012.
- [23] A. D. Dragan, S. Bauman, J. Forlizzi, and S. S. Srinivasa. Effects of robot motion on human-robot collaboration. In ACM/IEEE Intl. Conf. on Human-robot Interaction (HRI), 2015.
- [24] K. Dunfield, V. A. Kuhlmeier, L. OConnell, and E. Kelley. Examining the diversity of prosocial behavior: Helping, sharing, and comforting in infancy. *Infancy*, 16(3):227–247, 2011.
- [25] N. Eisenberg, R. A. Fabes, and T. L. Spinrad. Prosocial development. Handbook of child psychology, 1998.
- [26] N. J. Emery. The eyes have it: the neuroethology, function and evolution of social gaze. *Neuroscience & Biobehavioral Reviews*, 24(6):581–604, 2000.
- [27] E. Fehr and B. Rockenbach. Human altruism: economic, neural, and evolutionary perspectives. *Current opinion in neurobiology*, 14(6):784–790, 2004.
- [28] G. G. Fein. Pretend play in childhood: An integrative review. *Child development*, pages 1095–1118, 1981.
- [29] T. Fong, C. Thorpe, and C. Baur. Multi-robot remote driving with collaborative control. *Industrial Electronics, IEEE Transactions on*, 50(4), 2003.
- [30] J. H. Friedman and B. E. Popescu. Predictive learning via rule ensembles. The Annals of Applied Statistics, pages 916–954, 2008.
- [31] M. Fujita. Aibo: Toward the era of digital creatures. The International Journal of Robotics Research, 20(10):781–794, 2001.

- [32] M. C. Gombolay, R. A. Gutierrez, G. F. Sturla, and J. A. Shah. Decisionmaking authority, team efficiency and human worker satisfaction in mixed human-robot teams. *Robots: Science and Systems (RSS)*, 2014.
- [33] J. Gottlieb, P.-Y. Oudeyer, M. Lopes, and A. Baranes. Information-seeking, curiosity, and attention: computational and neural mechanisms. *Trends in cognitive sciences*, 17(11):585–593, 2013.
- [34] J. Gray, C. Breazeal, M. Berlin, A. Brooks, and J. Lieberman. Action parsing and goal inference using self as simulator. In *Robot and Human Interac*tive Communication, 2005. ROMAN 2005. IEEE International Workshop on, pages 202–209. IEEE, 2005.
- [35] D. B. Grimes, R. Chalodhorn, and R. P. Rao. Dynamic imitation in a humanoid robot through nonparametric probabilistic inference. In *Robotics: science and systems*, pages 199–206. Citeseer, 2006.
- [36] R. Groten, D. Feth, A. Peer, and M. Buss. Shared decision making in a collaborative task with reciprocal haptic feedback-an efficiency-analysis. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 1834–1839, 2010.
- [37] K. P. Hawkins, S. Bansal, N. N. Vo, and A. F. Bobick. Anticipating human actions for collaboration in the presence of task and sensor uncertainty. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2014.
- [38] B. Hayes and B. Scassellati. Effective robot teammate behaviors for supporting sequential manipulation tasks. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2013.
- [39] G. Hoffman. Evaluating fluency in human-robot collaboration. In HRI Workshop on Human Robot Collaboration, 2013.
- [40] G. Hoffman and C. Breazeal. Effects of anticipatory action on human-robot teamwork efficiency, fluency, and perception of team. In ACM/IEEE Intl. Conf. on Human-robot Interaction (HRI), 2007.
- [41] G. Hoffman and C. Breazeal. Effects of anticipatory perceptual simulation on practiced human-robot tasks. Autonomous Robots, 28(4):403–423, 2010.

- [42] M. L. Hoffman. Is altruism part of human nature? Journal of Personality and social Psychology, 40(1):121, 1981.
- [43] M. L. Hoffman. Interaction of affect and cognition in empathy. *Emotions*, cognition, and behavior, pages 103–131, 1984.
- [44] M. L. Hoffman. Empathy and prosocial behavior. Handbook of emotions, 3:440–455, 2008.
- [45] C. Huang and B. Mutlu. Learning-based modeling of multimodal behaviors for humanlike robots. In ACM/IEEE Intl. Conf. on Human-robot Interaction (HRI), 2014.
- [46] C.-M. Huang, M. Cakmak, and B. Mutlu. Adaptive coordination strategies for human-robot handovers. In *Robotics: Science and Systems (RSS)*, 2015.
- [47] S. Ivaldi, S. M. Anzalone, W. Rousseau, O. Sigaud, and M. Chetouani. Robot initiative in a team learning task increases the rhythm of interaction but not the perceived engagement. *Frontiers in neurorobotics*, 8, 2014.
- [48] N. Jarrassé, J. Paik, V. Pasqui, and G. Morel. How can human motion prediction increase transparency? In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 2134–2139, 2008.
- [49] J. Kärtner, H. Keller, and N. Chaudhary. Cognitive and social influences on early prosocial behavior in two sociocultural contexts. *Developmental Psychology*, 46(4):905, 2010.
- [50] B. Kenward and G. Gredebäck. Infants help a non-human agent. PloS one, 8(9):e75130, Jan. 2013.
- [51] H. Li, J.-J. Cabibihan, and Y. K. Tan. Towards an effective design of social robots. *International Journal of Social Robotics*, 3(4):333–335, 2011.
- [52] T. Liu, J. Wang, and M. Q.-H. Meng. Human robot cooperation based on human intention inference. In *Robotics and Biomimetics (ROBIO)*, 2014 *IEEE International Conference on*, pages 350–355. IEEE, 2014.
- [53] J. Mainprice and D. Berenson. Human-robot collaborative manipulation planning using early prediction of human motion. In *IEEE/RSJ Intl. Conf.* on Intelligent Robots and Systems (IROS), pages 299–306, 2013.

- [54] J. Mainprice, E. A. Sisbot, L. Jaillet, J. Cortés, R. Alami, and T. Siméon. Planning human-aware motions using a sampling-based costmap planner. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2011.
- [55] A. N. Meltzoff. Imitation and other minds: The like me hypothesis. Perspectives on imitation: From neuroscience to social science, 2:55–77, 2005.
- [56] A. N. Meltzoff. 'Like me': A foundation for social cognition. Developmental Science, 10(1):126–134, 2007.
- [57] H. Moll and M. Tomasello. Level 1 perspective-taking at 24 months of age. Br. J. Dev. Psychol., 24(3):603–613, 2006.
- [58] A. Moon, D. M. Troniak, B. Gleeson, M. K. Pan, M. Zheng, B. A. Blumer, K. MacLean, and E. A. Croft. Meet me where i'm gazing: how shared attention gaze affects human-robot handover timing. In ACM/IEEE Intl. Conf. on Human-robot Interaction (HRI), 2014.
- [59] B. Mutlu, A. Terrell, and C.-M. Huang. Coordination mechanisms in human-robot collaboration. In *HRI Workshop on Collaborative Manipulation*, 2013.
- [60] M. Myowa-Yamakoshi, C. Scola, and S. Hirata. Humans and chimpanzees attend differently to goal-directed actions. *Nature communications*, 3:693, 2012.
- [61] Y. Nagai and M. Asada. Predictive learning of sensorimotor information as a key for cognitive development. Proc. of the IROS 2015 Workshop on Sensorimotor Contingencies for Robotics, 2015.
- [62] N. Najmaei and M. R. Kermani. Prediction-based reactive control strategy for human-robot interactions. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 3434–3439, 2010.
- [63] S. Nikolaidis, R. Ramakrishnan, K. Gu, and J. Shah. Efficient model learning from joint-action demonstrations for human-robot collaborative tasks. In ACM/IEEE Intl. Conf. on Human-robot Interaction (HRI), 2015.
- [64] S. Nikolaidis and J. Shah. Human-robot cross-training: Computational formulation, modeling and evaluation of a human team training strategy. In ACM/IEEE Intl. Conf. on Human-robot Interaction (HRI), 2013.

- [65] P.-Y. Oudeyer. Intelligent adaptive curiosity: a source of self-development. in proceedings of the Fourth International Workshop on Epigenetic Robotics, 2004.
- [66] H. Over and M. Carpenter. Eighteen-month-old infants show increased helping following priming with affiliation: Research report. *Psychol. Sci.*, 20(10):1189–1193, 2009.
- [67] J. C. Park, D. S. Kim, and Y. Nagai. Developmental dynamics of RNNPB: New insight about infant action development. Lect. Notes Comput. Sci. (including Subser. Lect. Notes Artif. Intell. Lect. Notes Bioinformatics), 8575 LNAI:144–153, 2014.
- [68] U. Pattacini, F. Nori, L. Natale, G. Metta, and G. Sandini. An experimental evaluation of a novel minimum-jerk Cartesian controller for humanoid robots. *IEEE/RSJ 2010 Int. Conf. Intell. Robot. Syst. IROS 2010 - Conf. Proc.*, pages 1668–1674, 2010.
- [69] M. Paulus. The Emergence of Prosocial Behavior: Why Do Infants and Toddlers Help, Comfort, and Share? *Child Dev. Perspect.*, 8(2):77–81, 2014.
- [70] C. Perez-D'Arpino and J. Shah. Fast target prediction of human reaching motion for cooperative human-robot manipulation tasks using time series classification. In *IEEE Intl. Conf. on Robotics and Automation*, 2014.
- [71] J. Piaget. The theory of stages in cognitive development. 1971.
- [72] J. Piaget, M. Cook, and W. Norton. The origins of intelligence in children, volume 8. International Universities Press New York, 1952.
- [73] H. L. Rheingold. Little children's participation in the work of adults, a nascent prosocial behavior. *Child Development*, 53(1):114, 1982.
- [74] G. Rizzolatti and L. Craighero. The mirror-neuron system. Annu. Rev. Neurosci., 27:169–192, 2004.
- [75] P. Rochat. Object manipulation and exploration in 2-to 5-month-old infants. Developmental Psychology, 25(6):871, 1989.

- [76] A. Sagi and M. L. Hoffman. Empathic distress in the newborn. Developmental Psychology, 12(2):175, 1976.
- [77] K. Sakita, K. Ogawara, S. Murakami, K. Kawamura, and K. Ikeuchi. Flexible cooperation between human and robot by interpreting human intention from gaze information. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, volume 1, pages 846–851. IEEE, 2004.
- [78] J. Shah, J. Wiken, B. Williams, and C. Breazeal. Improved human-robot team performance using chaski, a human-inspired plan execution system. In ACM/IEEE Intl. Conf. on Human-robot Interaction (HRI), 2011.
- [79] E. A. Sisbot, A. Clodic, R. Alami, and M. Ransan. Supervision and motion planning for a mobile manipulator interacting with humans. In ACM/IEEE Intl. Conf. on Human-robot Interaction (HRI), 2008.
- [80] J. a. Sommerville and A. L. Woodward. Pulling out the intentional structure of action: The relation between action processing and action production in infancy. *Cognition*, 95(1):1–30, 2005.
- [81] J. A. Sommerville, A. L. Woodward, and A. Needham. Action experience alters 3-month-old infants' perception of others' actions. *Cognition*, 96(1):B1–B11, 2005.
- [82] J. R. Stevens and M. D. Hauser. Why be nice? psychological constraints on the evolution of cooperation. *Trends in cognitive sciences*, 8(2):60–65, 2004.
- [83] C. E. Taylor, A. J. Parker, S. F. Lau, E. C. Blair, A. Heninger, and E. Ng. Robot vacuum with internal mapping system, Sept. 28 2010. US Patent 7,805,220.
- [84] M. Tomasello, M. Carpenter, J. Call, T. Behne, and H. Moll. Understanding and sharing intentions: The origins of cultural cognition. *Behavioral and brain sciences*, 28(05):675–691, 2005.
- [85] E. Ugur, Y. Nagai, E. Sahin, and E. Oztop. Staged development of robot skills: Behavior formation, affordance learning and imitation with motionese. *IEEE Transactions on Autonomous Mental Development*, 7(2):119–139, 2015.

- [86] E. Ugur, E. Sahin, and E. Oztop. Affordance learning from range data for multi-step planning. In Int. Conf. on Epigenetic Robotics, 2009.
- [87] D. Verma and R. P. Rao. Goal-based imitation as probabilistic inference over graphical models. In Advances in neural information processing systems, pages 1393–1400, 2005.
- [88] F. Warneken, F. Chen, and M. Tomasello. Cooperative activities in young children and chimpanzees. *Child Dev.*, 77(3):640–63, 2006.
- [89] F. Warneken and M. Tomasello. Altruistic helping in human infants and young chimpanzees. *science*, 311(5765):1301–1303, 2006.
- [90] F. Warneken and M. Tomasello. Helping and cooperation at 14 months of age. *Infancy*, 11(3):271–294, 2007.
- [91] F. Warneken and M. Tomasello. The roots of human altruism. British journal of psychology (London, England : 1953), 100:455–471, 2009.
- [92] D. M. Wolpert, Z. Ghahramani, and M. I. Jordan. An internal model for sensorimotor integration. *Science*, 269(5232):1880, 1995.
- [93] A. L. Woodward. Infants selectively encode the goal object of an actor's reach. *Cognition*, 69(1):1–34, 1998.
- [94] A. L. Woodward. Infants' grasp of others' intentions. Curr. Dir. Psychol. Sci., 18(1):53–57, 2009.
- [95] C. Zahn-Waxler, M. Radke-Yarrow, E. Wagner, and M. Chapman. Development of concern for others. *Developmental psychology*, 28(1):126, 1992.

Chapter 7

Appendix

.1 Appendix A

In the following pages, we attached the logging questionnaire used by the participants of our user study during their interaction with the robot. As each participant performed a total of 6 tasks with the robot, the questionnaire is divided in 6 sections from "Task #1" to "Task #6". For each section and step, the user indicated who performed the action, with which object (object ID) and to what location (location ID).

Task Log *Required
1. Participant ID *
Practice task
2. STEP 1 Performed by: <i>Mark only one oval.</i>
RobotHuman
3. Object ID:
4. Target ID:
5. STEP 2 Performed by: <i>Mark only one oval.</i>
 Robot Human
6. Object ID:
7. Target ID:

Task #1

8.	STEP 1 Performed by: <i>Mark only one oval.</i>
	Robot
	Human
9.	Object ID:
10.	Target ID:
11.	STEP 2 Performed by: Mark only one oval.
	Human
12.	Object ID:
13.	Target ID:
14.	STEP 3 Performed by: <i>Mark only one oval.</i>
	C Robot
	Human
15.	Object ID:
16.	Target ID:
17.	STEP 4 Performed by: <i>Mark only one oval.</i>
	C Robot
	Human

18.	Object ID:
19.	Target ID:
Та	sk #2
20.	STEP 1 Performed by: <i>Mark only one oval.</i>
	RobotHuman
21.	Object ID:
22.	Target ID:
23.	STEP 2 Performed by: Mark only one oval.
	RobotHuman
24.	Object ID:
25	Targat ID:
25.	
26.	STEP 3 Performed by: <i>Mark only one oval.</i>
	RobotHuman
27.	Object ID:

28. Target ID:
29. STEP 4 Performed by: <i>Mark only one oval.</i>
RobotHuman
30. Object ID:
31. Target ID:
32. STEP 5 Performed by: <i>Mark only one oval.</i>
C Robot
Human
33. Object ID:
34. Target ID:
25 STED 6
Performed by: Mark only one oval.
Robot
Human
36. Object ID:
37. Target ID:
Task #3
38.

39.
40.
41.
42.
43.
44.
45.
46.
47.

48.	Object ID:
49.	Target ID:
Та	sk #4
50.	STEP 1 Performed by: <i>Mark only one oval.</i>
	RobotHuman
51.	Object ID:
52.	Target ID:
53.	STEP 2 Performed by: <i>Mark only one oval.</i> Robot Human
54.	Object ID:
55.	Target ID:
56.	STEP 3 Performed by: Mark only one oval. Robot Human
57.	Object ID:

58	Target ID:
59. \$ 	STEP 4 Performed by: Mark only one oval.
	RobotHuman
60. (Object ID:
61. ⁻	Target ID:
62. 9 	Step 5 Performed by: Mark only one oval.
	RobotHuman
63. (Object ID:
64. ⁻	Target ID:
65. \$ <i> </i>	STEP 6 Performed by: Mark only one oval. Robot Human
66. (Object ID:
67. ⁻	Target ID:
Tas	sk #5

68.	STEP 1 Performed by: Mark only one oval. Robot Human
69.	Object ID:
70.	Target ID:
71.	STEP 2 Performed by: Mark only one oval. Robot Human
72.	Object ID:
73.	Target ID:
74.	STEP 3 Performed by: Mark only one oval. Robot Human
75.	Object ID:
76.	Target ID:
77.	STEP 4 Performed by: Mark only one oval. Robot Human

78.	Object ID:
79.	Target ID:
Та	sk #6
80.	STEP 1 Performed by: <i>Mark only one oval.</i>
	Robot
	Human
81.	Object ID:
82.	Target ID:
83.	STEP 2 Performed by: <i>Mark only one oval.</i>
	Robot
	Human
84.	Object ID:
85.	Target ID:
86.	STEP 3 Performed by: <i>Mark only one oval.</i>
	Robot
	Human
87.	Object ID:

88. Target ID:
89. STEP 4 Performed by: <i>Mark only one oval.</i>
Human
90. Object ID:
91. Target ID:
92. STEP 5 Performed by: <i>Mark only one oval.</i>
RobotHuman
93. Object ID:
94. Target ID:
95. STEP 6 Performed by: <i>Mark only one oval.</i> Robot Human
96. Object ID:
97. Target ID:

Powered by

.2 Appendix B

In the following pages, we attached the subjective rating questionnaire used by the participants of our user study between each condition and at the end of their interaction with the robot. Between each condition, the participants answered a series of open ended questions about their interaction. They then rated the robot's performance using a mean-Likert rating. At the end of their interaction, the participants rated the 3 conditions one to another and answered few more questions.

Questionnaire for Human-Robot Collaboration Study

*Required

1. [Answered by the experimenter] Participant ID: * Number between 1 and 18

Part 1

Answer the following questions after completing the *first two tasks* with the robot.

2. [Answered by the experimenter] Experimental condition in Part 1: * Mark only one oval.



3. Based on your observation in the two tasks, briefly describe Rosie's strategy in assisting the tasks. *



4. Briefly describe your strategy in completing the tasks. *



10. Rosie and I worked efficiently together. *

On a scale from 1 to 7 please indicate your level of agreement with this statement. *Mark only one oval.*



11. The collaboration felt natural. *

On a scale from 1 to 7 please indicate your level of agreement with this statement. *Mark only one oval.*



Part 2

Answer the following questions after completing the *next two tasks* with the robot.

12. [Answered by the experimenter] Experimental condition in Part 2: *

\bigcirc	нн
\bigcirc	RRH
\bigcirc	RPH

Mark only one oval.

13. Based on your observation in the two tasks, briefly describe Rosie's strategy in assisting the tasks. *

.....

14. Briefly describe your strategy in completing the tasks. *

3 of 8



20. Rosie and I worked efficiently together. *

On a scale from 1 to 7 please indicate your level of agreement with this statement. *Mark only one oval.*



21. The collaboration felt natural. *

On a scale from 1 to 7 please indicate your level of agreement with this statement. *Mark only one oval.*



Part 3

Answer the following questions after completing the *next two tasks* with the robot.

22. [Answered by the experimenter] Experimental condition in Part 3: *

\bigcirc	нн
\bigcirc	RRH
\bigcirc	RPH

Mark only one oval.

23. Based on your observation in the two tasks, briefly describe Rosie's strategy in assisting the tasks. *

24. Briefly describe your strategy in completing the tasks. *



30. Rosie and I worked efficiently together. *

On a scale from 1 to 7 please indicate your level of agreement with this statement. *Mark only one oval.*



31. The collaboration felt natural. *

On a scale from 1 to 7 please indicate your level of agreement with this statement. *Mark only one oval.*



Additional questions

Please answer the following questions based on all three parts of the study.

32. Please rank the robot behaviors in the three parts of the study, in terms of your preference as a robot assistant. *

Mark only one oval per row.

Least preferred Second preferred Most preferred

Part 1	\bigcirc	\bigcirc	\bigcirc
Part 2	\bigcirc	\bigcirc	\bigcirc
Part 3		\bigcirc	

33. Do you prefer having control over when the robot helps you or letting the robot decide when it should help? *

Mark only one oval.

I prefer having control over when the robot helps

) I prefer the robot deciding when it should help

34. In two of the three parts of the study, the robot decided when it should help. Please describe the difference between the two based on your experience and indicate which one you preferred. *

Please refer to robot behaviors as Part 1, Part 2 or Part 3.

35.	How would you combine the different robot behaviors in an ideal robot assistant? *
-	
De	emographic Information
36.	Age: *
37.	. Gender: *
	Mark only one oval.
	Female
	Male
	Other

