Sergey Y. Yurish, Editor

IFSA

Advances in Robotics and Automatic Control

4



Advances in Robotics and Automatic Control

Book Series, Volume 4

S. Yurish
Editor

Advances in Robotics and Automatic Control

Book Series, Volume 4

IFSA International Frequency Sensor Association Publishing

S. Yurish, *Editor*
Advances in Robotics and Automatic Control, Book Series, Vol. 4

Published by IFSA Publishing, S. L., 2026

E-mail (for print book orders and customer service enquires):

ifsa.books@sensorsportal.com

Visit our Home Page on http://www.sensorsportal.com

Advances in Robotics and Automatic Control: Reviews, Vol. 4 is an open access book, which means that all content is freely available without charge to the user or his/her institution. Users are allowed to read, download, copy, distribute, print, search, or link to the full texts of the articles, or use them for any other lawful purpose, without asking prior permission from the publisher or the authors. This is in accordance with the BOAI definition of open access.

Neither the authors no IFSA Publishing, S.L. accept any responsibility or liability for loss or damage occasioned to any person or property through using the material, instructions, methods or ideas contained herein, or acting or refraining from acting as a result of such use.

ISBN: 978-84-09-95166-6
BN-20161119-XX
BIC: TJFM1

Contents

Preface **7**
Contributors **9**

1. Digitalization of Industry: How AGVs and COBOT Transform the Workplace and Bring New Challenges to OHS **11**
1.1. Introduction 11
1.2. Methodology 12
1.2.1. Bibliographic Review *12*
1.2.2. Bibliographic Search Criteria *13*
1.3. Results 13
1.3.1. AGVs versus Cobots *13*
1.3.2. Opportunities Associated with AGVs and Cobots *15*
1.3.3. Challenges Raised by AGVs and Cobots *16*
1.3.3.1. Risk of Interference and Proposed Solutions *16*
1.3.3.2. Risks Related to Movements and Proposed Solutions *23*
1.3.3.3. Data Management Risks and Proposed Solutions *24*
1.3.3.4. Ergonomic Risks and Proposed Solutions *25*
1.3.3.5. Psychosocial Risks and Proposed Solutions *26*
1.4. Discussion 27
1.5. Conclusions 29
Acknowledgments 30
Abbreviations 30
References 30

2. Sociability Modelling in Robot Motion **35**
2.1. Introduction 35
2.2 Related Work and Methodology 38
2.2.1. Notions of Predictability, Legibility, Trustability and Sociability *38*
2.2.1.1. Predictability *38*
2.2.1.2. Legibility *39*
2.2.1.3. Trustability *40*
2.2.1.4. Sociability *41*
2.2.2. Predictable, Legible, and Sociable Motion Planning *41*
2.2.2.1. Gradient Descent Methods *42*
2.2.2.2. Methods for Improving Predictability *43*
2.2.2.3. Methods for Improving Legibility *43*
2.2.2.4. Socially Aware Navigation *43*
2.3. Sociability Modelling and Quantification 44
2.3.1. Predictability *45*
2.3.2. Predictability Prime (Predictability') *45*
2.3.3. Legibility *46*
2.3.4. Trustability *48*
2.3.5. Sociability *51*
2.3.6. Modelling Limitations *52*

2.4. Generation of Acceptable Trajectories 53
2.4.1. Planners Analysis *55*
2.4.1.1 Determining the Best Planners *56*
2.4.1.2. Determining the Threshold Values *56*
2.4.1.3. Time Efficiency of Planners *57*
2.4.2. Defining the Different Types of Trajectories *57*
2.4.3. Execution of Trajectories *58*
2.4.3.1. (Highly) Predictable Trajectory *58*
2.4.3.2. Trustable Trajectory *59*
2.4.3.3. Sociable Trajectory *59*
2.4.4. Applying the Trust Region of Predictability – Trustably and Socially Predictable Trajectories *60*
2.4.4.1. Generation of Trajectories *61*
2.4.5 Limitations *62*
2.5. Experimental Design and Creation of Social Intelligence 63
2.5.1. Human-Dependent Experiments (User Study) *64*
2.5.2. Human-Independent Experiments *66*
2.5.2.1. Trustability *67*
2.5.2.2. Sociability *68*
2.5.3. Creation of Social Intelligence – Training a Classifier to Recognise Socially Predictable Trajectories *70*
2.5.3.1. Dataset *70*
2.5.3.2. Model Architecture *71*
2.5.3.3. Model Training *71*
2.6. Results and Discussions 72
2.6.1. Human-Dependent Experiments (User Study) *72*
2.6.1.1. Experiments with Live Demonstrations *72*
2.6.1.2. Experiments with Pre-Recorded Demonstrations *74*
2.6.2. Human-Independent Experiments *78*
2.6.2.1. Trustability *78*
2.6.2.2. Sociability *81*
2.6.3. Creation of Social Intelligence – Training a Classifier to Recognise Socially Predictable Trajectories *83*
2.6.3.1. Deep Learning-Based Model Training *83*
2.6.3.2. Evaluation *83*
2.7. Future Work 85
2.7.1. Advantages of Dreamer *85*
2.7.2. Modified Dreamer Architecture *86*
2.8. Conclusions 89
Acknowledgements 90
References 90

Preface

Robotics and automatic control are reshaping modern industry and society. Autonomous machines are increasingly deployed in workplaces and public environments where they must operate alongside people. Their success depends not only on efficiency, precision, and autonomy, but also on safety, reliability, and the ability to behave in ways humans can understand and trust. Addressing these human-centred requirements is therefore essential to the responsible development of robotic technologies.

Autonomous systems now operate in spaces that were once exclusively human, placing new demands on their design. Safe collaboration requires more than collision avoidance: it involves dependable sensing and control, understandable behaviour, and attention to people's physical and psychological well-being. These considerations increasingly connect robotics research with ergonomics, occupational safety, and human factors.

The 4th volume of *Advances in Robotics and Automatic Control* presents two complementary studies of the opportunities and challenges created by closer human–robot cooperation. Although the chapters differ in scope and methodology, both consider how robotic systems can be integrated into human environments without compromising safety or confidence.

Chapter 1, '*Digitalization of Industry: How AGVs and COBOT Transform the Workplace and Bring New Challenges to OHS*' focuses on automated guided vehicles (AGVs) and collaborative robots (cobots) in Industry 4.0. Drawing on a review of the scientific literature, it examines productivity and operational benefits alongside emerging occupational health and safety risks. Particular attention is given to collisions, movement-related hazards, data management, ergonomic concerns, and psychosocial effects. The chapter also considers risk assessment, preventive measures, employee training, and the importance of involving workers in technological change.

Chapter 2, '*Sociability Modelling in Robot Motion*', explores how robots can communicate their intentions through movement and behave in socially acceptable ways. It investigates predictability, legibility, trustability, and sociability, and proposes methods for quantifying these

properties and generating appropriate robot trajectories. User studies and simulation-based experiments evaluate the proposed approach, while a deep-learning classifier is trained to recognise socially predictable motion. The results provide a basis for further work towards more socially intelligent robotic systems.

The chapters approach this transition from two complementary perspectives. The first investigates the conditions under which robotic technologies enter workplaces; the second considers how robots can make their actions intelligible through motion. This distinction matters because technically safe behaviour may still be difficult for people to anticipate or trust.

Together, these chapters emphasise that advances in robotics cannot be evaluated solely through technical performance. Human safety, comfort, perception, and trust are also crucial to the successful adoption of automated systems across industrial and everyday settings.

Future progress will require interdisciplinary cooperation and evaluation in realistic environments. Advances in sensing, intelligent control, and motion planning must be considered alongside human expectations and working practices. Addressing these factors from the earliest design stages can support robotic systems that are not only capable and efficient, but also dependable and genuinely human-centred.

The book is intended for researchers, engineers, and developers in robotics, automatic control, industrial automation, intelligent systems, and human–robot interaction. It should also be useful to graduate students, university teachers, occupational health and safety specialists, and professionals responsible for implementing robotic solutions.

This volume has been written by seven authors from Canada and the United Kingdom. I am grateful to them for contributing their research and expertise. I hope the book will encourage further interdisciplinary research and support the development of robotic systems that are both technologically advanced and attentive to human needs

Sergey Y. Yurish

Editor
IFSA Publishing *Barcelona, Spain*

Contributors

Adel BADRI
Université du Québec à Trois-Rivières (UQTR), Canada

Mohamed Naceur BEN AZIZA
Université du Québec à Trois-Rivières (UQTR), Canada

Hajer JEMAI
Université du Québec à Trois-Rivières (UQTR), Canada

Gaoyong LUO
Department of Computing, Imperial College London,
London SW7 2AZ, UK

Haiwei LUO
Department of Computing, Imperial College London,
London SW7 2AZ, UK

Yaqing LUO
Department of Computing, Imperial College London,
London SW7 2AZ, UK

Jin ZHANG
Department of Computing, Imperial College London,
London SW7 2AZ, UK

1.

Digitalization of Industry: How AGVs and COBOT Transform the Workplace and Bring New Challenges to OHS

Mohamed Naceur Ben Aziza, Adel Badri and Hajer Jemai

1.1. Introduction

As the modernization of industry advances, digitization, automation, information/communication technologies and increasingly robotics become major business assets. The fourth industrial revolution, often called Industry 4.0, is now the predominant reality (Fig. 1.1).

Robots, cobots and AGVs have been iconic symbols of technological advancement for decades, often equated with artificial intelligence and even development of human intelligence. Since the fourth wave of the Industrial Revolution was recognized in 2011, these technologies have come to the forefront and are being integrating more and more into the workplace, especially in industrial settings.

This is occurring for various reasons, mainly the drive to reduce production costs and increase factory productivity and product quality, but also because of workplace safety concerns.

AGVs and cobots have brought obvious improvements to industrial productivity, autonomy and product quality [1]. These devices are used to carry out tasks that are repetitive or dangerous or require high precision, and they can operate under harsh or extreme conditions and collaborate with workers to improve overall performance.

Mohamed Naceur Ben Aziza
Université du Québec à Trois-Rivières (UQTR), Canada

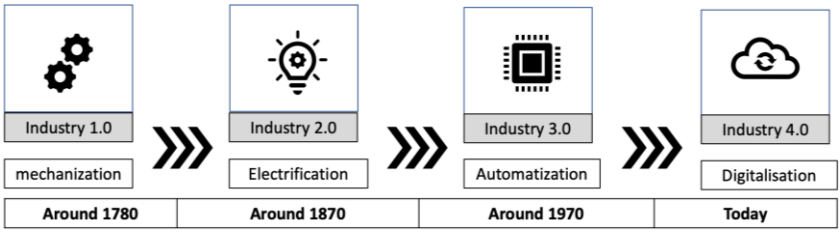


Fig. 1.1. Industrial revolutions over the centuries.

However, growing use of AGVs is posing multiple challenges to worker health and safety [2]. The rapid and sometimes radical changes in the workplace do not necessarily bring a concomitant adaptation of managers and workers to the new reality. On the contrary, resistance to change seems to typify the behavior of many actors in industrial settings and poses a barrier to be overcome to ensure a smooth transition towards Industry 4.0.

For OHS purposes, workplace adaptation to change requires adjustment and updating of safety and preventive measures, as well as modification of procedures, managerial methods, and human behavior. In the absence of this exercise, workers' health and safety suffer a decline.

In the present chapter, we review studies focused on AGVs and cobots regarding OHS. We examine the security measures in use or proposed and summarize the challenges and the opportunities for improvement. We conclude the chapter with observations and a few recommendations.

1.2. Methodology

1.2.1. Bibliographic Review

We queried the broad multidisciplinary database Scopus using keywords relating to collaborative robotics and autonomous guided vehicles, with the search limited to titles and abstracts. The selection includes articles published during the 2017–2022 period. We chose to focus on recent articles because the research topic is current, the number of articles published in this field having rebounded over this period.

1.2.2. Bibliographic Search Criteria

To ensure retrieval of relevant and significant documents, we conducted our search using the English-language terms most used in the literature as synonyms of autonomous guided vehicles, collaborative robots, and occupational health and safety, including their acronyms: (TITLE-ABS-KEY (("collaborative robot" OR "cobot") AND ("OHS" OR "safety" OR "security" OR "health")) OR TITLE-ABS-KEY (("autonomous intelligent vehicle" OR "AIV" OR "autonomous guided vehicle" OR "AGV") AND ("OHS" OR "safety" OR "security" OR "health"))) AND PUBYEAR> 2017 AND (LIMIT-TO (LANGUAGE, "English")) AND (LIMIT-TO (SUBJAREA, "ENGI"))

Among the publications thus identified, we selected 40 based on relevance to our research interest, especially regarding opportunities and challenges associated with aspects of OHS according to Prisma methodology (Fig. 1.2).

To obtain the most up-to-date results possible, we decided to focus on the most recent publications and therefore concentrated on those published between 2017 and 2022. The first selection was then made by choosing publications in engineering, in line with the main research area of this study. Only scientific articles written in English were selected.

Next, articles dealing with cybersecurity related to the use of cobots and AGV were excluded, since this was not the focus of our research. We wanted to concentrate much more on occupational health and safety aspects. After reviewing the remaining articles, we kept only the most relevant ones.

1.3. Results

1.3.1. AGVs versus Cobots

As their name suggests, AGVs have electronics that enable them to function without a human driver. Their use has grown rapidly over recent decades, especially in the field of materials transport and logistics and in highly modern work environments and production systems. AGVs are perhaps the most iconic representatives of the worldwide shift to Industry 4.0 and integration of new technologies to improve

productivity, reduce costs and meet the demands of sustainable development (OHS, ecological, etc.). The question of safety is front and center in discussions of the transition towards Industry 4.0. The introduction of new technologies into the work environment inevitably brings emergent OHS risks [1].

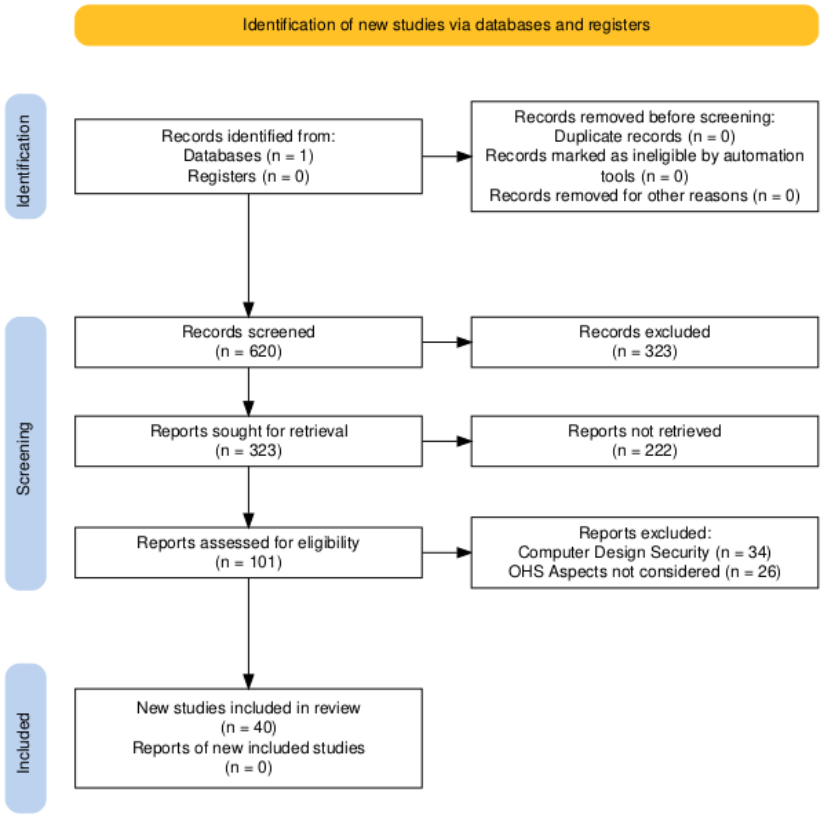


Fig. 1.2. Prisma methodology flowchart.

In contrast, the use of robots in industrial settings is hardly recent. The first productive electric robot was introduced in the 1970s, and robots began to appear on automobile assembly lines soon after. However, since 2010, a paradigm shift has taken place, a so-called fourth wave of the Industrial Revolution, in which more and more robots function in partial or total collaboration with human workers. Not surprisingly, such collaboration raises concerns, relating mainly to safety. Safeguards

based on unambiguous separation between spaces where humans move and those where robots function, primarily fixed barriers, luminous curtains, or motion detectors, are no longer sufficient. Collaboration requires workers and robots to occupy the same work zone at least partially. Under such conditions, a very different philosophy must be applied to safety management. Published studies of the integration of cobots focus mainly on optimization, efficiency and economics, in some cases casting doubt on the legitimacy of some safety concerns [3]. An overview of possible negative effects of human-cobot interactions and the effectiveness of protective or corrective measures has been published [4].

Robotization does not eliminate the necessity for close human presence. No robotic system is fully autonomous. For example, robots remain incapable of timely and adequate response to unforeseeable situations. Claims that humans will be entirely safe in the presence of cobots are unfounded, given that the relevant safety standards are very new or still in development. Numerous accidents in robotized work areas occur during non-routine operations such as programming, maintenance, testing, configuration, or personalization. Cobots need to be designed intrinsically for safety, and these designs as well as any new functions need to be tested rigorously.

Emerging needs and possible trends in safety testing procedures for cobotic applications have been considered in an overview of the relevant standards to be validated, with a description of a new interdisciplinary approach based on safety testing skills and protocols [5]. Although robots are replacing humans for tasks that are tedious, unsanitary, dangerous, and/or delicate, their generalized use currently runs into difficulties in validating the safety of specific applications. The human-cobot interface can be improved by integrating various technological innovations into the cobots, such as vision, voice recognition, augmented reality glasses and even bio-signal detection.

1.3.2. Opportunities Associated with AGVs and Cobots

Integration of AGVs and cobots into an industrial setting can allow major savings in time and energy. Replacing manually operated or semi-automated vehicles with complementary autonomous transport is usually more efficient, especially when several vehicles must operate simultaneously. Having cobots perform routine tasks that are repetitive

and/or tedious or take place under harsh environmental conditions also improves workers physical and mental health in addition to productivity and efficiency [2]. The economic benefits come largely from the ability of cobots to complete such tasks faster and more reliably than humans do [6].

The use of cobots in a workplace alongside humans represents several advantages, one of these being economic when human workers are thus replaced. A cobot can perform tirelessly the tasks of several workers during a shift, which is even more advantageous in productions involving multiple work crews. This can lead to considerable reductions of labor costs and hence increases in productivity as well as more consistent product quality. By taking over repetitive and demoralizing tasks, cobots decrease the overall ergonomic wear endured by humans and thus increase worker productivity. They are often regarded as playing an indispensable role in decreasing the exposure of humans to dangerous work conditions, for example, reducing the risk of accidents and injury associated with handling of heavy loads and thus allowing an overall reduction of work-related anxiety and stress [4]. Cobots are generally designed to stop before any involuntary contact potentially harmful to a human worker can occur. Some are equipped with additional sensors to detect gaseous or ambient pollution [7].

1.3.3. Challenges Raised by AGVs and Cobots

1.3.3.1. Risk of Interference and Proposed Solutions

From an OHS perspective, the risk of collision has received the most attention in the literature on human/machine interactions since this event is most likely to cause injury. The problem of AGV safety has been studied in terms of traffic control methods, on-vehicle devices, sensing technologies, human behavioral modification, and so on. Collisions between humans and cobots is a no less serious concern. In this case, the challenge is to establish a precise collision threshold suited to the workers and to output requirements [8].

Close and frequent contact with cobots may be unsafe for their users and lead directly to injury. Electronic fences or barriers or other industrial standards have provided passive safety, of which the limitations are known, such as low flexibility and safety in real time. To improve the

safety of human-cobot interactions, flexible tactile sensors have been designed to allow cobots to perceive contact pressure or force [9].

The goal of introducing cobots into an industrial workplace is, of course, to increase productivity, but this will not be achieved if restrictions associated with worker safety and wellbeing become obstacles. A cobot must not be introduced into a workplace without thorough preparation. A maturity model score sheet has been integrated into a tool for evaluating safety readiness for cobot implementation [8]. Cobot technologies represent to some extent a tradeoff between safety and efficiency, and lack of knowledge about safety during their introduction can create physical risks (collision), technological risks (cybersecurity) or psychosocial risks (work-related stress). In addition, the vagueness of OHS legislation complicates its translation into concrete measures [8].

The management of mobile workers and AGV traffic in a controlled industrial setting is achieved generally under two constraints: time and safety. Achieving harmony between mechanized material transport and human activity in the same space can be a daunting mission, especially when a plurality of vehicles is involved. Several collision avoidance techniques and strategies are currently in use, including detection of walking people and display of transport vehicle intentions. A technique of coordinating movement based on speed profiles and communicative behaviors as vehicles and workers pass each other has been described [10]. In experiments, this approach gave better avoidance of collisions than did the current techniques of collision avoidance, including the model predictive control approach.

In a study of the possible benefits of integrating AGVs into ports equipped with roll-on/roll-off terminals, where loading and unloading operations still depend on humans, the goal was to reduce human-error-associated risks and heighten safety [11]. Based on evaluations of the impact on productivity, profitability and the environment using the various simulators developed, automation of port operations using AGVs led to several gains, in addition to suggesting how to deal with the increasing labor shortage in this sector.

Entirely autonomous robotic transport of pallets in a factory using a master/slave pair of robots has been studied [12]. In this case the two robots work like the arms of a fork-lift, without any physical connection between them. To ensure free access to the space underneath the pallets,

a compact training unit was designed. The mobile robot is more compact, evolvable, and safe to use, making it very practical for factories with limited space. Each of the robots is equipped with safety modules and sensors.

The impact of the model representing the environment on AGV performance has been studied using a perfect circular robot [13]. A multi-objective cost function was defined to consider path length and complexity and the minimal distance around obstacles. Three types of representation, namely quadrilateral, irregular triangle, and variable-sized irregular triangle were used to model the environment while applying an inflation layer to the discretized view. Testing scenarios using different methods of cellular decomposition and a single inflation layer size show that a coarse and nearly constant triangular mesh is adequate for a single robot of unchanging dimensions in an unchanging environment. The tests also showed that variable triangular mesh representation sizes may be the best choice for factories with changing configurations and different sized robots, because of the inflation layer effect. In terms of safety and performance (e.g., calculation time), use of an irregular triangular mesh remains suitable for different types and complexities of environment (e.g., multi-obstacle, corridor, etc.). This also allows determination of the optimal tradeoff between AGV path length and complexity and the minimal distance relative to obstacles encountered on the way.

LiDAR (light detection and ranging) is a distance detection and estimation technology that is proving efficient for processing environmental complexities. Optimal working conditions such as constant lighting, paved roads, stable weather, and otherwise non-demanding surroundings can seldom be taken for granted, notably in agriculture. In a study of its applicability in irregular environments such as farms, LiDAR was found effective not only in navigation but also for avoidance of obstacles and collisions, and thus to maintain a high level of safety [14]. It can thus play an effective role in guiding driverless transport systems such as AGVs. While challenges remain in terms of accuracy, coherence, and precision, the results obtained under complex working conditions are promising. The quite reliable collision avoidance obtained indoors remains to be reproduced outdoors.

A qualitative and quantitative analysis of various configurations of sensors used in industry has led to a systematic method of inquiry into the degree of precision of cameras as position sensors for AGVs and the

trajectory monitoring problem [15]. The aim of this study was to optimize AGV energy efficiency and autonomy over longer distances while ensuring greater safety for people working in the same space.

A novel safety system for AGVs has been presented wherein all movement of the vehicle is stopped as soon as an obstacle is detected in its path [16]. Using ultra-broad-band technology, this system differs from other real-time localization approaches by not requiring installation of any materials on the factory infrastructure. Sensors on board the vehicle ensure all detections of objects or people. This simplifies the use of AGVs, especially in dynamic environments. The system has shown good precision in tests under real conditions, and positioning errors are distributed more uniformly around an obstacle.

The problem of detecting collisions with cobots has been addressed using a virtual force sensor based on neural networks [17]. The goal of implementing robotic cells is to reduce if not eliminate human presence in digitalized hazardous production zones. Cobots were designed initially to reproduce human operations such as pulling, pushing, pressing, and raising/lowering. This made them inherently threatening to humans because of their high power and speed of movement. To make them safe for collaboration with workers in the same space, they must be re-engineered. Refinements tested to date include reduced load capacities, lower speed of operation and limited impact power. These do not eliminate the risk of collision. To achieve this, cobot design must include failsafe means of immediate reduction of speed as soon as an intrusion into the work zone is detected, in addition to constant monitoring of the intersection of cobot and human trajectories. The authors mention that cobots meeting ISO safety standards 10218 and ISO/TS 15066 should be able to interact safely with workers, although timely detection of cobotic threats to humans is generally quite expensive.

During the past decade, a novel apparatus called V2SOM (variable stiffness safety-oriented mechanism) was developed to ensure the safety of human-cobot physical interactions in unfamiliar environments [6]. These authors studied scenarios in which the cobot controls its power based on two safety criteria, namely impact force and the risk of head injury. In addition, a low-stiffness mode of operation was proposed to absorb safely any potential bumping into a human. These experiments serve as a reminder that from an OHS perspective, some tasks are

difficult to automate fully, human intervention being unavoidable in these cases.

The same authors later presented a simulation and experimental evaluation of their apparatus under normal working conditions in the case of uncontrolled collision between cobots and humans [18]. Their study of the problem of worker safety confronted with the dynamics of high-capacity cobots confirmed that human-robot collaboration can lead to serious injury in the event of a collision if the cobot does not possess the performance capabilities of humans. As soon as a collision is underway, a spontaneous gentle shift to a low-stiffness mode must be activated to absorb the impact safely.

In another way of improving the safety of cobots, visual perception of human gestures is combined with tactile perception of physical contact with the human [19]. The cobot is programmed to recognize and interpret, based on input from cameras and pressure sensors. Again, the risk of collisions remains, requiring new detection methods (e.g., observers of extended states, vibrational analysis, finite-time perturbation controllers, energy controllers…). Of course, the cobot is never truly conscious of the human position and intention, nor can it distinguish between intentional touching and accidental contact. To meet these challenges, a multi-objective system must be designed whereby the cobot "knows" and recognizes routine human actions including all intended touching and thus how to avoid collisions. These authors believe that sharing workspaces by humans and robots decreases productivity, unlike [6] who emphasize the importance of human presence.

Starting again from the assumption that cobotic is inherently risky, another study has focused on developing a framework for improving the reliability and safety of Industry 4.0 technologies [7]. Since OHS awareness is growing, cobotics is now subject to strict safety standards. Any cobot equipped with dangerous tools, moving quickly, or transporting large loads must be guided continuously by a properly trained human and kept at an appropriate distance from workers.

Integrating a cobot into an industrial setting raises several technical and managerial challenges, both productivity-related and OHS-related. One of these is getting safety evaluated during the feasibility study as well as in the application design phase. All restrictions imposed by safety standards must be respected, such as zero risk of contact with human

heads. Researchers have proposed a monitoring system that allows prediction of failure and evaluation of reliability factors using machine learning [20]. This could help the manufacturer of the robot by suggesting adaptations throughout the automation project to ensure safety at all stages. The manufacturer must also maintain and update all internal safety evaluation documents on a permanent basis. The equipment operators should be involved from the beginning of a collaborative robotics implementation project, to establish trust and adjust quickly to perturbations brought into the workplace with the new cobotic application.

Another concern has been the repeatability of cobotic test procedures. Conditions that may lead to unreliable results have been identified in an interlaboratory study, which also showed that heightening the risk of collisions increases the risk of biomechanical stress among workers [21]. In this case, a model that governs interaction to control the robot within safety rules is necessary. Safety standards such as described in ISO/TS 15066 cover modalities of interaction such as monitoring operation to allow safety shutoff, manual guiding to transmit commands to the robot, monitoring speed and separation distances, and limiting robot mechanical power. Nevertheless, identifying the safety requirements in specific fields of application remains difficult, since these standards do not describe these requirements exhaustively for the purposes of cobotics. Users must navigate through uncertainties as they interpret the applicable directives and standards. The authors of this study also cite the lack of practical trial procedures for evaluating compliance with requirements in a collaborative robotic system. These problems are addressed somewhat by providing cobots with active means of protection (tactile, sensors of torque, pressure, speed, reach, etc.) as well as passive means (adaptation of cobot shape, gripping device, cutting tools, of the piece to be processed, of any other apparatus involved in the work process, use of protective materials [e.g., skins] on the robot) and ensuring passive compliance at the design stage. To prevent collisions, measurement of the load at critical points must be provided to evaluate and validate the reduction of this risk. The user must repeat the test procedures whenever any change is brought to the conditions of application of a cobot.

A surface material having tactile sensitivity like human skin has been proposed to allow safer interaction between humans and cobots [22]. This solution allows at the very least a decrease in the impact of the

collision between the two and may thus make the system compliant with safety requirements. However, to reduce the risk of injury substantially, the cobot must be able to anticipate a collision. In other words, the skin of the robot also should detect rapid and large changes in contact force. According to these authors, human-cobot interaction can be improved by incorporating multiple detection functions and building in mechanisms of safety evaluation at several levels.

Experiments with the use of antenna-borne proximity sensors to stop the cobot as soon as a collision seems imminent suggest that avoidance of collisions improves as the number of sensors and their closeness (meters to centimeters) to the cobot is increased [23]. The safety of a multi-DoF cobot with variable joint rigidity has been evaluated for application to medical diagnostics [24]. For example, in Doppler sonography, the pressure applied to the patient by the handler of the sensing device must remain optimal. The authors conclude that this may be achieved by using such a system.

In view of the safety standards set forth in ISO 15066 regarding the use of industrial robotic arms in applications requiring human-machine interaction, [25] have proposed a semi-rigid cobot arm design that uses a virtual sensor of force plus stiffness control to allow detection of imminent collisions and timely responsiveness. They recommend integrating this aspect of safety beginning at the design phase. Cobot arms often include flexible mechanisms and sensors of force to detect and absorb impact from external sources. These may include safety features such as elastic actuators.

Most commercial cobots use the technique of collision detection with protective shutoff after the unwanted human-cobot contact has begun. An approach based on virtual sensors that calculates the peak collision force and pressure expected to result from a given human-cobot interaction have been developed using an analytical model of contacts [26]. By predicting the force and pressure before the collision has occurred, this method allows the guidance of the cobot in the management of its task, according to a collision safety threshold. Including pressure as a parameter implies taking contact surface area into account in the evaluation of the danger associated with the collision. By conducting collision tests under different undesirable scenarios, cobot practical safety can be examined. The proposed virtual sensor of collisions can be used jointly with the control system, not just to keep the

human away from the robot but also to prevent the robot from moving at unreduced speed towards a collision with the human.

A motion algorithm developed to allow a multi-robot system to avoid obstacles and maintain stable mobility uses self-adapting planning of trajectories [27]. Various collision avoidance algorithms have been tested in simulation with mobile obstacles [28]. In these dynamic environments, cobots equipped with large sets of sensors detect human movement and the position of objects inside a workspace representing risk to workers.

A review of published studies of issues raised by the integration of cobots into industrial settings concludes with a proposal to examine novel approaches based on virtual reality or more broadly on extended reality, as tools for testing the safety of collaborative robotics [29]. Although replacing humans with robots or cobots for high-risk tasks is inevitable and generally desirable, it will remain essential to monitor and even anticipate emergent physical and psychological risks that can and likely will arise from this development. The design of cobots can have an impact on the way humans perceive them, and hence on manual collaborations with them. The role that "understanding" of human actions by robots and vice versa will play in the acceptability of robotic actions to humans will soon come to light. In principle, collaboration between humans and cobots will be limited abruptly by aspects of worker safety and comfort, some of which will also be emergent.

1.3.3.2. Risks Related to Movements and Proposed Solutions

Stability has a major impact on the safety aspects of an AGV. Under variable centroid conditions, AGV stability becomes challenging. In a study of independent wheel drive and independent four-wheel drive, safety and stability problems due to drive actuator fault or poor efficiency were examined [30]. Models were then developed to support the design of an optimal controller that considers the conditions of centroid variability. In simulations, the models gave results very close to the theoretical objective.

The case of autonomous high-speed surface vehicles, a type of AGV being used more and more on the ground in numerous industrial and mining environments, has also been examined [31]. In this study, a

predictive model of AGV stabilization control at high speeds is presented, one that considers ground topography, primarily its curvature and angle of incline. Safe driving constraints such as zero-tilting moment point, zero moment point, and the lateral safety corridor are all considered to ensure stability and avoid collisions with obstacles, stationery or mobile.

1.3.3.3. Data Management Risks and Proposed Solutions

Data plays a preponderant role in Industry 4.0. Real-time management of data, gathering and analyzing massive data, safeguarding of data, and connection and safe interconnection of cobots and AGVs all may be exposed to breaches of integrity. Safety and reliability AGVs and cobots can depend on data quality, which can be compromised accidentally or by malevolent intention. The cybersecurity of human-machine collaborations is essential, and its maintenance faces constant threats [2].

A Cloud-based solution for reducing delays in processing raw data received by AGVs and drones has been proposed [32]. Such delays will obviously have an impact on safety since a problem detected but not solved in time can lead to major workplace accidents when several AGVs are involved. The proposed remote detection platform, which must be installed directly on AGV dashboards, consists essentially of a peripheral monitoring system. Gains of time and communication costs as well as high performance and coherence were noted in experiments with this system.

In any operation involving a human-cobot interaction, at least some real-time safety must be guaranteed. An ontological approach to solving industrial scenarios for safety applications in cyber-physical production systems has been presented [33]. Force or torque sensors and infrared/visible cameras are used to detect imminent collisions.

Without real-time detection of collisions, attempts to plan collision-free routes are fruitless. Among the real-time strategies proposed is ellipsoid modelling to ensure a minimal level of safety by changing speed [34]. The effectiveness of this approach has been validated in simulations and experiments. A danger index with ellipsoids is used to measure quantitatively the degree of safety to detect collisions rapidly. If a collision anticipated, the cobot slows down or waits, thus lowering the probability of injury.

An attempt to find a reliable and cost-effective real-time solution to be integrated into a cobot safety system has been described in which humans are detected, identified and localized as they approach the operational zone, and the information is relayed directly to the cobot control system [35]. One of the challenges facing developers of highly reliable safety systems is the requirement of obtaining real-time responsiveness in changing environments without loss of cobot productivity. To guarantee safe cooperation between humans and robots, it is generally necessary to detect humans, to estimate their position and to avoid obstacles when changing course. The authors cited above conclude that short reaction time, reliability, and high capacities for detection and identification are crucial requirements for a safety system appropriate for cobots.

1.3.3.4. Ergonomic Risks and Proposed Solutions

Studies of biomechanical interactions between humans and cobots show at least modest economic benefits of sensors and safety cages or fences, which may be cumbersome as well as expensive [36]. Although cobots offer a wide range of opportunities to improve the ergonomics and efficiency of manual labor, cooperative human-robot interactions themselves create risks of musculoskeletal trauma or injury.

A framework has been proposed to guide ergonomic practitioners through all steps of designing and evaluating manual workstations [37]. These authors believe that ergonomic intervention is the best strategy for preventing work-related musculoskeletal trauma and therefore propose human-cobot collaboration as a way of improving work conditions in manufacturing. A strong collaboration framework will have a positive impact on productivity and flexibility, and a net positive effect on new job creation rather than accelerate replacement of workers. By assisting with or taking charge of the most demanding physical tasks, robots decrease the ergonomic problems that result from physical and cognitive stress in the workplace, and thus improve safety, quality, and productivity.

A preliminary design of a mechatronic interface has been proposed that could allow integration of different cobots and mobile manipulators [38]. Though robots relieve factory workers of repetitive non-ergonomic tasks, the standards (or guidelines) applicable to evaluating emergent ergonomic risks are not always easy to interpret. Integration of cobots

and mobile manipulators increases the risk of accidental contact between humans and devices whether the robot is moving or not, notwithstanding the use of sensors to map the workspace and detect the presence of persons therein. Torque and force sensors detect undesirable contacts and thus reduce the damage potentially sustained by humans. However, the safety regulatory framework for mobile manipulators is vague, and potentially dangerous situations need to be anticipated before such robotic applications are implemented. Engineers must recommend complementary measures adapted to the specific scenario in which the system will be operated.

1.3.3.5. Psychosocial Risks and Proposed Solutions

The safety and reliability of the workspace depend on the operator maintaining full confidence in the performance and consistency of the mechanical teammate. A degradation of this trust is likely to compel human operators to allocate increased cognitive resources to monitoring and validating the accuracy of the robot's actions. This will increase mental fatigue and psychological risks (stress, anxiety, depression…) A cobot must be able to maintain a physical separation, a safe distance. To achieve this, a psychological safety field (PSF) model has been proposed, based on experimentation at different speeds and minimal distances from the head, chest and abdomen, to obtain the ordinary surface equation of psychological stress as a function of minimal speed and distance using data adjustment [39]. Using this model, the researchers were able to calculate the specific separation distance at which the cobot must stop in order not to increase the psychological risk to the human workers. However, they note that despite these calculations and the associated precautionary measures, the psychological risks persist.

In the application of collaborative robotics, analysis of psychological risks is generally necessary to identify possible dangers, estimate OHS risks and identify appropriate risk mitigation strategies. A novel risk analysis tool and methodology have been developed to support the management of change following the introduction of cobots into an Industry 4.0 capable factory [40]. Such change has complex consequences because it involves numerous parameters, some controllable and others not, for example, acceptance by workers of cohabiting with cobots.

Psychosocial factors must be considered in any analysis of the integration of new technologies into a business if the decision makers are to be fully appraised of the emergent risks inherent in the project. To facilitate the introduction of cobots and deal adequately with the concern of striking an acceptable balance between human safety and cobotic productivity, the workers must be involved in the process beginning at the design phase and throughout cobot integration and utilization. This increase in their participation, sensitization and training have a positive effect on their wellbeing and on their acceptance of working alongside robots [8].

1.4. Discussion

Industry is currently undergoing several changes that are due not only to the evolution of market demand but especially to the integration of more and more technologies into manufacturing businesses. The main goal of this development is to improve profit margins and productivity as well as quality and consistency while creating opportunities to increase product diversity and offer customization. As consumers demand ever-increasing sophistication and personalization, mass production becomes decreasingly able to keep up. It has become crucial for businesses in the present era of industrialization to diversify and increase their adaptability, to respond quickly to changing needs. In this sense, human capacities are limited and do not adequately support such responsiveness. In the search for effective solutions to meet these challenges, Industry 4.0 has emerged, through the convergence of digital communications and numerical technologies, to drive the necessary evolution and update the workplace to current modern reality. Amidst this rapid change, the place of the human worker in manufacturing is being redefined.

AGVs and cobots have become icons of the new industrial era. These devices allow automation of two of the most onerous industrial activities, namely production/assembly line tasks and inventory management. Rapid and precise vehicular movement without direct participation of workers is an easy sell in many businesses, as are robots that assist workers with monotonous, tedious, or dangerous operations to achieve greater efficiency and reliability. These advantages do not come without OHS risks, associated mainly with collisions between the cobot and

workers but also with other aspects of this integration of human labor into technological progress (Fig. 1.3).

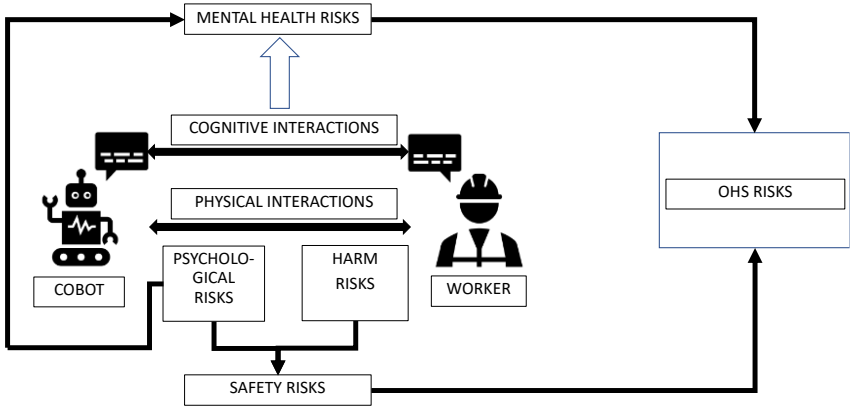


Fig. 1.3. Associations between challenges to worker safety and health (physical and mental) and the integration of cobots and AGVs into industrial settings.

Striking the ideal balance between the autonomy of new robotic technologies and the conscientization of workers obliged to collaborate with these devices in the same space raises several challenges. To improve the safety of such a venture, rules have been defined with the aim of eliminating risks by designing equipment appropriately and with safety features such as sensors and accessible stop buttons, providing informative markings and warning panels, educating and training production staff, programmers and service personnel through practical exercises, and using personal protective wear or devices [22]. In parallel, several researchers concerned with the OHS implications of these developments have produced a set of recommendations. Risk analysis and evaluation of workplace safety must remain ongoing and thorough [1].

Fine-tuning, inspection, and preventive as well as predictive maintenance will remain necessary [4]. Ergonomics must be built into equipment from the design stage, and the possibilities of complementing or replacing the human input into a task must be evaluated [3]. Before integrating new equipment into any industrial setting, measures aimed at reducing OHS risks must be incorporated into the project. Analysis must

be broad enough to identify the risks and dangers associated with the activities of autonomous technologies [4].

In the present chapter, which summarizes the first step of a broader project, we have presented a bibliographic review of the field to show how most researchers have been addressing OHS concerns associated with AGVs and cobots in the workplace, primarily in industrial settings. Based on our examination of the published literature, we conclude the following:

- Researchers have shown much more interest in the physical safety challenges than in the mental or psychological impact of the rapid changes taking place continually in the workplace. This latter impact needs more attention, given the human tendency to resist change of any kind;
- Most of the research focuses on fragmented solutions that address very specific problems rather than build intrinsic safety into autonomous equipment starting at the design phase. This appears as a disregard for the fundamental principle of eliminating dangers at the source;
- Cybersecurity, a major and universal challenge, does not appear to be a research priority in the literature focused on AGVs and cobots;
- Training and qualification of personnel in the new industrial reality remain inconspicuous as a research area, despite their importance in the emerging digital era. The role of personnel qualification in reducing the risks inherent in this transition can hardly be overstated. Proper training and education of workers can only make their collaborations with autonomous mechanical agents safer.

1.5. Conclusions

The contributions of AGVs and cobots to factory productivity, product quality and worker autonomy are obvious. However, the OHS challenges associated with their growing use are multiple. The rapid and sometimes radical changes occurring in the workplace do not always lead to a quick adaptation of managers and workers to the new reality.

OHS in general is still receiving insufficient attention from researchers. To serve its purpose of uncovering effective solutions to real problems in the current workplace, OHS research requires multidisciplinary

analysis covering design, integration, and utilization of new autonomous technologies, with closer proximity to real situations. In other words, more experimental study involving designers and users of AGVs and cobots in different economic sectors is needed.

Acknowledgments

The authors thank the Université du Québec à Trois-Rivières (UQTR), Mitacs Canada and Natural Sciences and Engineering Research Council of Canada (NSERC) for their financial support.

Abbreviations

The following abbreviations are used in this chapter:

AGV - Autonomous Guided Vehicle;

Cobot - Collaborative robot;

OHS - Occupational Health and Safety.

References

[1]. A. Badri, B. Boudreau-Trudel, A. S. Souissi, Occupational health and safety in the Industry 4.0 era: A cause for major concern?, *Safety Science*, Vol. 109, 2018, pp. 403-411.

[2]. M. N. Ben Aziza, A. Badri, F. Chihi, OHS-related risks in an Industry 4.0 manufacturing plant, in Handbook of Smart Materials, Technologies, and Devices: Applications of Industry 4.0 (C. M. Hussain, P. Di Sia, Eds.), *Springer International Publishing*, Cham, 2020, pp. 1-20.

[3]. S. Jocelyn, et al., Robotique collaborative: Évaluation des fonctions de sécurité et retour d'expérience des travailleurs, utilisateurs et intégrateurs au Québec, *IRSST*, 2017 (in French).

[4]. A. Pauliková, Z. Gyurák Babeľová, M. Ubárová, Analysis of the impact of human–cobot collaborative manufacturing implementation on the occupational health and safety and the quality requirements, *International Journal of Environmental Research and Public Health*, Vol. 18, Issue 4, 2021, 1927.

[5]. M. Valori, et al., Validating safety in human–robot collaboration: Standards and new perspectives, *Robotics*, Vol. 10, Issue 2, 2021, 65.

[6]. Y. Ayoubi, et al., V2SOM: A novel safety mechanism dedicated to a cobot's rotary joints, *Robotics*, Vol. 8, Issue 1, 2019, 18.

[7]. K. Aliev, D. Antonelli, Proposal of a monitoring system for collaborative robots to predict outages and to assess reliability factors exploiting machine learning, *Applied Sciences*, Vol. 11, Issue 4, 2021, 1621.
[8]. N. Berx, et al., Assessing system-wide safety readiness for successful human–robot collaboration adoption, *Safety*, Vol. 8, Issue 3, 2022, 48.
[9]. K. Yang, et al., Implementation of a sponge-based flexible electronic skin for safe human–robot interaction, *Micromachines*, Vol. 13, Issue 8, 2022, 1344.
[10]. M. Zhang, et al., Coordination between connected automated vehicles and pedestrians to improve traffic safety and efficiency at industrial sites, *IEEE Access*, Vol. 10, 2022, pp. 68029-68041.
[11]. S. H. Park, et al., Automatic guided vehicles introduction impacts to roll-on/roll-off terminals: Simulation and cost model analysis, *Journal of Advanced Transportation*, Vol. 2022, 2022, 6062840.
[12]. G. Li, et al., A master-slave separate parallel intelligent mobile robot used for autonomous pallet transportation, *Applied Sciences*, Vol. 9, Issue 3, 2019, 368.
[13]. A. Meysami, et al., Investigating the impact of triangle and quadrangle mesh representations on AGV path planning for various indoor environments: With or without inflation, *Robotics*, Vol. 11, Issue 2, 2022, 39.
[14]. M. Reger, J. Stumpenhausen, H. Bernhardt, Evaluation of LiDAR for the free navigation in agriculture, *AgriEngineering*, Vol. 4, Issue 2, 2022, pp. 489-506.
[15]. D. Puppim de Oliveira, W. P. N. Dos Reis, O. Morandin Junior, A qualitative analysis of a USB camera for AGV control, *Sensors*, Vol. 19, Issue 19, 2019, 4181.
[16]. L. Zamora-Cadenas, I. Velez, J. E. Sierra-Garcia, UWB-based safety system for autonomous guided vehicles without hardware on the infrastructure, *IEEE Access*, Vol. 9, 2021, pp. 96430-96443.
[17]. M. Czubenko, Z. Kowalczuk, A simple neural network for collision detection of collaborative robots, *Sensors*, Vol. 21, Issue 12, 2021, 4235.
[18]. Y. Ayoubi, et al., Safe pHRI via the variable stiffness safety-oriented mechanism (V2SOM): Simulation and experimental validations, *Applied Sciences*, Vol. 10, Issue 11, 2020, 3810.
[19]. F. Mohammadi Amin, et al., A mixed-perception approach for safe human-robot collaboration in industrial automation, *Sensors*, Vol. 20, Issue 21, 2020, 6347.
[20]. S. K. L. Andersson, et al., Experienced challenges when implementing collaborative robot applications in assembly operations, *International Journal of Automation Technology*, Vol. 15, Issue 5, 2021, pp. 678-688.
[21]. A. Scibilia, et al., Analysis of interlaboratory safety related tests in power and force limited collaborative robots, *IEEE Access*, Vol. 9, 2021, pp. 80873-80882.

[22]. G. Pang, et al., Development of flexible robot skin for safe and natural human-robot collaboration, *Micromachines*, Vol. 9, Issue 11, 2018, 576.
[23]. M. Geiger, C. Waldschmidt, 160-GHz radar proximity sensor with distributed and flexible antennas for collaborative robots, *IEEE Access*, Vol. 7, 2019, pp. 14977-14984.
[24]. J. S. Sandoval Arévalo, et al., On the design of a safe human-friendly teleoperated system for Doppler sonography, *Robotics*, Vol. 8, Issue 2, 2019, 29.
[25]. S.-H. Yen, et al., Development of a virtual force sensor for a low-cost collaborative robot and applications to safety control, *Sensors*, Vol. 19, Issue 11, 2019, 2603.
[26]. H. Shin, et al., A virtual pressure and force sensor for safety evaluation in collaboration robot application, *Sensors*, Vol. 19, Issue 19, 2019, 4328.
[27]. D. Engelbrecht, N. Steyn, K. Djouani, Adaptive virtual impedance control of a mobile multi-robot system, *Robotics*, Vol. 10, Issue 1, 2021, 19.
[28]. G. Chiriatti, et al., Adaptive obstacle avoidance for a class of collaborative robots, *Machines*, Vol. 9, Issue 6, 2021, 113.
[29]. S. B. i. Badia, et al., Virtual reality for safe testing and development in collaborative robotics: Challenges and perspectives, *Electronics*, Vol. 11, Issue 11, 2022, 1726.
[30]. Q. Zhang, W. Liu, P. Liu, Fault-tolerant control of automated guided vehicle under centroid variation, *IEEE Access*, Vol. 10, 2022, pp. 68995-69009.
[31]. K. Liu, et al., Model predictive stabilization control of high-speed autonomous ground vehicles considering the effect of road topography, *Applied Sciences*, Vol. 8, Issue 5, 2018, 822.
[32]. C.-K. Tsung, F.-S. Chang, X.-Y. Liu, On the construction of an edge-based remote sensing framework: The applications on automated guided vehicles and drones, *Electronics*, Vol. 11, Issue 7, 2022, 1034.
[33]. S. O. B. Islam, et al., Exploiting visual cues for safe and flexible cyber-physical production systems, *Advances in Mechanical Engineering*, Vol. 11, Issue 12, 2019, 1687814019897228.
[34]. C.-C. Chan, C.-C. Tsai, Collision-free speed alteration strategy for human safety in human-robot coexistence environments, *IEEE Access*, Vol. 8, 2020, pp. 80120-80133.
[35]. D. Gradolewski, et al., A distributed computing real-time safety system of collaborative robot, *Elektronika ir Elektrotechnika*, Vol. 26, Issue 2, 2020, pp. 4-14.
[36]. R. Behrens, et al., A statistical model to determine biomechanical limits for physically safe interactions with collaborative robots, *Frontiers in Robotics and AI*, Vol. 8, 2021, 683706.
[37]. A. Colim, et al., Towards an ergonomic assessment framework for industrial assembly workstations—A case study, *Applied Sciences*, Vol. 10, Issue 9, 2020, 3048.

[38]. F. Vitolo, et al., Mobile robots and cobots integration: A preliminary design of a mechatronic interface by using MBSE approach, *Applied Sciences*, Vol. 12, Issue 1, 2022, 419.
[39]. B. Liu, et al., Cobot motion planning algorithm for ensuring human safety based on behavioral dynamics, *Sensors*, Vol. 22, Issue 12, 2022, 4376.
[40]. M. Pantano, et al., Novel approach using risk analysis component to continuously update collaborative robotics applications in the smart, connected factory model, *Applied Sciences*, Vol. 12, Issue 11, 2022, 5639.

2.

Sociability Modelling in Robot Motion

Haiwei Luo, Jin Zhang, Yaqing Luo and Gaoyong Luo

2.1. Introduction

From harvesting crops to assembling automobiles and delivering medications, modern robotics is the driving force behind these diverse tasks, and the deployment of robotic systems can now enhance productivity, safety, and flexibility in various industries such as manufacturing, agriculture, and healthcare [1]. With technological advances, robots are no longer confined to industrial applications and have now become part of the habitats and workspaces of humans, such as self-driving cars. Furthermore, the emergence of humanoid robots is accelerating the deployment of robotic systems in everyday human environments, as their humanlike appearance naturally evokes expectations of human capabilities during task execution [2]. Robots can now be used to assist industrial workers in factories as co-workers (cobots), support care attendants in nursing homes or hospitals, help customers in a supermarket, or do the housework at home. Particularly service, medical, or social robots [3], which are machines designed by humans to perform tasks automatically or under the control of a computer program, and equipped with sensors, actuators, and a processing unit, and carrying out specific functions, have integrated mechanical and electronic engineering for crafting tangible components and computational systems for information processing and decision-making, allowing them to interact with humans [4]. In some cases, they can operate by responding to or learning from human movements through supervising or training their behaviours during deployment [5]. As such, determining a positively perceived interaction for socially accepted robot behaviour in human environments is very

Haiwei Luo
Department of Computing, Imperial College London, London SW7 2AZ, UK

important, where legibility means that an interacting human is able to understand the robots' intentions [6].

Other crucial aspects of robot behaviour include predictability and trustability. The predictability of a robot's actions is a key factor in successful human–robot interactions (HRIs) [7], where humans can anticipate what the robot will do. The trustability of robot interactions, often referred to as human–robot trust, has emerged as a new research area that is gaining attention. However, a more inclusive definition of trust, encompassing all aspects of human–robot trust, requires further investigation [8].

On the other hand, the social context of HRIs, which is defined in diverse ways, is critical for the design, evaluation, and automatic generation of appropriate robot behaviour [9]. This variability makes it challenging to draw connections between related work on understanding and modelling human sociability in robot motion.

Particularly in real-time and dynamic environments, generating socially accepted motions and aligning robot behaviour with social expectations remain a challenge, due to the complex mechanisms and reasoning processes that require the incorporation with human intelligence. For example, deploying multimodal artificial intelligence (AI) to learn from large datasets, identify patterns, and apply logical inference, could allow robots to engage in natural social interactions, ensuring that social behaviour is appropriate and safe. Therefore, this study aims to address the challenge by investigating HRIs focusing on a social context, in which a robot operates in close proximity to humans. For the robot to blend naturally into the social environment, its motion must not only achieve the task but also adhere to social norms in a way that is both easily understood (legible) and reassuring to those around it (trustable). Our research [10] has shown that these two properties are the essential pillars of robot sociability. Both legibility and trustability hinge on the robot's way of expressing its intent through motion. Together, they define sociability as the ability of a robot to communicate its intentions in a human-like way that is, in turn, socially appropriate. The legibility component allows humans to clearly interpret the robot's intended actions, while the trustability component is required for humans to feel comfortable and safe during interactions, based on their interpretation of robot intent. In this sense, sociability arises not only from what the robot intends to do, but how it communicates that intent to humans. As a result, it is important to ensure the intent is conveyed in a way that can be

anticipated and interpreted by humans, to create motion that is socially predictable; our study [11] demonstrated that such trajectories were recognisable by humans. By making intent interpretable, clear, and comforting, robots can foster human confidence and acceptance, enabling genuinely sociable behaviour in HRI.

Furthermore, social intelligence in robots arises from the ability to model, recognise, and reproduce aspects of human social behaviour [12]. In HRI, social behaviour is not only expressed through verbal communication, but also through movement, spatial awareness, intent signalling, and the ability to adapt to others [13]. For robots to behave socially, they must therefore be capable of identifying these behaviours and generating actions that are consistent with human expectations. This makes the recognition of socially predictable motion an important step towards creating social intelligence in robotic systems. For this purpose, we propose not only a method to model and quantify human socialness onto robots for generating socially predictable trajectories by addressing the issues identified and discussed, but also a model-free algorithm for identifying this type of motion.

However, social intelligence requires a robot not only to recognise and interpret human behaviour, but also to anticipate how its own actions will be perceived and to adapt those actions accordingly. This motivates the modelling of sociability through robot motion that is both legible and trustable. By incorporating human perceptions and expectations into these two properties, the sociability model captures aspects of human social behaviour and provides a basis for generating motions that humans can interpret, anticipate, and accept.

To this end and for future work, we aim to further build social intelligence by integrating the proposed sociability model with a learning-based framework that enables robots to evaluate the social consequences of their actions and adapt their behaviour accordingly. A deep model-based reinforcement learning approach thus will be developed to utilise the sociability model to generate socially acceptable trajectories in real-time, dynamic environments. Specifically, a combination of deep reinforcement learning (DRL) and hierarchical sensory processing based on a fast deep learning model [14] previously developed could be deployed, in which DRL combines function approximation with objective optimisation to map states and actions to rewards. By constructing an agent that learns through trial and error, as well as adjusting its actions based on feedback from the environment, it

enables robots to simulate human capabilities such as sociability. This framework could be implemented using a Dreamer-style architecture [15, 16], with the sociability model serving as part of the reward model and an actor–critic method used to learn a policy for low-latency action selection. In this way, the robot could anticipate how its motion may be perceived and select actions that are not only task-effective, but also legible, trustable, predictable, and in turn, socially predictable. Ultimately, this approach could facilitate more natural and effective HRI by supporting the development of robots that function not merely as tools, but as intelligent, supportive, and collaborative entities capable of fostering greater human confidence, acceptance, and trust.

2.2 Related Work and Methodology

2.2.1. Notions of Predictability, Legibility, Trustability and Sociability

2.2.1.1. Predictability

Predictability in robot motion has been formally defined [17] as the extent to which the robot's motion matches with the expectation of the human observer, for a specified goal. If we consider actions to be goal-oriented, where many studies in psychology suggest that humans naturally perceive actions in this manner [18], then we can explain the inference occurring in predictability as "goal-to-action".

This has been defined [17] as the inference mapping goals $\mathcal{G}$ to trajectories Ξ, where $\mathcal{I}_{\mathcal{P}}$ models the observer's expectation:

$$\mathcal{I}_{\mathcal{P}}: \mathcal{G} \rightarrow \Xi \tag{2.1}$$

Furthermore, predictable motion is defined [17] as one where the trajectory from an initial point S to a goal point G, denoted as $\xi_{S \rightarrow G}$, matches the following inference:

$$\mathcal{I}_{\mathcal{P}}(G) = \xi_{S \rightarrow G} \tag{2.2}$$

The optimality of the trajectory inference is evaluated by the predictability score, affected by the trajectory's cost, which penalises the trajectory's length, as it is established that humans expect a direct

trajectory to the goal. Shorter trajectories are considered to be more "efficient", thus the predictability score can be interpreted as a measure of the "efficiency" of a trajectory. Formally, the predictability score of a trajectory ξ is defined [17] as a normalised value between 0 and 1:

$$\text{predictability}(\xi) = \exp\big(-C(\xi)\big), \tag{2.3}$$

where C denotes the cost function, modelling the observer's expectation, and can be defined in terms of the trajectory length.

2.2.1.2. Legibility

Legibility in robot motion is also formally defined as the degree to which a human observer is able to infer the robot's intent from its motion [17]. In the context of goal-oriented actions [18], legible motion can be understood as "action-to-goal", where the robot expresses its intended goal through its motion.

This has been defined [17] as the inference mapping the snippets of trajectories, from all trajectories to the goal:

$$\mathcal{I}_{\mathcal{L}}: \Xi \rightarrow \mathcal{G} \tag{2.4}$$

Moreover, legible motion is [17] defined as one which allows for quick and confident inference of the correct goal G by the observer, after observing a snippet of the trajectory $\xi_{S \rightarrow Q}$, from an initial point S to the configuration at a time t, where $Q = \xi(t)$:

$$\mathcal{I}_{\mathcal{L}}\big(\xi_{S \rightarrow Q}\big) = G \tag{2.5}$$

The optimality of the goal inference across an entire trajectory is evaluated using the legibility score, tracking the assigned probability of the actual goal G^* across the trajectory denoted by $P\big(\, G^* \mid \xi_{S \rightarrow \xi(t)} \,\big)$. This is weighted using $f(t)$ which assigns more weight to the earlier segments of the trajectory, by employing for example $f(t) = T - t$, where T is the duration of the entire trajectory. This ensures the weight decreases over time; as t increases, $f(t)$ decreases, assigning lower weights to later segments of the trajectory. The score is normalised by dividing $\int f(t)dt$ for each weighted assigned probability of G^* across the trajectory. Formally the legibility score of a trajectory ξ is given by [17]:

$$\text{legibility}(\xi) = \frac{\int P\left(G^* | \xi_{S \to \xi(t)}\right) f(t)\, dt}{\int f(t)\, dt} \tag{2.6}$$

To compute $P\left(G^* \mid \xi_{S \to \xi(t)}\right)$, C can be used as shown below [17], which calculates the inference of a goal G, given a trajectory snippet $\xi_{S \to Q}$:

$$P\left(G \mid \xi_{S \to Q}\right) \propto \frac{\exp\left(-C(\xi_{S \to Q}) - C(\xi^*_{Q \to G})\right)}{\exp\left(-C(\xi^*_{S \to G})\right)} P(G), \tag{2.7}$$

where ξ^* denotes an optimal trajectory.

2.2.1.3. Trustability

Unlike predictability and legibility, the notion of trustability lacks a clear definition in the context of robot motion. Insights from psychology and ideas of trust already established in robotics provide us a foundation for its formalisation.

Psychological models of trust [19, 20] identify three key factors in shaping belief in another agent from the perspective of an observer: benevolence, competence, and integrity. Benevolence, the extent to which the observer perceives the agent as caring and harmless, is the most critical factor when humans evaluate unfamiliar agents such as robots.

In robotics, trust has historically been framed in terms of the avoidance of harm or hostility towards humans [21], echoing Asimov's laws [22]. Benevolence thus dominates, as without clear signals that a robot lacks harmful intent, no initial trust can form, particularly given humans' tendency to fear the unknown [23]. Trust has also been linked to the willingness to be vulnerable [24, 25], which is encouraged when robots exhibit benevolent or socially aware behaviours, such as acknowledging human presence through movement cues [21, 25]. Experiments further show that humans trust robots more when they display submissive moment strategies that account for human actions, rather than dominant ones [25].

Moreover, studies have shown that initial trust in the agent has a lasting effect, whereby early impressions strongly influence long-term trust, even after subsequent violations [20, 26]. Therefore, agents must signal

the mentioned factors that influence trust early in the interaction with the observer.

2.2.1.4. Sociability

Sociability in robot motion is commonly defined in the context of socially aware navigation. For example, it has been suggested that sociability is achieved when humans become comfortable [27] or when robots adhere to cultural conventions [28]. A key way of enforcing these social conventions in robot motion is through Hall's theory of proxemics [29], which concerns the use and preservation of interpersonal space during social interaction. In human–robot interaction (HRI), this is commonly represented as a personal or safety zone around the human that the robot should avoid violating unless the task explicitly requires close interaction. This ensures that the robot does not simply avoid physical collision, but also respects the human's psychological comfort and perceived safety.

However, it should be stressed that there still lacks a formal and generic definition for sociability in robot motion, which could encompass robot arm movements.

2.2.2. Predictable, Legible, and Sociable Motion Planning

The notions of predictability, legibility, and sociability have been discussed in the field of robotics motion planning. For instance, Ngo et al. [30] proposed a novel potential field method that ensured the legibility property of motion, compared with traditional field methods that only guaranteed predictability. This approach is beneficial in human-centred environments, since intent-expressive motion could be generated.

Furthermore, Bodden et al. [31] proposed a new intent-expressive objective function, that splits motion into two separate components – one for reaching the goal and the other that allows for legibility. This proposal allows for the objective function to simultaneously optimise for direct motion (predictable) and legible motion.

A machine learning-based approach in generating legible motion which is less dependent on pre-set tasks has been proposed by Busch et al. [6], where a legible policy is learnt through reinforcement learning.

Nikolaidis et al. [32] expanded on the idea of legibility [17] by incorporating observer perspective. They first introduced a viewpoint model, projecting trajectory and goals onto the observer's view plane, to generate viewpoint-based legible motion. Secondly, they introduced an occlusion model, capturing the change in legibility when parts of the trajectory are hidden from the observer. By taking into account occlusion in different segments of the trajectory, occlusion-based legible motion can then be generated.

2.2.2.1. Gradient Descent Methods

In Dragan's thesis [33], predictable and legible motion have been generated by utilising the gradient descent update rule for trajectory ξ:

$$\xi_{i+1} = \xi_i - \frac{1}{\eta} A^{-1} \nabla_{\xi_i} \mathcal{U}, \tag{2.8}$$

where $\nabla_{\xi_i} \mathcal{U}$ is the Euclidean gradient and η is a coefficient to adjust the step-size. A is a matrix that has off-diagonal non-zero elements, which establishes a relationship between the current time point, and the previous and next ones, by satisfying:

$$\langle \xi_1, \xi_2 \rangle_A = \xi_1^\top A \xi_2, \tag{2.9}$$

which adapts from the Euclidean inner product as defined as:

$$\langle \xi_1, \xi_2 \rangle_I = \int \xi_1(t)^\top \xi_2(t)\, dt \tag{2.10}$$

The cost function C is set as the integral of the squared velocities within the configuration space. The predictability gradient has been derived as the negative acceleration (without considering obstacles), and thus the most predictable motion is attained by minimising C, which occurs when the negative acceleration is 0.

Moreover, the legibility gradient has been derived as the derivative of normalised weighted probability inference of the actual goal across the trajectory, with respect to trajectory ξ (without accounting for obstacles):

$$\nabla \text{Legibility}(t) = K \frac{\partial P}{\partial \xi}(\xi(t), t), \tag{2.11}$$

where

$$P(\xi(t), t) = P\left(G^* \middle| \xi_{S \to \xi(t)}\right) f(t),\ K = \frac{1}{\int f(t)dt}$$

2.2.2.2. Methods for Improving Predictability

To generalise to new situations, where the robot can adapt to new initial and goal positions, we can employ the learning from demonstration method to further improve predictability. To optimise over a Hilbert space of trajectories, model-free approaches [33] can be used through generalising the adaptation process, such that trajectory adaptation can be formalised as a Hilbert norm minimisation problem.

A method of encouraging humans to adapt to the robot, known as familiarisation, which exposes humans to repeated robot motions, can be employed to increase predictability over time. However, inherently unnatural motions would never achieve high predictability, regardless of exposure.

2.2.2.3. Methods for Improving Legibility

It has been observed that sometimes when a trajectory is optimised, it can become arbitrarily unpredictable [33], because the definition of goal inference [17] is unable to capture the way humans make inferences in unpredictable situations. Dragan's solution was to introduce a trust region of predictability [33], which constrains legibility optimisation to only when predictability is high enough. This constraint is defined with respect to the cost function C, that bounds the domain of trajectories. Thus, the legibility model is only trustable if it lies inside the trust region:

$$C[\xi] \leq \beta \tag{2.12}$$

In the context of Yu et al.'s trust model [20], the utilisation of the trust region ensures display of competence and integrity for a robot, as the purpose of the trust region is to limit a human's confusion regarding the robot's trajectory.

2.2.2.4. Socially Aware Navigation

Eirale et al. [34] proposed using an additional cost map to encode the social constraints in navigation, that incorporates information such as the positions of people and goals. By using a DNN (deep neural network), a

social cost map is trained to recognise and encode social obstacles such as a group of people or a queue. When integrated with a discrete path planner alongside the local and global cost maps, a socially acceptable path is generated, where for example, the plan allows the navigation of the robot around a group of people and towards the back of a queue, ensuring adherence to social norms.

On the other hand, Singamaneni et al. [35] reviewed socially aware robot navigation, highlighting the importance of human comfort, trust, proxemics, intention communication, and context-dependent evaluation.

Although both of these studies of literature mainly concern robot navigation, these principles are also relevant to manipulators operating near humans. A research gap remains in quantitatively evaluating such social behaviours and in generating manipulator trajectories that exhibit them.

2.3. Sociability Modelling and Quantification

For robots to create social intelligence through modelling sociability to generate socially predictable trajectories, we begin by establishing that a plan (trajectory) generated by a planner consists of many waypoints, or local goals that the robot arm must traverse via, from the initial point in order to reach the goal point. We discretise time by treating each waypoint as a discrete time unit, where the index in the list of waypoints directly determine the time. This is formalised as:

$$Q_n = \xi[n], \tag{2.13}$$

where Q_n describes the nth waypoint on trajectory ξ.

In particular Q_0 describes the initial point and Q_{N-1} describes the goal point.

We then compute the scores based on the assumption that the path between each adjacent pair of trajectories to be a straight line. We also use the trajectory length as the cost function C. This means that the following holds for a (partial) trajectory between two waypoints:

$$C\left(\xi_{Q_1 \to Q_2}\right) = C\left(\xi^*_{Q_1 \to Q_2}\right), \tag{2.14}$$

where $\xi^*_{Q_1 \to Q_2}$ denotes the optimal trajectory from waypoint Q_1 to waypoint Q_2. More generally:

$$C\left(\xi_{Q_n \to Q_{n+1}}\right) = C\left(\xi^*_{Q_n \to Q_{n+1}}\right), \text{for } n \geq 0,\ n \in Z \qquad (2.15)$$

2.3.1. Predictability

To compute the predictability score of a trajectory, we can follow Equation (2.3). However, it is noted that the generated trajectory is not a continuous line, but rather waypoints that form a series of line segments or partial trajectories; we therefore propose a discretised predictability score as:

$$\text{predictability}_d(\xi) = \exp\left(-\sum_{n=0}^{N-2} C\left(\xi_{Q_n \to Q_{n+1}}\right)\right), \qquad (2.16)$$

where N is the number of waypoints, inclusive of the initial and goal points.

2.3.2. Predictability Prime (Predictability′)

The predictability score is proficient at evaluating multiple trajectories with the same initial and goal points, as the trajectory with the highest predictability score would be the one with the lowest cost (length). However, this score is not very useful in evaluating trajectories with different initial and goal points, since some trajectories could be naturally longer due to the distance between the initial and goal points, yielding a lower predictability score, even if they still remain predictable. Hence, we propose a modified predictability score known as the predictability prime (predictability′) score that takes into account the shortest possible trajectory, given an initial point p_{init} and goal point p_{goal}. This is the straight-line distance between the two points. Formally, the predictability prime score measures the amount of deviation between the actual trajectory and the shortest possible trajectory, and is defined as follows:

$$\text{predictability}'(\xi) = \exp\left(-\left(C(\xi) - C_{\min}(\xi)\right)\right), \qquad (2.17)$$

where $C(\xi)$ is the trajectory distance, and $C_{\min}(\xi)$ is the shortest possible trajectory distance, given by the Euclidean distance between p_{init} and p_{goal} in 3D:

$$C_{\min}(\xi) = \left\| p_{\text{goal}} - p_{\text{init}} \right\|_2 =$$
$$= \sqrt{\left(x_{\text{goal}} - x_{\text{init}}\right)^2 + \left(y_{\text{goal}} - y_{\text{init}}\right)^2 + \left(z_{\text{goal}} - z_{\text{init}}\right)^2} \quad (2.18)$$

Note that $C_{\min}(\xi)$ evaluates to the same value as $C(\xi^*)$, since the optimal trajectory is the straight-line trajectory between two points.

Furthermore, Equations (2.16) and (2.17) can be combined to compute the discretised predictability prime score:

$$\text{predictability}'_d(\xi) =$$
$$= \exp\left(-\left(\left(\sum_{n=0}^{N-2} C\left(\xi_{Q_n \to Q_{n+1}}\right)\right) - C_{\min}\left(\xi_{Q_0 \to Q_{N-1}}\right)\right)\right), \quad (2.19)$$

with the computation procedure illustrated in Fig. 2.1.

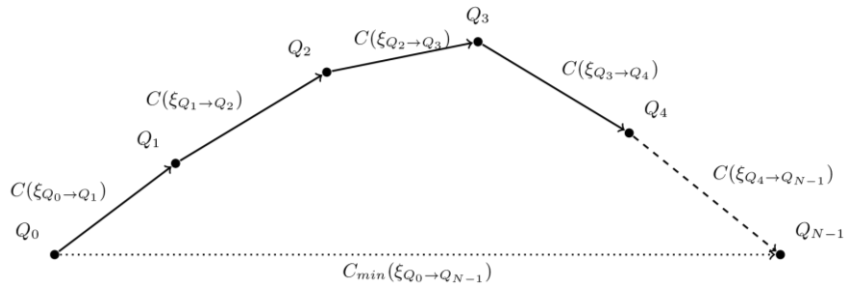


Fig. 2.1. The computation of the discretised predictability and predictability prime scores.

2.3.3. Legibility

To compute the legibility, we first follow Equation (2.6) to calculate the legibility score, as well as Equation (2.7) to derive the probability of the goal inference of a trajectory snippet. In this study, we assume the number of possible goals to be one, as we are only evaluating the legibility of a trajectory after a plan towards the goal that has been generated. This means that it is certain that the trajectory would go towards the intended goal, resulting in $G^* = G$ and $P(G) = 1$ in Equations (2.6) and (2.7). For the weighting function $f(t)$, we refer to Dragan et al.'s example [17] of setting $f(t) = T - t$ to give more weight to earlier parts of the trajectory.

For the quantifying purpose, it is required to discretise the legibility score for computation. The numerator of Equation (2.6) can be discretised as follows:

$$\sum_{n=0}^{N-1} P\big(\, G \mid \xi_{S\to Q_n} \big) f[n], \tag{2.20}$$

where $f[n] = N - n$ and $\xi_{S\to Q_n}$ represents the trajectory snippet from the initial point to the waypoint that denotes the end of the trajectory snippet. For example, $\xi_{S\to Q_0}$ can be understood as the 0th trajectory snippet, which spans from the initial point to the waypoint that marks the end of the 0th trajectory. In this context, the goal point is the same as Q_{N-1} and N is the number of trajectory snippets in the trajectory.

The denominator of (2.6) can be discretised as shown below:

$$\begin{aligned}\sum_{n=0}^{N-1}(N-n) &= (N-0) + (N-1) + \cdots + 2 + 1 = \\ &= \sum_{n=1}^{N} n = \frac{1}{2}N^2 + \frac{1}{2}N\end{aligned} \tag{2.21}$$

Thus, the discretised legibility score is:

$$\text{legibility}_d(\xi) = \frac{\sum_{n=0}^{N-1} P\big(\, G|\xi_{S\to Q_n}\big) f[n]}{\frac{1}{2}N^2 + \frac{1}{2}N} \tag{2.22}$$

$P\big(\, G \mid \xi_{S\to Q_n}\big)$ can be computed using (2.7), which can be rewritten using our C_{min} notation as:

$$P\big(\, G \mid \xi_{S\to Q_n}\big) \propto \frac{\exp\big(-C(\xi_{S\to Q_n}) - C_{min}(\xi_{Q_n\to G})\big)}{\exp(-C_{min}(\xi_{S\to G}))} P(G) \tag{2.23}$$

Due to the assumption that $P(G) = 1$, we can drop that term:

$$P\big(\, G \mid \xi_{S\to Q_n}\big) \propto \frac{\exp\big(-C(\xi_{S\to Q_n}) - C_{min}(\xi_{Q_n\to G})\big)}{\exp(-C_{min}(\xi_{S\to G}))} \tag{2.24}$$

The issue with computing the legibility score directly is that we only know of the proportional relationship for $P\big(\, G \mid \xi_{S\to Q_n}\big)$, and cannot calculate the exact value for $P\big(\, G \mid \xi_{S\to Q_n}\big)$. We overcome this problem by proposing a discretised relative legibility score:

$$\text{relative legibility}_d(\xi) = \frac{\sum_{n=0}^{N-1} \tilde{P}\left(G|\xi_{S\to Q_n}\right) f[n]}{\frac{1}{2}N^2+\frac{1}{2}N}, \qquad (2.25)$$

where $\tilde{P}\left(G \mid \xi_{S\to Q_n}\right)$ is computed as an equation instead and shown in Fig. 2.2:

$$\tilde{P}\left(G \mid \xi_{S\to Q_n}\right) = \frac{\exp\left(-C(\xi_{S\to Q_n})-C_{min}(\xi_{Q_n\to G})\right)}{\exp(-C_{min}(\xi_{S\to G}))} \qquad (2.26)$$

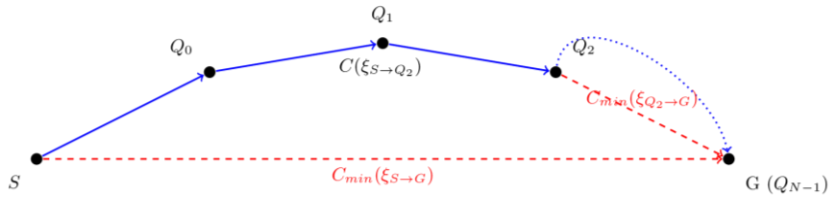


Fig. 2.2. Computation of $\tilde{P}\left(G \mid \xi_{S\to Q_2}\right)$.

2.3.4. Trustability

Next, we model trustability in robot motion based on two key principles:

1). Benevolence as the core component of trust

Inspired by the notion of benevolence in the trust model [20], where research [21-25] has shown to be the dominating factor for humans to develop trust for a robot, trustability can thus be modelled as the clearness of the display of benevolent behaviours from robots. When robot benevolence is modelled as non-adversarial robot intent, a robot would appear to be trustable if it does not display harmful intentions towards humans. This is done through the acknowledgement of human presence and clear indication that the robot's ultimate goal is not the human observer, via its trajectory.

2). Importance of initial trust formation

Studies [20, 21] have highlighted the critical role of first impressions in developing long-term trust, where trust is more likely to be sustained over time if it is established early in the interaction. We wish to model this in robot motion by the speed of trust formation, meaning that a robot is more trustable if it can establish trust quickly through its motion.

We now define trustability based on these two principles:

Definition 1. Trustability is the measure of speed at which a human can confidently infer that the robot's intended motion is non-adversarial, away from the human observer.

Intuitively, untrustability is defined as:

Definition 2. Untrustability is the measure of speed at which a human can confidently infer that the robot's intended motion is adversarial, towards the human observer.

Since the formalisation of untrustability is more intuitive, we will derive the formalisation of trustability from that of untrustability. As inspired by Dragan et al.'s formalisation of legibility [17], we first formalise untrustability by the measure of the robot's speed of establishing clear intent towards the human observer through motion. Untrustable motion is defined as one that allows quick and confident inference of the goal as the human observer H, after observing a snippet of the trajectory $\xi_{S\to Q}$, from an initial point S to the configuration at time t, where $Q = \xi(t)$:

$$\mathcal{J}_{\mathcal{U}}\left(\xi_{S\to Q}\right) = H \tag{2.27}$$

We propose the untrustability score to be between 0 and 1, which is an evaluation of the optimality of the goal inference as the human, as:

$$\text{untrustability}(\xi) = \frac{\int P\left(H|\xi_{S\to\xi(t)}\right) f(t)\,dt}{\int f(t)\,dt}, \tag{2.28}$$

where H denotes the location of the human observer, or more specifically, the observer's origin of visual perception, modelled as a point denoting the centre of the eyes. Thus, $P\left(H \mid \xi_{S\to\xi(t)}\right)$ denotes the probability inferring that the goal across the trajectory is the human observer. Also, we employ the defined weighting function $f(t) = T - t$ to model the violation of initial trust.

The discretised untrustability score is given as:

$$\text{untrustability}_d(\xi) = \frac{\sum_{n=0}^{N-1} P\left(H|\xi_{S\to Q_n}\right) f[n]}{\frac{1}{2}N^2 + \frac{1}{2}N}, \tag{2.29}$$

where $P\left(H \mid \xi_{S \to Q_n}\right)$ is the probability of inferring the goal as H, given the trajectory snippet $\xi_{S \to Q_n}$.

Furthermore, the discretised relative untrustability score is:

$$\text{relative untrustability}_d(\xi) = \frac{\sum_{n=0}^{N-1} \tilde{P}\left(H \middle| \xi_{S \to Q_n}\right) f[n]}{\frac{1}{2}N^2 + \frac{1}{2}N}, \quad (2.30)$$

which can be computed via:

$$\tilde{P}\left(H \mid \xi_{S \to Q_n}\right) = \frac{\exp\left(-C(\xi_{S \to Q_n}) - C_{min}(\xi_{Q_n \to H})\right)}{\exp\left(-C_{min}(\xi_{S \to H})\right)} \quad (2.31)$$

Conversely, the trustability score is the evaluation of the optimality of the goal inference as not the human observer ($\overline{H}$):

$$\text{trustability}(\xi) = \frac{\int P\left(\overline{H} \middle| \xi_{S \to \xi(t)}\right) f(t)\, dt}{\int f(t)\, dt}, \quad (2.32)$$

where the weighting function $f(t) = T - t$ is used to model the importance in the establishment of initial trust.

Similarly, the discretised trustability score is:

$$\text{trustability}_d(\xi) = \frac{\sum_{n=0}^{N-1} P\left(\overline{H} \middle| \xi_{S \to Q_n}\right) f[n]}{\frac{1}{2}N^2 + \frac{1}{2}N}, \quad (2.33)$$

where the discretised relative trustability score can be deduced as:

$$\text{relative trustability}_d(\xi) = \frac{\sum_{n=0}^{N-1} \tilde{P}\left(\overline{H} \middle| \xi_{S \to Q_n}\right) f[n]}{\frac{1}{2}N^2 + \frac{1}{2}N}, \quad (2.34)$$

which is computed via:

$$\tilde{P}\left(\overline{H} \mid \xi_{S \to Q_n}\right) = 1 - \tilde{P}\left(H \mid \xi_{S \to Q_n}\right) = \\ = 1 - \frac{\exp\left(-C(\xi_{S \to Q_n}) - C_{min}(\xi_{Q_n \to H})\right)}{\exp\left(-C_{min}(\xi_{S \to H})\right)} \quad (2.35)$$

2.3.5. Sociability

Legibility and trustability are two of the major factors contributing to sociability, due to the fact that both are essential for conveying a comprehensive and well-rounded representation of robot intent. Legibility ensures that the robot's goal is clearly communicated, while trustability assures the human observer that the robot's intent is non-adversarial. Our encoding of sociability would thus be influenced by both factors, where we encourage sociable behaviour through a sociability score that prioritises the clear and rapid display of robot intent. By modelling trustability as a core component of sociability, we also encode the idea of proxemics [29], as a trustable trajectory respects the human observer's personal space by prioritising motion away from the human observer, thereby enforcing a comfort region around the human. Thus, we define sociability as follows:

Definition 3. Sociability is the measure of the speed of clear robot intent conveyance towards the goal and away from the human observer, expressed as a weighted sum of legibility and trustability.

While the specific type of legibility and trustability scores (e.g., discretised or relative) used has minimal impact on the computation of the sociability score, it is important that both scores are of the same type to ensure a meaningful and consistent computation. Formally, we propose the sociability score as follows:

$$\text{sociability}(\xi) = \alpha \cdot \text{legibility}(\xi) + (1 - \alpha) \cdot \text{trustability}(\xi), \qquad (2.36)$$

where α is a value between 0 and 1, defined as the legibility weight, representing the importance of legibility in determining sociability. An α value greater than 0.5 would prioritise the trajectory being clear of its intent towards the goal, rather than the intent away from the observer. In this study, we set the α value to 0.5 by default, so that both legibility and trustability hold equal importance in computing a sociable trajectory.

As shown in Fig. 2.3, if both trajectories have the same sociability score, then ξ_1 has a higher α value compared to ξ_2, as ξ_1 focuses on expressing intent towards the goal, rather than expressing intent away from the human observer.

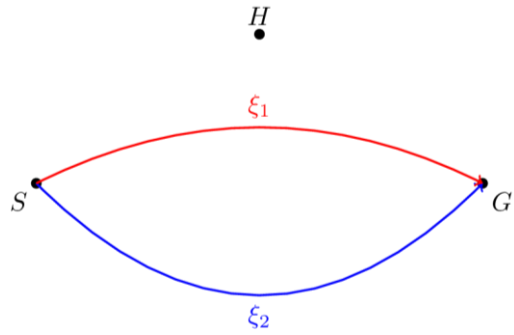


Fig. 2.3. Comparison of α in sociable trajectories.

2.3.6. Modelling Limitations

The first modelling limitation arises from the formalisation process of both legibility and trustability, where they currently lack a completely coherent Bayesian setup. We take cost-exponential ratios, then normalise them in an arbitrary way, resulting in what seems to be a heuristic index, rather than a principled score. The main contribution, however, is that the proposal of these scores does allow us to compare between different trajectories, and investigate which ones yield more of the described characteristics. Nonetheless, future work could involve constructing a coherent probabilistic formulation using explicit candidate intentions, priors, an observation model, and normalisation across the full goal set.

Furthermore, the lack of consideration of velocity, acceleration, and jerk in robot motion modelling is also a key limitation. Trustability in particular, hinges on the style of motion, which cannot be fully captured by the current geometric relationship between the end effector trajectory and the observer. For example, a trajectory may receive a high trustability score despite exhibiting motion that a human could perceive as fast, intrusive, or threatening. Future models should therefore also incorporate other characteristics of robot motion, in order to better encode perceived trust in humans. Moreover, considering the motion of the entire arm could also improve the model.

One of our modelling limitations is the single goal assumption in the formalisation of the discretised (relative) legibility score. This deviates from Dragan et al.'s interpretation [17] of legibility as a goal inference problem with multiple candidate goals. Our assumption instead can be seen as a measure of compatibility with a known goal, rather than a full

goal-disambiguation process. The main advantage of this approach is the simplification of the computational process and the legibility model, as well as providing a foundation for the formalisation of the trustability score, where there is also assumed to be one human observer. Nonetheless, the extension to multiple candidate goals is relatively straightforward. Future models that discretise legibility could compute the compatibility of each trajectory snippet with every candidate goal, then normalise these compatibilities over the entire candidate goal set. This also poses an opportunity to consider trustability with regard to multiple human observers.

Moreover, the model is also limited by trustability's narrow definition, as benevolent motion affected by first impression, although this has been justified as some of the most important factors in the relevant literature. However, the introduction of trustably predictable trajectories later on (Section 2.4.4) does incorporate other components of trust such as competence and integrity. Nonetheless, the proposal of the trustability score does not constitute a comprehensive model of trust, but rather focuses on the main components of trust in robot motion. The trustability score could be improved in the future, by incorporating other dimensions of perceived trust in HRI.

Another limitation of trustability modelling is the direct reliance on the legibility weighting scheme to encode the importance of first impressions. Not only does this not guarantee the representation of the asymmetric severity of motion directed towards a human, it also does not explicitly account for risk or motion dynamics. Future work should therefore incorporate risk-sensitive penalties and dynamic factors to better model human responses to potentially threatening motion.

Due to the relatively simple model of sociability as a result of factors such as the narrow definition of trustability and the single goal assumption of legibility, this model enables the fast computation of scores, which provides the potential for generating trajectories in real-time and dynamic environments. This is one of the major reasons why we decided to proceed with these modelling choices.

2.4. Generation of Acceptable Trajectories

We generate trajectories that are predictable, trustable, legible, or sociable, using sampling-based planners [36]. Due to its random nature,

repeated planning queries with these planners, from the same start and goal, can yield different paths. Our strategy samples a set number of candidate plans and selects the highest one that surpasses the score-specific threshold. If no candidate meets the threshold, we fall back to a greedy sampling loop until a path candidate does. We denote this as the acceptable trajectory generation algorithm.

For each trajectory generation cycle, we optimise for a single chosen score using the same planner. We also compute the scores by treating end effector (final joint on the robot's arm) positions over the motion as the trajectory waypoints.

We formally present the acceptable trajectory generation algorithm in pseudocode below (Algorithm 1):

Algorithm 1. Acceptable trajectory generation algorithm

```
function generate_best_plan_from_samples(move_group,
goal_pose, num_samples, eyes_point, threshold, planner)
    move_group.set_pose_target(goal_pose)
    move_group.set_planner(planner)
    best_plan ← None
    best_score ← 0.0
    for i ← 1 to num_samples do
        current_plan ← move_group.plan()
        if current_plan.validity = True then
            current_score ← compute_score(current_plan.plan,
move_group, eyes_point)
            if current_score > best_score then
                best_plan ← current_plan
                best_score ← current_score
    while best_score < threshold do
        current_plan ← move_group.plan()
        if current_plan.validity = True then
            current_score ← compute_score(current_plan.plan,
move_group, eyes_point)
            if current_score > best_score then
                best_plan ← current_plan
                best_score ← current_score
    return best_plan
end function
```

Note that the mentioned compute_score function calculates the score of the relevant score type for the plan or trajectory. The *eyes_point* variable denotes the modelled location of the human observer.

2.4.1. Planners Analysis

To determine the threshold values for each score, as well as the best planner for generating trajectories optimised for each score, a planners analysis is conducted:

- Setup – Fix a start and goal pose. For each planner, generate 1000 valid plans without execution, to keep state fixed;
- Computation – Compute the discretised predictability prime (predictability′), discretised relative legibility, discretised relative trustability, and sociability scores;
- Results – Record per planner, the min, max, and mean scores.

The results are shown in the table below (Table 2.1).

Table 2.1. Planners analysis results for each planner.

Planner	Predictability′			Legibility			Trustability			Sociability		
	Min	Max	Mean	Min	Max	Mean	Min	Max	Mean	Min	Max	Mean
SBL	0.3622	0.9824	0.9740	0.6932	0.9905	0.9855	0.1571	0.5238	0.1580	0.5714	0.6458	0.5718
EST	0.7555	0.9740	0.9677	0.9158	0.9890	0.9853	0.1568	0.3746	0.1574	0.5710	0.6452	0.5714
LBKPIECE	0.6766	0.9645	0.9602	0.8770	0.9852	0.9849	0.1565	0.4057	0.1577	0.5707	0.6414	0.5713
BKPIECE	0.3407	0.9620	0.9522	0.6868	0.9900	0.9840	0.1305	0.5267	0.1600	0.5517	0.6452	0.5720
KPIECE	0.3405	0.9579	0.9468	0.6880	0.9898	0.9834	0.1304	0.5256	0.1607	0.5518	0.6410	0.5721
RRT	0.5431	0.9540	0.9451	0.8098	0.9896	0.9836	0.1299	0.4506	0.1601	0.5513	0.6385	0.5719
RRTConnect	0.4653	0.9504	0.9424	0.7651	0.9895	0.9829	0.1304	0.4787	0.1581	0.5521	0.6447	0.5722
RRTstar	0.4506	0.9846	0.9628	0.7543	0.9863	0.9847	0.1483	0.4852	0.1573	0.5655	0.6442	0.5710
TRRT	0.3457	0.9940	0.9851	0.6776	0.9915	0.9860	0.1369	0.5324	0.1567	0.5590	0.6355	0.5713
PRM	0.9868	0.9889	0.9879	0.9864	0.9865	0.9865	0.1561	0.1564	0.1563	0.5713	0.5715	0.5714
PRMstar	0.9889	1.0000	0.9975	0.9830	0.9880	0.9873	0.1487	0.1596	0.1581	0.5659	0.5738	0.5727

All values are normalised to four decimal places (d.p.). Similarly to the table, where score types (i.e., discretised, relative) are omitted for conciseness, we will also omit specifying the score types in the rest of the section, as the type has already been clarified previously.

The mean metric has been chosen over the median metric as most planners bias towards highly predictable trajectories. This meant that

legible, trustable, and sociable trajectories would be rarer outliers. The mean metric captures the "anomalous" trajectories, whilst the median metric inherently suppresses their impact.

The minimum and maximum values are recorded as the scores for each trajectory and are expected to cluster closely around either the minimum or maximum. If the scores cluster near the maximum, setting a threshold slightly below this value is safe, as new trajectories, including "generic" ones, will likely satisfy this threshold. Conversely, if the scores cluster around the minimum, a high threshold would be infeasible. In this situation, the goal of the threshold value is to filter out all the "generic" trajectories and only accept the more creative trajectories that yield higher scores than usual.

2.4.1.1 Determining the Best Planners

From the results above (Table 2.1), we can conclude that the best planner for each score is as follows:

- Predictability′ – PRMstar;
- Trustability – KPIECE;
- Sociability – RRTConnect.

PRMstar and KPIECE have been chosen as the most predictable and trustable planners respectively, because these planners yielded the highest mean values for predictability′ and trustability respectively.

While PRMstar yielded the highest mean sociability score, its trajectories achieved this largely through high legibility, with trustability remaining low. The trajectories' sociability scores also tightly clustered near the mean. In contrast, RRTConnect not only achieved a high mean but also produced trajectories with significantly higher maximum sociability, indicating balanced optimisation of both legibility and trustability. Hence, RRTConnect was selected as the most sociable planner.

2.4.1.2. Determining the Threshold Values

After further scrutinisation of the results (Table 2.1), the threshold values for each scoring metric have been deduced as follows:

- Predictability′: 0.97 (highly predictable), 0.8 (predictable);
- Trustability: 0.3;
- Sociability: 0.6.

We first determined a threshold value for predictability′ that describes a highly predictable trajectory, which is very close to being a straight line. This was initially set to 0.99, around the mean value for the highly predictable PRM planners, but was adjusted downwards to 0.97 after trials revealed the PRM planners rarely produced trajectories above 0.98 in reality. We then derived a general threshold value for predictability′ of 0.8, representing a predictable trajectory in which the overall path is directed towards the goal. This value was selected as it lies approximately ¾ of the way between the minimum and maximum scores observed across all the planners.

The trustability and sociability thresholds were chosen around midway between their mean and maximum values since the mean values for both scores were relatively close to their respective minimum values. This ensures that non-generic trajectories could be generated reliably within a reasonable time frame.

2.4.1.3. Time Efficiency of Planners

Although not a primary metric, planning time showed clear differences across planners when scaled. Star-based planners (RRTstar, PRMstar) were substantially slower, requiring around 5 s per trajectory, while non-star planners were about 12 times faster. This discrepancy is especially important for human-dependent experiments, where excessive delays in trajectory generation are impractical.

2.4.2. Defining the Different Types of Trajectories

Thus, we formally define predictable, highly predictable, trustable, and sociable trajectories as those that have a corresponding score higher than the specified threshold value.

Definition 4. A predictable trajectory, or a trajectory that is said to be predictable, is one with a (discretised) predictability′ score equal to or more than 0.8.

Definition 5. A highly predictable trajectory, or a trajectory that is said to be highly predictable, is one with a (discretised) predictability′ score equal to or more than 0.97.

Note that all highly predictable trajectories are automatically considered a predictable trajectory.

Definition 6. A trustable trajectory, or a trajectory that is said to be trustable, is one with a (discretised relative) trustability score equal to or more than 0.3.

Definition 7. A sociable trajectory, or a trajectory that is said to be sociable, is one with a sociability score equal to or more than 0.6.

When comparing trajectories, we are referring directly to the trajectory's relevant score. For example, one trajectory is said to be more predictable than another if it has a higher (discretised) predictability′ score.

2.4.3. Execution of Trajectories

We generate different types of trajectories by leveraging the best planner and the threshold value determined for each corresponding score.

2.4.3.1. (Highly) Predictable Trajectory

We apply the acceptable trajectory generation algorithm for the generation of a (highly) predictable trajectory, using the predictability prime (predictability′) score. As established in our planners analysis, the PRMstar planner is chosen, with the threshold set to 0.97. The number of (non-greedy) samples is set to 5, balancing computation time (around 30 s if no fallback to greedy method) and trajectory optimality.

Fig. 2.4 shows the generation of a (highly) predictable trajectory, visualised as a chain of red spheres representing the waypoints. The robot is displayed in its starting configuration, where the location and the direction of the start of the final joint (end effector) is the initial pose. Moreover, the intended goal pose is denoted as a green arrow. The trajectory achieved a predictability′ score of 0.977 and closely matched with the straight-line path between the start and the goal.

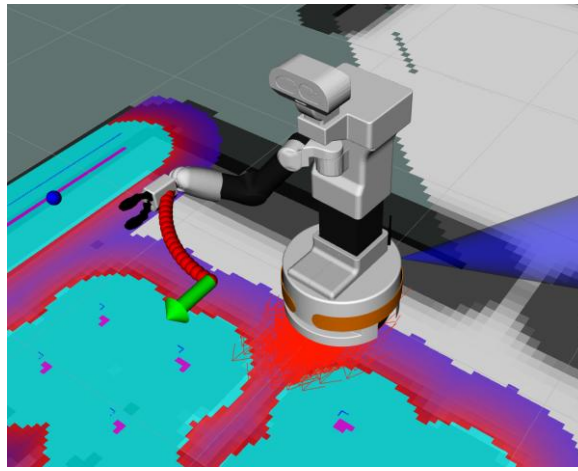


Fig. 2.4. The generation of a (highly) predictable trajectory.

2.4.3.2. Trustable Trajectory

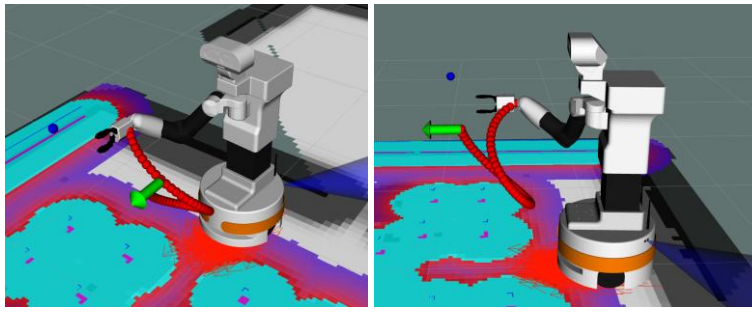
The acceptable trajectory generation algorithm is used to generate a trustable trajectory, where the trustability score is computed. The KPIECE planner is used with a threshold of 0.3 and 60 (non-greedy) samples, exploiting the non-star planner's faster performance to maintain around 30 s of runtime.

Fig. 2.5 shows the resulting trustable trajectory, where the human observer position is denoted by the blue sphere. This trajectory achieves a trustability score of 0.529, which can be identified as one of the highest possible trustability scores from the planners' analysis. This high score is supported by the visualisation of the waypoints, where the initial snippets of the trajectory strongly convey the arm's intended movement away from the human observer.

2.4.3.3. Sociable Trajectory

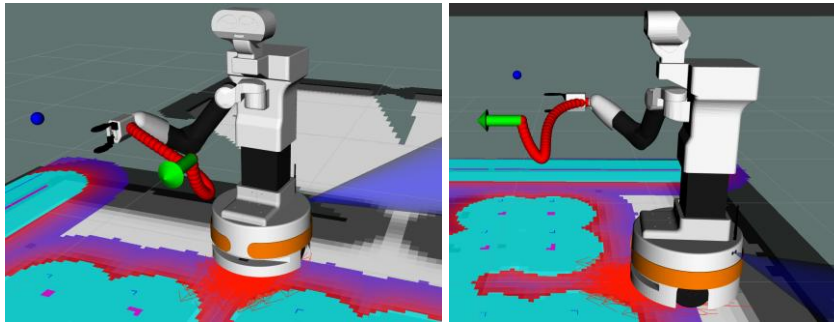
A sociable trajectory is produced by optimising the sociability score with the acceptable trajectory generation algorithm. The threshold is decided to be 0.6 and the planner is set to RRTConnect. The number of (non-greedy) samples is set to 60, to yield a runtime similar to the trustable case.

Fig. 2.6 shows the generated sociable trajectory. A sociability score of 0.645 has been achieved, among the highest obtainable ones from the planners analysis. Early segments of the trajectory in particular can be seen simultaneously conveying intention away from the human observer and towards the actual goal.



(a) (b)

Fig. 2.5. The generation of a trustable trajectory: (a) Front view, (b) Side view.



(a) (b)

Fig. 2.6. The generation of a sociable trajectory: (a) Front view, (b) Side view.

2.4.4. Applying the Trust Region of Predictability – Trustably and Socially Predictable Trajectories

Later experiments show that when only the trustability scores and sociability scores are individually optimised to generate trustable and sociable trajectories, the resulting trajectories are often excessively long

and thus unpredictable, where humans struggle to interpret and understand the robot's intent. Dragan's thesis [33] noted a similar problem when optimising for legibility, where the legibility score is unable to capture how humans make inferences in unpredictable situations. We hence apply Dragan's proposed trust region of predictability in the context of trustability and sociability, to constrain trustability and sociability optimisation to the space where predictability is high enough. By ensuring that the robot conveys its intent efficiently, it would result in more intuitive human understanding of the robot conveying trustable and sociable intent, from motion that can be anticipated. We denote trustable trajectories and sociable trajectories within the trust region as trustably predictable and socially predictable trajectories respectively.

In the context of Yu et al.'s trust model [20], applying the trust region now also encourages the display of competence and integrity for a robot, other elements of trust as shown at later experiments to also hold notable significance. This is because predictable motion makes the robot appear more familiar, goal-oriented, reliable, and behaviourally consistent, thereby reinforcing the perception that the robot is both competent and acting with integrity.

We now formally define trustably and socially predictable trajectories as follows:

Definition 8. A trustably predictable trajectory, or a trajectory that is said to be trustably predictable, is one that is simultaneously trustable and predictable, with a (discretised relative) trustability score equal to or more than 0.3 and a (discretised) predictability′ score equal to or more than 0.8.

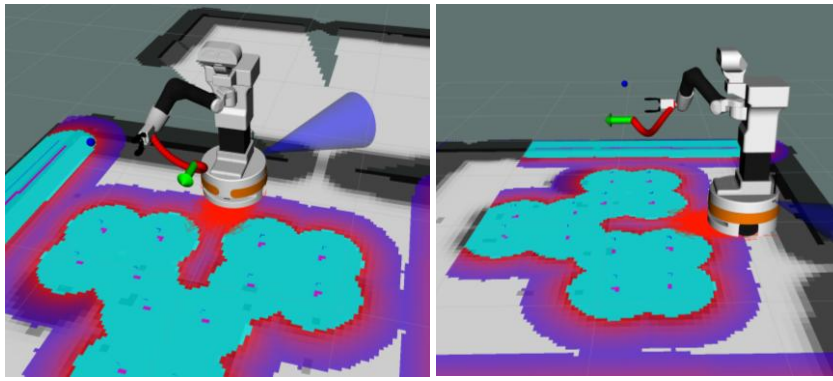
Definition 9. A socially predictable trajectory, or a trajectory that is said to be socially predictable, is one that is simultaneously sociable and predictable, with a sociability score equal to or more than 0.6 and a (discretised) predictability′ score equal to or more than 0.8.

2.4.4.1. Generation of Trajectories

To generate trustably and socially predictable trajectories, we directly employ the greedy approach from the acceptable trajectory generation algorithm and select the first plan that satisfies both conditions (predictability and trustability/sociability thresholds) in the optimisation

process. This is due to the sample space already being heavily reduced from the predictability constraint.

We also utilise the previously stated planners for generating trustable and sociable trajectories, for generating the trustably and socially predictable trajectories respectively. For example, for generating a trustably predictable trajectory, we employ the KPIECE planner, with a threshold of 0.3 for trustability and 0.8 for predictability′. Fig. 2.7 shows such a trajectory.



(a) (b)

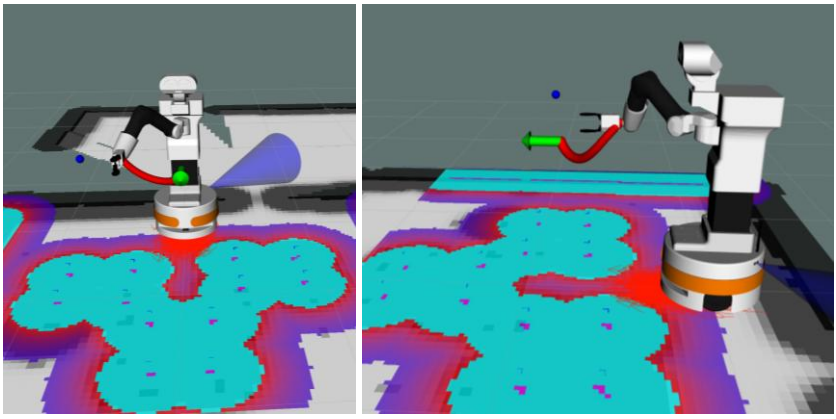
Fig. 2.7. The generation of a trustably predictable trajectory:
(a) Front view, (b) Side view.

For generating a socially predictable trajectory, we use the RRTConnect planner. Fig. 2.8 shows such a trajectory.

2.4.5 Limitations

One of the major limitations of the acceptable trajectory generation algorithm is that the proposed thresholds for determining predictable, trustable, sociable, trustably predictable, and socially predictable trajectories relied on empirically derived values from a specific planner analysis under a specific scenario. Consequently, these thresholds should not be interpreted as universal constants, as their exact values may vary with the robot platform, start-goal configuration, observer position, environmental layout, and other contextual factors. Future work should evaluate the sensitivity of the reported results to variations in these

thresholds and investigate their calibration across a broader range of robots, environments, and human observer configurations. It could also involve establishing psychometrically calibrated decision boundaries, defined as the threshold at which a specified proportion of participants judge a trajectory to be predictable, trustable, or sociable.



(a) (b)

Fig. 2.8. The generation of a socially predictable trajectory: (a) Front view, (b) Side view.

2.5. Experimental Design and Creation of Social Intelligence

Next, we evaluate whether the newly proposed scores (trustability and sociability) align with human expectations and capture their intended characteristics. Two experiment types are conducted:

1) **Human-dependent experiments (user study):** Assess whether participants perceive trajectories with higher scores as more trustable/sociable;
2) **Human-independent experiments (exclusively in simulation):** Test whether the scores respond consistently with their intended characteristics to different human observer and goal configurations.

The final part of this section presents the design of a classifier for recognising socially predictable trajectories. This classifier represents an initial step toward implementing social intelligence in robot motion, by

enabling the robot to identify trajectories that exhibit the proposed social characteristics. Since socially predictable trajectories have later been shown to also be recognisable by human participants, the classifier is intended to emulate, in a computational form, the human ability to identify predictable motion that is socially appropriate.

For clarification, the discretised predictability′, discretised relative legibility, and discretised relative trustability scores are simply referred to predictability′, legibility, and trustability for the rest of the section.

2.5.1. Human-Dependent Experiments (User Study)

Participants interact with the either the real robot, the Gazebo simulation, or a pre-recorded video of the Gazebo simulation. The setup places the human diagonally to the robot's right, seated and facing the robot (setup shown in Fig. 2.9).

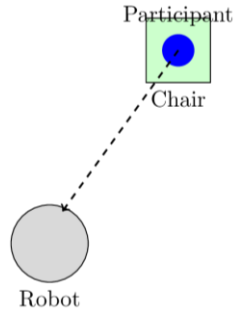


Fig. 2.9. A 2D view of the environment for the human-dependent experiments.

Two types of experiments are conducted, where one evaluates the trustability score, and the other evaluates the sociability score. Each experiment involves a unique participant and consists of two rounds, with each round consisting of five iterations. During each iteration, the participant is shown a pair of trajectories, presented sequentially, and is asked "Which trajectory out of the pair do you think is more sociable/trustable?" with the possible responses being:

- "First trajectory";
- "Second trajectory";
- "Unable to tell" (N/A).

During the first round of the experiment, the participant is not provided with any definition of sociability/trustability, motivating them to judge the trajectories based on their own interpretations. However, before the second round, aside from the formal definitions of sociability/trustability we defined previously, they are also given the more intuitive definitions:

- **Trustability:** "How early or quickly you can confidently identify that the robot's intended goal is not the human observer, i.e., the robot isn't going for you or attacking you";
- **Sociability:** "How fast you are able to confidently infer the robot's intended goal and that this goal is not the human observer, i.e., it's clear where the robot is going for, and it's not going for you".

After each round, the participant explains their reasoning behind choosing the trajectory that was more sociable/trustable.

Each iteration presents a pair of trajectories: a highly predictable baseline and a trustable/sociable alternative, both sharing the same start and goal poses. For trajectories with no trust region of predictability enforced, these are generated through the default acceptable trajectory generation algorithm for the respective score. In contrast, for trustably and socially predictable trajectories generation, we use the specific trajectory generation method defined for trajectories bound by the trust region instead (detailed in Section 2.4.4.1). The highly predictable trajectory is considered the control and is the one that matches the user's typical expectations the most, and thus should be considered as the one that is less sociable or trustable. The order of presentation of the trajectory out of the pair is also randomised to avoid bias.

For experiments with pre-recorded demonstrations, we follow the general guideline described above, but instead show the participants videos of trajectories generated in simulation. It should be noted that we still attempt to simulate the live generation of trajectories as much as possible, so the participant will never be shown the same video over an experiment. We however do enforce that during each iteration, the participant would see a pair of trajectories, where one is the highly predictable baseline, and the other is a trustable/sociable trajectory with a relevant score higher that the corresponding threshold value. The trajectories within the pairs are thus guaranteed to look distinct.

For experiments with live demonstrations, participants could either be asked to judge trajectories generated on the real robot or in simulation.

For these experiments, trajectory generation is adapted to a time-limited sampling approach, where instead of a fixed number of samples, candidates are generated until a one-minute time limit has been reached. The highest-scoring trajectory out of the candidates is selected. This avoids exposing the participants to long waits, while also preventing bias from differences in generation time between highly predictable (common) and sociable/trustable (rarer) trajectories.

Due to the time-limited sampling approach, as well as the slowness of computation on the real robot, the acceptable trajectory generation algorithm may not always produce a trajectory that meets the threshold for trustability or sociability scores ("anomalous" trajectories cannot be generated in time). As a result, the highly predictable trajectory may sometimes appear more trustable or sociable than the one intended to be so. Therefore, we consider the trajectory with the higher score directly in each pair to be the more sociable or trustable one.

It should be noted that only experiments conducted on the real robot using the live demonstration approach directly follow the setup shown in Fig. 2.9. For the experiments conducted in simulation, whether using pre-recorded videos or live trajectory generation, the participants are shown a third-person view, as shown in Fig. 2.10. This view provides a full view of both the trajectory and the experimental environment, since it is difficult to observe the full extent of the trajectory from a first-person perspective. The participants are instructed to imagine themselves seated at the chair position shown in Fig. 2.9, specifically the chair closest to the robot arm, while facing the body of the robot. This corresponds to a diagonal seating position relative to the robot.

2.5.2. Human-Independent Experiments

To test whether the proposed scores behave as intended under different conditions, we evaluate them using a fixed highly predictable trajectory (same start, goal, and similar waypoint structure), while varying only the theoretical goal point and/or human observer location ("eyes point"). We distinguish the goal of the highly predictable trajectory by referring to it as the "goal pose" and the theoretical goal as a "goal point".

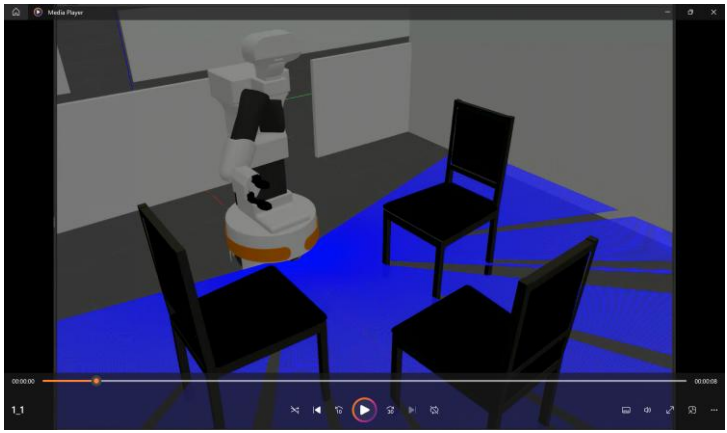


Fig. 2.10. An example of the experimental setup in simulation (pre-recorded demonstrations).

These experiments are designed to compare whether the computed scores align with human expectation of what the scores should be, based on the locations of the human observer and the goal point.

2.5.2.1. Trustability

For the experiment evaluating trustability, the location of the theorised goal point becomes a controlled variable, due to its independence in the determination of the trustability score. The trustability score is determined, given an eyes point. This is the only piece of data required to be noted down for each test case.

The eyes point positions serve as test cases for this experiment. To simplify and ensure fairness, the z-component (height) of each eyes point is kept constant, even though some cases assume a seated observer, while others assume the observer is standing. The test cases are as follows:

1. Sitting on the chair, diagonally to the right of the robot;
2. Sitting on the chair, directly in front of the robot;
3. Sitting on the chair, diagonally to the left of the robot;
4. Standing behind the end effector of the robot (initial pose);
5. Standing behind the theoretical path that forms between the initial pose and goal pose of the highly predictable trajectory;

6. Standing directly at the initial pose (location of eyes point is set as the initial pose);
7. Standing in the way of the highly predictable trajectory (location of eyes point is set as a point in the middle of the trajectory);
8. Standing at the goal pose (location of eyes point is set as the goal pose).

Fig. 2.11 is a visual representation of all of the test cases, denoted as T, where each point represents a possible eyes point. The robot is assumed to be facing downwards, thus T1 is the chair located diagonally to the right of the robot.

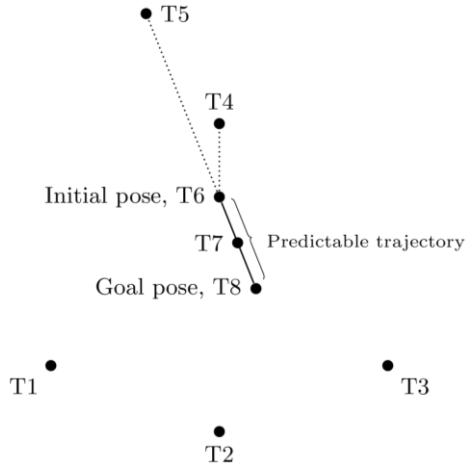


Fig. 2.11. A 2D representation of the trustability test cases.

2.5.2.2. Sociability

In the experiment designed to evaluate sociability, we vary the location of the eyes point (affecting trustability) and the theorised goal point (affecting legibility), both of which directly affects the computation of the sociability score.

The test cases aim to evaluate whether the sociability score has the ability to reflect varying degrees of legibility and trustability. These are stated below:

1. Legible and trustable;
2. Legible and not trustable (untrustable);
3. Not legible (illegible) and trustable;
4. Neither legible nor trustable (illegible and untrustable).

For the simplification of the experiment, we decide to reduce the number of different locations of the eyes point and the theoretical goal point. This is achieved by using the relationship that if the goal point G is the same as the eyes point H, then for a trajectory ξ:

$$\text{legibility}(\xi) = \text{untrustability}(\xi), \tag{2.37}$$

as per the definition of the scores (2.6), (2.28). Due to the notion that untrustability is how "not trustable" a trajectory is, where it is implied that:

$$\text{untrustability}(\xi) \propto \frac{1}{\text{trustability}(\xi)}, \tag{2.38}$$

we can thus conclude that if $G = H$:

$$\text{legibility}(\xi) \propto \frac{1}{\text{trustability}(\xi)} \tag{2.39}$$

Due to this relationship, it is possible to pick a point which simultaneously conveys legibility and untrustability, as well as a point that simultaneously conveys illegibility and trustability, for a known trajectory. Since the generated trajectory is highly predictable, the point that represents the goal pose would be both legible and untrustable. We denote this point as LUT. The eyes point for test case 5 (T5) proposed in the trustability experiment is a good candidate for a point that is both trustable and illegible, as its location would make it appear that the robot conveys clear intent away from both the actual goal and the human observer. We denote this point as ILT. We now define the test cases with their corresponding locations for the eyes point and the goal point below:

1. Goal point – LUT | Eyes point – ILT;
2. Goal point – LUT | Eyes point – LUT;
3. Goal point – ILT | Eyes point – ILT;
4. Goal point – ILT | Eyes point – LUT.

2.5.3. Creation of Social Intelligence – Training a Classifier to Recognise Socially Predictable Trajectories

As a step towards creating social intelligence in robots, we train a binary classifier to identify socially predictable trajectories from raw end effector waypoints over trajectories, without access to any pre-computed scores (e.g., the sociability score). Since socially predictable trajectories would later be shown to be recognisable by the human participants (Section 2.6.1.2), this classifier is intended to emulate, in a computational form, the human intelligence to reliably identify sociable trajectories. Notably, this problem is harder than the one presented to humans in the human-dependent experiments. This is due to the fact that humans were only required to compare between a pair of trajectories, which is a relative judgement aided by a comparison. On the other hand, the proposed classifier is required to determine whether a single trajectory is socially predictable without reference to another trajectory.

2.5.3.1. Dataset

A balanced dataset of 250 trajectories is collected, split evenly between positive examples, those that are socially predictable, and negative examples. The trajectories overall consist of three classes: socially predictable (50 %), sociable but not predictable (sociable only) (25 %), and predictable but not sociable (highly predictable) (25 %). Trajectories for each class are collected using their corresponding planner, where RRTConnect is utilised for socially predictable and sociable only trajectories, and PRMstar for highly predictable trajectories. These trajectories are generated using the proposed method, which has been detailed in Section 2.4.4.1.

Each trajectory is verified against its intended class before being saved, ensuring label integrity. This is due to the possibility of generating a sociable trajectory that also happened to be predictable, and thus socially predictable. The raw end effector waypoints and the eyes point are stored for each trajectory alongside a hard binary label, where socially predictable trajectories are assigned label 1, and all others label 0.

It is also notable that the dataset consists of trajectories all beginning and ending at the same pose. Despite this, the number of waypoints in each trajectory could differ, as trajectories may take longer or shorter paths to reach the goal.

The dataset is split into 80 % training (200 trajectories), 10 % validation (25 trajectories), and 10 % test (25 trajectories), with stratification applied to preserve the class distribution across all splits.

2.5.3.2. Model Architecture

Due to the deep structure consisting of multiple input, output, and hidden layers, deep learning-based models can process data effectively and learn various representations with their strong capacity in dealing with high-dimensional data. In particular, convolutional neural networks (CNNs) are feed-forward neural networks which use filters and pooling layers, and can automatically detect spatial hierarchies of features with relatively few parameters. Thus, we propose a 1D (one-dimensional) CNN as the classifier, which takes the raw waypoint sequence as the main input. This design choice is motivated by the need to preserve the temporal ordering of waypoints, which is fundamental to both the legibility and trustability scores, and indirectly to the sociability score. In particular, earlier waypoints are weighted more heavily via the weighting function, and a model designed to operate on ordered sequences can learn this structure directly from the training signal.

The input to the network is a tensor of shape (4, 100), where the waypoints are padded to a fixed length of 100. Three of the four channels correspond to the x, y, and z end-effector position of each waypoint. The fourth channel is the Euclidean distance from each waypoint to the eyes point, which provides the network with direct access to the geometric signal underlying the scores (e.g., trustability).

The encoder consists of three 1D convolutional layers with 32, 64, and 128 filters respectively, each followed by batch normalisation and ReLU activation. Global average pooling is applied to collapse the sequence dimension, producing a fixed-size representation for each trajectory. A two-layer MLP (multi-layer perceptron) classification head with dropout then produces a binary logit output, trained with binary cross-entropy loss.

2.5.3.3. Model Training

The model is trained in two phases. In the first phase, the model is trained on the training set using the Adam optimiser with a learning rate of 1×10^{-3} and weight decay of 1×10^{-4}. The validation set is used solely

for epoch selection, where the epoch with the lowest validation loss is identified as the optimal number of epochs. Moreover, early stopping is applied with a patience of 20 epochs to prevent overfitting. In the second phase, the model is retrained from scratch on the combined training and validation set for exactly the optimal number of epochs determined. This ensures the final model benefits from the maximum available data while keeping the test set completely untouched throughout both phases.

2.6. Results and Discussions

This section details both the results from the experiments and the discussions. The final part of this section reports the evaluation results from deep learning-based classifier, aiming to recognise a socially predictable trajectory.

2.6.1. Human-Dependent Experiments (User Study)

This subsection shows the results from the user study, conducted to evaluate human perceptions of the proposed scores. Our results in general demonstrate the proficiency of the proposed scores, particularly under the predictability constraint of the trust region. However, the limited participant sample size restricts the generalisability of the conclusions, and the experiments could benefit from the recruitment of more participants in the future.

2.6.1.1. Experiments with Live Demonstrations

For the experiments with trajectories generated live, we conducted a user study focused on evaluating the trustability score. This is for both experiments conducted on the simulator and on the real robot. As a result, the participants would be expected to watch pairs of trajectories generated in real time, then judge which one appears to be more trustable.

The results in the table below (Table 2.2) show how many of the pairs of the trajectories the participants were able to correctly identify as the more trustable one. The last column records the participant's definition of trustability before they were informed of our definition of trustability in the second round. The rest of the tables in Section 2.6.1 share a similar structure.

Table 2.2. A summary of the findings from the trustability experiment (live demonstrations).

Participant	Platform	Round 1 (Participant's Own Definition)	Round 2 (Participant Given Definition)	Participant's Own Definition of Trustability in First Round
P1	Real robot	1/3 correct (2 N/A)	4/5 correct	Legibility
P2	Real robot	0/3 correct (2 N/A)	3/4 correct (1 N/A)	Predictability ("Smoothness")
P3	Simulation	1/1 correct (4 N/A)	1/3 correct (2 N/A)	Benevolence
P4	Simulation	1/5 correct	4/4 correct (1 N/A)	Legibility

The results show that the participants who initially did not align with our definition of trustability experienced a significant improvement in their ability to identify the more trustable trajectory once they were informed of our definition. This is evident in the increased number of correctly identified trajectories, as well as the decrease in instances where participants were unable to distinguish between the two trajectories (reported as N/A). These findings suggest that our encoding of the robot's benevolence within the trustability score aligns well with human perception, by effectively and proficiently capturing this important social quality.

Moreover, participant 1 also reported that they determined the more trustable trajectory in round 2 based on benevolent movement cues such as the initial downward dip of the end effector. This subtle acknowledgement of the human's presence suggests that trajectories with higher trustability scores may implicitly convey social signals of benevolence, highlighting the score's potential to capture such cues.

However, the poor performance throughout the first round of most experiments also demonstrates that it was not always intuitive for the participants to associate benevolence as a core component of trustability. This finding has been reinforced by the fact that only one out of the four participants in the first round judged trustability based on our definition – benevolence. Participants 1 and 4 also reported that they viewed competence and integrity to be important factors in determining perceived trust in the context of the trust model [20], as they emphasised the importance of the robot displaying clear intent towards the goal in perceiving trust. Participant 2 also shared similar views in the other components of trust being significant, as the preference for more predictable trajectories signalled the desire for more consistent and reliable behaviour.

As previously discussed, the time-limited sampling approach in generating trajectories would result in trajectories intended to be trustable, to not be truly trustable, where the trustability score of the trajectory falls below the trustability score threshold of 0.3. In practice, the participants were mostly comparing between two highly predictable trajectories, to determine which one was more trustable. This was especially the case for the experiments conducted on the real robot, where computation is slower, resulting in a lower chance of a truly trustable trajectory being generated. Overall, this is reflected in many studies, where the pair of trajectories was too similar to each other. This problem could be resolved by using pre-recorded demonstrations, which will be discussed next.

2.6.1.2. Experiments with Pre-Recorded Demonstrations

As mentioned previously, by leveraging pre-recorded videos of trajectories in experiments, the intended trustable/sociable trajectories can be guaranteed to be over their defined thresholds for their respective scores. Since we have greater control of the trajectories shown to the participants, we could fall back to the default acceptable trajectory generation algorithm, which would result in generated trustable/sociable trajectories with scores close to the maximum values, as determined from the planners analysis. This would allow the participants to be more confident with their decisions, resulting in very few reports that the participants could not distinguish between the pairs of trajectories (reported as N/A), as the two trajectories within each pair are now very distinct.

The results below (Table 2.3 and Table 2.4) from the trustability and sociability experiments reveal a new problem, where optimising for solely the trustability or the sociability score resulted in unpredictable trajectories, where the participants could not understand any intent from the robot, let alone trustable or sociable intent. This has been demonstrated from the poor results from both experiments. We hence conduct experiments with the trust region of predictability applied, where we show the participants trustably and socially predictable trajectories. The results from these experiments show a great improvement in the humans' ability to recognise the trustable/sociable trajectories.

Trustability

These results were initially surprising, as even though we have significantly increased the trustability score of the intended trustable

trajectory compared to the ones in the live demonstrations, the participants performed worse. Not only did most participants struggle to identify the trustable trajectory in the first round, this was also mostly replicated in the second round, even after they were given our definition of trustability.

Table 2.3. A summary of the findings from the trustability experiment (pre-recorded demonstrations).

Participant	Round 1 (Participant's Own Definition)	Round 2 (Participant Given Definition)	Participant's Own Definition of Trustability in the First Round
P5	1/5 correct	0/5 correct	Safety
P6	0/5 correct	0/5 correct	Predictability/Efficiency
P7	3/5 correct	5/5 correct	Predictability
P8	0/5 correct	2/5 correct	Predictability

However, the way most participants perceived trustability (in the first round) reveal the reasoning processes behind the poor performance. Even though the trustable trajectories had a high trustability score (on average around 0.48), their predictability′ scores were very low (on average around 0.3). The participants in general reported that they struggled to understand the robot's intent from these highly unpredictable trajectories, and thus decided to choose the highly predictable trajectory as the one they could trust. Although excessive movements from the unpredictable trajectories provided more opportunity to signal trustable intent (resulting in a higher trustability score), they also crucially hindered the perception of trust. This was from clear signals of the robot's incompetence, which is another key component in the trust model [20]. This explains the disappointing performance in the second round, even after the participants were told our definition of trustability.

We could also now understand why the participants performed worse in this experiment compared to the one with live demonstrations. Even though the pairs of trajectories were very similar in the live demonstrations, and there rarely existed a trajectory that was truly trustable, both of the trajectories within the pairs were however both highly predictable. As a result, the trust region [33] was not violated, and the participants could sometimes interpret trustable intent from certain

trajectories, despite being minor. This motivates using trustably predictable trajectories for generating trustable motion, as these would convey trustable intent that could be interpreted.

Sociability

Table 2.4. A summary of the findings from the sociability experiment (pre-recorded demonstrations).

Participant	Round 1 (Participant's Own Definition)	Round 2 (Participant Given Definition)	Participant's Own Definition of Sociability in the First Round
P9	0/5 correct	0/5 correct	Humanness/Predictability
P10	0/5 correct	0/5 correct	Predictability
P11	1/5 correct	0/5 correct	Humanness
P12	0/5 correct	0/5 correct	Humanness/Predictability

The sociable trajectories also suffered from a similar problem, where the sociability scores were very high (at around 0.65) and the predictability′ scores were very low (at around 0.4). Predictability seems to be the crucial limiting factor in the motion's socialness, as the participants performed even worse in the sociability experiment than the trustability experiment, where the number of sociable trajectories correctly identified in total is one, over all four participants. This time, the participants also noted a movement present in all sociable trajectories that caused them to determine them as not sociable. These trajectories all involved the end effector rotating more than 360 degrees during motion, which most humans concluded to not only be unpredictable and difficult to understand, but also physically impossible to be replicated on humans, and thus inhuman. Although in the context of the sociability score, these movements provided more chances in conveying sociable intent, in directing the general motion away from the human observer and towards the goal, they nonetheless heavily violated the perceived socialness by exhibiting unnatural behaviour which confused the human observers.

The participants also reported the definition of sociability to be complicated and not intuitive, meaning that even after being informed of our definition of sociability, the participants still struggled to have a clear understanding of which trajectory we intended to be social. This is reflected from no improvement in performance in the second round

across all the participants. Despite this, it is not considered a major concern, as our main goal is to propose scores that implicitly captured social characteristics, not definitions that intuitively described how we capture these characteristics.

All of this also motivates employing socially predictable trajectories for generating sociable motion that humans could interpret more intuitively.

Trustably Predictable Trajectories

The results shown in Table 2.5 demonstrate that the participants were much more competent in choosing the trustable trajectory when they are trustably predictable, as opposed to only being trustable. Similarly to the experiment conducted on live demonstrations, the participants performed better in general after being informed of our definition of trustability, which demonstrates that trustability does encode the described characteristics well. Since there were no reports of N/A, we can also conclude that the participants were more confident in identifying trustably predictable trajectories.

Table 2.5. A summary of the findings from the trustably predictable trajectories experiment (pre-recorded demonstrations).

Participant	Round 1 (Participant's Own Definition)	Round 2 (Participant Given Definition)	Participant's Own Definition of Trustability in the First Round
P13	1/5 correct	5/5 correct	Predictability
P14	2/5 correct	5/5 correct	Predictability ("Stability")
P15	4/5 correct	3/5 correct	Inhumanness/Comfort

In general, the participants noted predictability to be an important factor in determining perceived trust. This justifies our approach of applying the trust region of predictability to enforce a hard constraint on how predictable a trajectory must first be, in order to encode the importance of competence and integrity.

Socially Predictable Trajectories

The results in Table 2.6 reveal very significant improvements from the sociability experiment.

Table 2.6. A summary of the findings from the socially predictable trajectories experiment (pre-recorded demonstrations).

Participant	Round 1 (Participant's Own Definition)	Round 2 (Participant Given Definition)	Participant's Own Definition of Sociability in the First Round
P16	3/5 correct	5/5 correct	Positive impression ("Prettiness")
P17	4/5 correct	2/5 correct	Humanness/Intelligence
P18	3/5 correct	2/5 correct	Inhumanness/Comfort

Since most participants were able to reliably and intuitively identify the sociable trajectory in the first round of the experiment using their own interpretation, this shows that socially predictable trajectories are proficient in encoding how humans naturally perceive socialness in robot motion.

Although different participants gave different reasoning processes for their choice, the fact that they were all able to confidently choose the socially predictable trajectory demonstrates that these trajectories indeed exhibit universal social behaviour seen in interactions. This also highlights the significance of our contribution, as the definition of sociability has been shown to be very subjective, and can even be interpreted in opposing ways, with some participants associating it with more human-like behaviour and others with less human-like behaviour. Nonetheless, our proposal of socially predictable trajectories could be widely accepted by humans with diverse views of sociability.

2.6.2. Human-Independent Experiments

2.6.2.1. Trustability

The table below (Table 2.7) displays different trustability scores captured for the eight different locations of the human observer (eyes point).

The trustability scores fall within the expected range in each test case, thus signifying the trustability score does capture its described characteristics well.

Table 2.7. The results for the evaluation of trustability in different environments.

Test Case (T)	Trustability (3 d.p.)
1	0.155
2	0.023
3	0.049
4	0.292
5	0.310
6	0.264
7	0.089
8	0.028

Analysis by groups

- **Seats in front of robot (T1–T3).** Trustability was low overall, which is expected since the arm is moving forwards. The early segment of the trajectory headed most directly towards T2, yielding the lowest score. T1 yielded the highest score as distance to the point from the trajectory grew consistently over the trajectory. Fig. 2.12 and Fig. 2.13 give examples of test case 2 and 3 respectively.

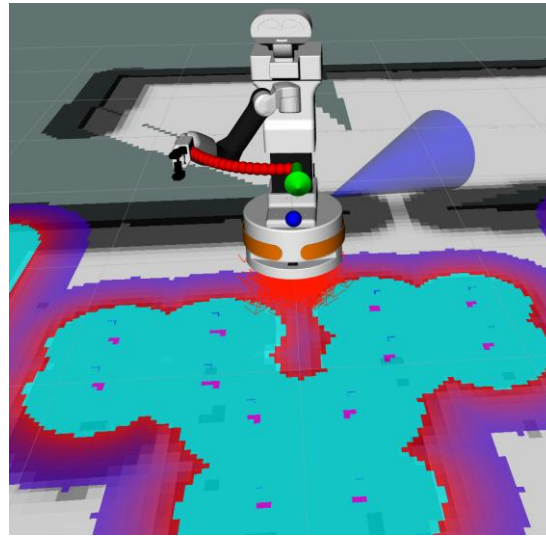


Fig. 2.12. Experiment where test case 2 is evaluated for trustability.

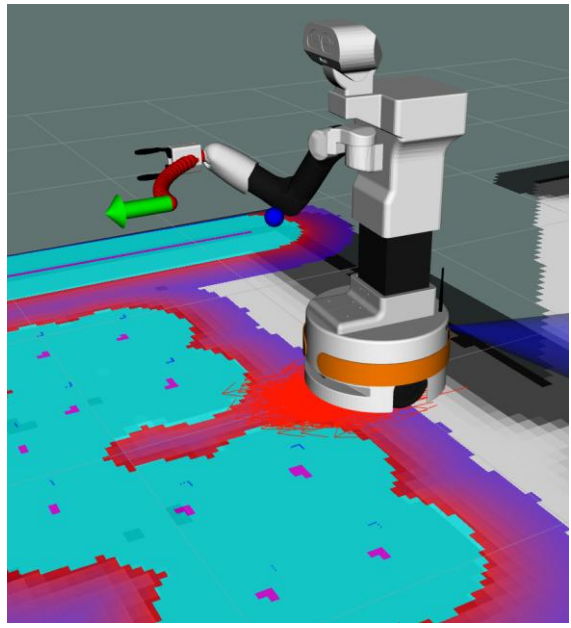


Fig. 2.13. Experiment where test case 3 is evaluated for trustability.

- **Points behind the robot (T4–T5).** The group exhibited high trustability overall, as the trajectory indicated a general movement away from these points. As expected, a point directly opposite the motion direction (T5) scored slightly higher than one merely behind the initial pose (T4).
- **Theoretical points on trajectory (T6–T8).** T6 was the most trustable since the entire trajectory moved away from the initial pose. T8 was the least trustable because the entire trajectory acted towards it. T7 had a relatively low trustability score since the trajectory's conveyed intent was clearly towards the point in the earlier segment of the trajectory.

Cross-group analysis

- **Away-from-initial-point vs. away-from-observer:** T4–T5 were more trustable than T6, suggesting motion can be trusted more if the human observer is off the trajectory completely, as opposed to being at the initial pose.
- **Perspective matters:** The highly predictable trajectory was not perfectly predictable, meaning there was a slight curve. This unexpectedly caused T2 to be the least trustable instead of T8 (goal),

as the early curve in the trajectory caused more benevolence to be conveyed towards the goal.

- **Abnormally long and unpredictable trajectory required for high trustability:** Even the most trustable configuration (T5) barely exceeded the threshold and remained far below the upper range from the planners analysis. This shows the score rewards longer, less predictable paths that create more opportunities to show intent away from the human observer, while short, highly predictable motion that consistently moves away is undervalued. This bias has been reduced by employing the trust region of predictability in generating trustably predictable motion.

2.6.2.2. Sociability

Sociability was evaluated across four test cases combining different levels of legibility and trustability, using a legibility weight of 0.5. The test cases varied the positions of the goal point (affecting legibility), as well as the human observer's eyes point (affecting trustability). Fig. 2.14 gives an example of test case 1. The results are shown in the table below (Table 2.8).

Table 2.8. The results for the evaluation of sociability in different environments.

Test Case (T)	Legibility (3 d.p.)	Trustability (3 d.p.)	Sociability (3 d.p.)
1	0.985	0.301	0.643
2	0.985	0.026	0.506
3	0.698	0.298	0.498
4	0.698	0.026	0.362

Overall, the results demonstrate that the sociability score, using the default legibility weight of 0.5, fairly captures variations in both legibility and trustability values. This is evident since test cases 2 and 3 yield similar sociability scores, and the differences between test cases 1 and 2/3, and between test cases 4 and 2/3 are around the same, indicating that neither score (legibility or trustability) disproportionately influences the sociability score. Although the magnitude of the legibility score for legible trajectories is higher than the trustability score for trustable trajectories, their respective ranges remain consistent. As the sociability score primarily serves as a comparative metric between trajectories,

maintaining equal ranges for legibility and trustability scores is sufficient to ensure fair weighting. Furthermore, this consistency indirectly confirms that legibility and trustability measure comparable dimensions of intent expression, eliminating the need for additional weighting adjustments in calculating sociability. Despite this, the equal weighting remains a design choice rather than one derived from an explicit utility or risk model. Alternative linear and nonlinear formulations, including Pareto-based, maximum penalty, or risk-constrained approaches, may better capture interactions between legibility and trustability. Future work should therefore examine and validate different weighting and aggregation schemes.

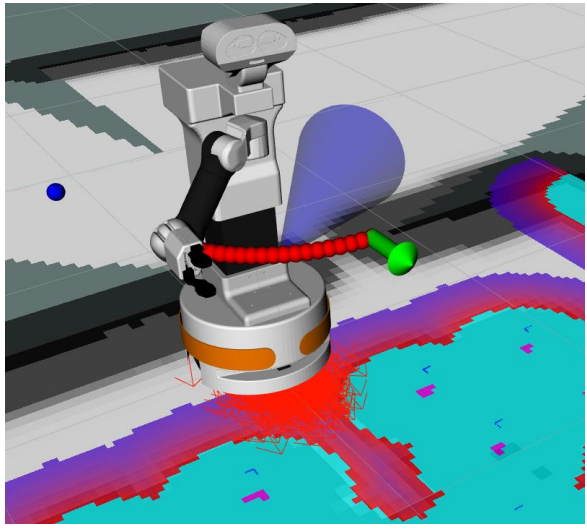


Fig. 2.14. Experiment where test case 1 is evaluated for sociability.

However, it should be noted that these conclusions are drawn from the assumption of a highly predictable trajectory, which constrains the range of the legibility and trustability scores. Since sociability is a measure of the proficiency of robot intent expression, evaluating the effectiveness of the sociability score in highly unpredictable trajectories is not only difficult, but also not very meaningful, as robot intent becomes difficult or even impossible to interpret under highly unpredictable conditions. This also justifies the employment of the trust region of predictability in generating socially predictable trajectories.

2.6.3. Creation of Social Intelligence – Training a Classifier to Recognise Socially Predictable Trajectories

As shown, humans are capable of identifying sociable motion from socially predictable trajectories. Therefore, successfully building a binary classifier that demonstrates the same ability could thus be concluded to show the presence of some form of social intelligence.

2.6.3.1. Deep Learning-Based Model Training

The training and validation curves below (Fig. 2.15) show well-behaved learning in the first phase of training. Both train and validation loss decrease rapidly in the first 10 epochs and stabilise below 0.1 thereafter. The validation loss is noisier than the training loss, which is expected given the small validation set of 25 samples, where a single misclassification would produce a visible spike. Importantly, there is no divergence between training and validation loss, indicating no significant overfitting. The accuracy curves also reflect the same findings. The best epoch was identified as epoch 93 with a validation loss of 0.0051, at which point early stopping was triggered at epoch 113.

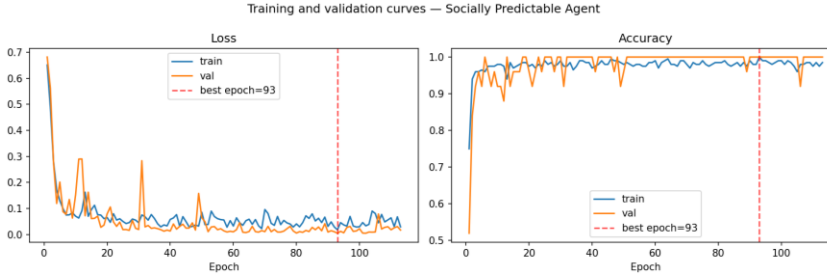


Fig. 2.15. Training and validation curves.

We thus retrained the model with 93 epochs after combining the training and validation datasets, and saved the resulting model for evaluation.

2.6.3.2. Evaluation

The classifier achieves perfect classification on the held-out test set of 25 trajectories, with accuracy, precision, recall, F1-score, ROC-AUC,

and average precision all equal to 1.0. The confusion matrix confirms zero misclassifications in either direction, where all 12 socially predictable and all 13 non-socially predictable trajectories were correctly identified. This shows that the decision boundary between socially predictable and non-socially predictable trajectories is distinct and easily learnable.

Table 2.9 shows the predicted probability that each trajectory in the test set is socially predictable, grouped by the trajectory's true class. This could be understood as how confident the classifier is that a presented trajectory is indeed socially predictable. Despite being trained on hard binary labels, the classifier assigns varying probabilities across the three trajectory classes.

Table 2.9. The predicted probability of being socially predictable on the test set.

Class	Number of Trajectories	Mean (4 d.p.)	Standard Deviation (4 d.p.)	Median (4 d.p.)
Socially predictable	12	0.9649	0.0646	0.9926
Sociable only (not predictable)	9	0.0007	0.0009	0.0004
Highly redictable (not sociable)	4	0.0023	0.0034	0.0004

As expected, the mean and median probabilities for the socially predictable class are close to 1.0, signifying the classifier's confidence in identifying a truly socially predictable trajectory. The near 0 probabilities for the negative classes also demonstrate the classifier's confidence in correctly identifying a trajectory as not being socially predictable. However, the fact that the sociable only class receives a lower mean probability than the highly predictable class shows that predictability contributes more to being socially predictable. This is expected as a trajectory that is sociable but not predictable violates the trust region constraint, making it geometrically more distinct from being socially predictable, than a trajectory that is (highly) predictable but not sociable.

Overall, these results demonstrate that socially predictable trajectories, as formalised through the sociability and predictability′ thresholds, constitute a geometrically learnable property of robot trajectories. The

deep learning-based classifier successfully replicates the formal definition from raw waypoints alone, without access to any pre-computed scores, representing a computational mechanism for recognising socially predictable motion, and thus characteristics of social behaviour.

2.7. Future Work

As mentioned previously, for future work, we can integrate the proposed sociability score as a reward function in the context of a deep model-based reinforcement learning approach, such as Dreamer [15, 16]. This section begins with the justification of using a Dreamer-style architecture, and then focuses on discussing the modifications that can be made on the Dreamer architecture to adapt to the HRI context of robot motion.

2.7.1. Advantages of Dreamer

Dreamer is particularly suitable for this purpose because socially intelligent robot behaviour must be learned in environments characterised by high-dimensional sensory observations, continuous action spaces, complex human-robot dynamics, and limited opportunities for safe physical exploration, where action execution is expected to be in real-time. It learns a latent world model from collected experience, where it is then possible to use trajectories imagined within this model to train an actor-critic policy [15]. During execution, the trained policy selects an action based on the current latent state inferred from the robot's observations.

Unlike model-free reinforcement learning methods, which generally require the agent to learn primarily through repeated interactions with the actual environment, Dreamer learns a compact latent world model from the robot's previous observations and actions stored in a replay buffer, while the policy is trained using trajectories generated through latent imagination [15, 37]. Consequently, multiple possible robot actions and their predicted robot motion trajectories, and thus the social consequences (through the sociability model), can be evaluated without requiring every candidate behaviour to be executed around human participants. This characteristic has the potential to improve sample efficiency, reduce the amount of costly real-world interaction data

required, and limit human exposure to unsafe, unpredictable, or socially inappropriate exploratory actions.

Furthermore, DreamerV2 has demonstrated that a latent world model can represent the complex continuous dynamics of a simulated humanoid robot and support behaviours such as standing and walking from visual observations [38]. Although this does not by itself establish the suitability of Dreamer for social HRI, it provides evidence that latent world models can represent high-dimensional embodied dynamics. This capability is important for socially interactive robots, whose behaviour may depend on multimodal observations of human movement, interpersonal distance, environmental context, and the robot's own physical state.

Once the actor has been trained using imagined trajectories generated by the learned world model, it can select actions directly from the inferred latent state without requiring computationally expensive online trajectory optimisation (planning). This property could support the low-latency action selection required in real-time and dynamic human environments.

Nevertheless, the safety and social acceptability of the resulting behaviour would depend on the accuracy of the learned world model, the safety of the learned policy, and the validity of the sociability reward. These challenges could be addressed through uncertainty-aware world modelling, explicit safety constraints, and preference-based reinforcement learning. This motivates the need to modify the Dreamer architecture.

2.7.2. Modified Dreamer Architecture

The proposed architecture will use deep learning to learn both the world model and the policy [38]. The world model will be based on a Recurrent State-Space Model (RSSM), consisting of an encoder, a dynamics model, a reward model, and a decoder. The world model learning procedure is displayed in Fig. 2.16. First, modality-specific encoders will process and fuse high-dimensional observations of the human, with the robot's internal state, including its joint angles and previous movements [39]. These observations will then be transformed into a compact latent-state representation. The robot's actions will include physical body or arm movements.

Building on this representation, the latent dynamics model will be learned through self-supervised learning from collected interaction trajectories stored in the replay buffer. This model will predict how the latent interaction state changes in response to the robot's actions. To learn the 'oracle' reward model, the proposed sociability score will first be used to estimate the reward associated with each robot motion trajectory. This provides a dense reward signal and helps address the sparse-reward problem. However, because the sociability score may not fully capture human perceptions, a residual reward will be learned through preference-based reward learning using human preferences over pairs of interaction trajectories. Incorporating the residual reward will refine the sociability score by capturing aspects of social acceptability that are not represented by the proposed sociability model. Therefore, the 'oracle' reward combines both the reward derived from the sociability score and the residual reward. To learn the latent reward model, we predict the reward in latent space from the 'oracle' reward, by using a DNN (deep neural network).

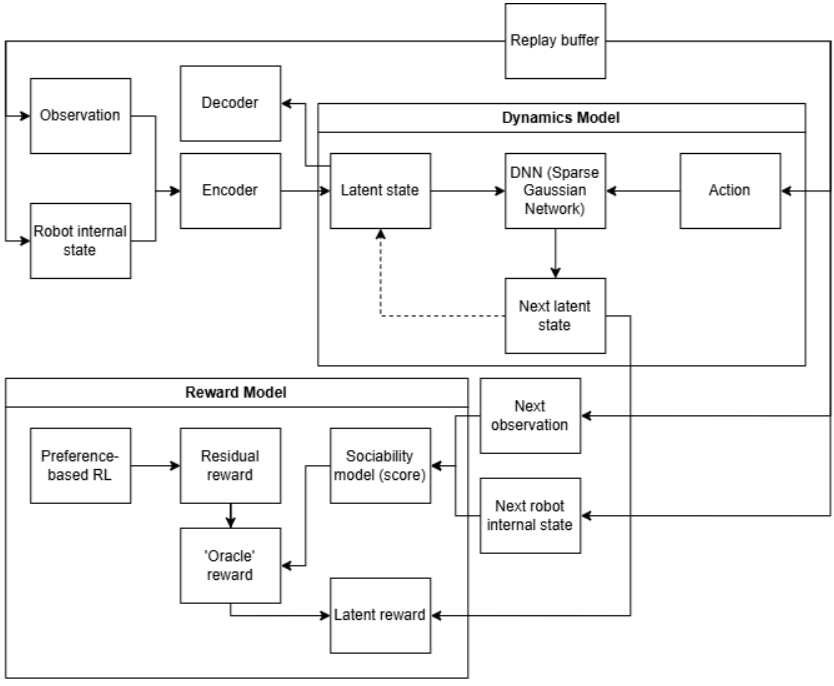


Fig. 2.16. Learning procedure of the world model.

As the policy will map the learned latent state to the robot's physical actions, it thus will be learned using an actor–critic method. Specifically, the critic predicts the expected discounted return of a latent state and the actor maximises the critic's prediction through policy-gradient updates, enabling the policy to optimise task performance and perceived sociability. During training, the dynamics model predicts ("imagines") future latent states under actions proposed by the actor, where the critic estimates the expected discounted return of each latent state using the predicted rewards. An overview of the policy-learning loop is illustrated in Fig. 2.17. Importantly, the use of a learned actor is preferred to online planning because it enables low-latency action selection for real-time interaction during deployment.

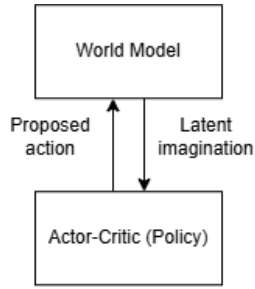


Fig. 2.17. Learning procedure of the policy.

To ensure that the learned policy is safe for interaction with humans, a safe model-based reinforcement learning approach, such as SafeDreamer's background safety-reward planning with Lagrangian relaxation method, BSRP-Lag, could be adopted [40, 41]. A cost model will represent safety constraints associated with each interaction, including minimum human-robot distance, proximity to obstacles, excessive motion velocity, and potential collisions. The safe actor will then be trained to maximise the expected reward and policy entropy while minimising the expected safety cost. An Augmented Lagrangian method could be used to enforce the specified cost constraints during policy learning [42].

Moreover, Gaussian-process-inspired probabilistic modelling could be integrated into the latent dynamics model through a fast sparse Gaussian network [14, 43] for real-time processing and fast learning (as shown in Fig. 2.16). This enables the robot to predict not only the next states, but

also their uncertainty about those predictions, resulting in an uncertainty-conditioned policy, where the robot adapts its behaviour depending on the confidence. This is essential for reliable and safety-critical decision making, as robots can identify unfamiliar situations and respond more cautiously, when the policy is trained over a distribution of "imagined" rollouts [44, 45]. Furthermore, Gaussian process models naturally distinguish between aleatoric and epistemic uncertainty, enabling effective learning in low-data regimes, which reduces the amount of interaction data required for training [46]. These properties are particularly valuable in social HRI, where data collection is costly, human behaviour and the environment are variable, and inappropriate actions during exploration can have significantly negative consequences.

Therefore, during deployment, the robot can use the learned world model to infer its current latent state and the learned policy to select a suitable action in real-time as the interaction evolves, as illustrated in Fig. 2.18.

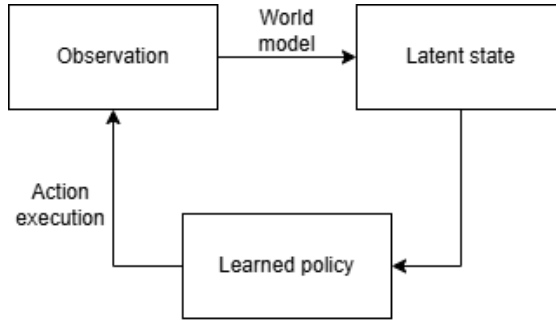


Fig. 2.18. Use of world model and learned policy during deployment.

2.8. Conclusions

This chapter presents a novel method to model socialness in the context of human–robot interactions (HRIs), through deriving a sociability score which is characterised by the legibility and trustability of robot motion. By modelling the importance of benevolence and the establishment of initial trust in fostering long-term trust, it has enabled the quantification of sociability with the integration of both legibility and trustability by deriving a procedure for score computation. An acceptable trajectories generation algorithm is also developed, where the scores of the sociable

and trustable trajectories would exceed the defined threshold values. Furthermore, to generate predictable trajectories, a modified predictability score known as the predictability prime (predictability′) score has been developed to evaluate the predictability of a trajectory independent of others. By applying the trust region of predictability, trustably and socially predictable trajectories are thus generated with social characteristics that humans can identify and interpret reliably. The experiments conducted have demonstrated the effectiveness of the proposed method, with the modelling of sociability and trustability in robots validated at the same time. Furthermore, through training a deep learning-based classifier to identify socially predictable trajectories, we aim to make significant strides in giving robots social intelligence by replicating humans' ability to recognise such complex motion in real-time and dynamic environments. For future work, we will leverage a deep model-based reinforcement learning approach which utilises our proposed scores to generate socially acceptable trajectories in real-time and dynamic environments, where hierarchical sensory processing system based on a fast deep learning model can be deployed to build more intelligent and autonomous robots.

Acknowledgements

Technical support and insightful discussions from Dr. Matteo Leonetti and his team at King's College London are acknowledged.

References

[1]. J. T. Licardo, M. Domjan, T. Orehovacki, Intelligent robotics – A systematic review of emerging technologies and trends, *Electronics*, Vol. 13, Issue 3, 2024, 542.

[2]. S. Behnke, Humanoid robots - From fiction to reality?, *Künstliche Intelligenz*, Vol. 22, Issue 4, 2008, pp. 5-9.

[3]. J. Amirian, M. Abrini, M. Chetouani, Legibot: Generating legible motions for service robots using cost-based local planners, *arXiv*, 2024, arXiv:2404.05100.

[4]. C. Lichtenthäler, A. Kirsch, Legibility of robot behavior: A literature review, *HAL*, 2016, hal-01306977.

[5]. D. L. McPherson, D. R. Scobee, J. Menke, A. Y. Yang, et al., Modeling supervisor safe sets for improving collaboration in human robot teams, in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018, pp. 861-868.

[6]. B. Busch, J. Grizou, M. Lopes, F. Stulp, Learning legible motion from human–robot interactions, *International Journal of Social Robotics*, Vol. 9, Issue 5, 2017, pp. 765-779.
[7]. L. Naendrup-Poell, L. Onnasch, Predictive robot eyes enhance attentional guidance in cooperative human-robot interaction, *Scientific Reports*, Vol. 15, Issue 1, 2025, 32661.
[8]. Z. R. Khavas, A review on trust in human-robot interaction, *arXiv*, 2021, arXiv:2105.10045.
[9]. S. Thompson, K. Candon, M. Vazquez, The social context of human-robot interactions, *arXiv*, 2025, arXiv:2508.13982.
[10]. H. Luo, J. Zhang, Y. Luo, G. Luo, Sociability modelling in robot motion by integrating legibility and trustability, in *Proceedings of the 6th Winter IFSA Conference on Automation, Robotics & Communications for Industry 4.0/5.0/6.0 (ARCI)*, 2026, pp. 28-34.
[11]. H. Luo, J. Zhang, Y. Luo, G. Luo, Sociability modelling in robot motion for generating socially predictable trajectories, *Machines*, Vol. 14, Issue 7, 2026, 718.
[12]. C. Breazeal, Toward sociable robots, *Robotics and Autonomous Systems*, Vol. 42, Issue 3-4, 2003, pp. 167-175.
[13]. T. Fong, I. Nourbakhsh, K. Dautenhahn, A survey of socially interactive robots, *Robotics and Autonomous Systems*, Vol. 42, Issue 3-4, 2003, pp. 143-166.
[14]. G. Luo, H. Luo, J. Zhang, Y. Luo, Real-time deep learning for leaks detection and demands prediction of compressed medical gas systems, in *Proceedings of the Intelligent Systems Conference (IntelliSys)*, 2023, pp. 503-518.
[15]. D. Hafner, T. Lillicrap, J. Ba, M. Norouzi, Dream to control: Learning behaviors by latent imagination, *arXiv*, 2020, arXiv:1912.01603.
[16]. D. Hafner, J. Pasukonis, J. Ba, T. Lillicrap, Mastering diverse control tasks through world models, *Nature*, Vol. 640, 2025, pp. 647-653.
[17]. A. D. Dragan, K. C. Lee, S. S. Srinivasa, Legibility and predictability of robot motion, in *Proceedings of the 8th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 2013, pp. 301-308.
[18]. G. Csibra, G. Gergely, 'Obsessed with goals': Functions and mechanisms of teleological interpretation of actions in humans, *Acta Psychologica*, Vol. 124, Issue 1, 2007, pp. 60-78.
[19]. D. H. McKnight, N. L. Chervany, Trust and distrust definitions: One bite at a time, in Trust in Cyber-Societies (C. Castelfranchi, Y.-H. Tan, Eds.), *Springer*, Berlin, 2001, pp. 27-54.
[20]. M. Yu, M. Saleem, C. Gonzalez, Developing trust: First impressions and experience, *Journal of Economic Psychology*, Vol. 43, 2014, pp. 16-29.
[21]. P. A. Hancock, D. R. Billings, K. E. Schaefer, Can you trust your robot?, *Ergonomics in Design*, Vol. 19, Issue 3, 2011, pp. 24-29.
[22]. I. Asimov, Runaround, *Astounding Science Fiction*, Vol. 29, Issue 1, 1942, pp. 94-103.

[23]. K. Riezler, The social psychology of fear, *American Journal of Sociology*, Vol. 49, Issue 6, 1944, pp. 489-498.
[24]. M. Körber, E. Baseler, K. Bengler, Introduction matters: Manipulating trust in automation and reliance in automated driving, *Applied Ergonomics*, Vol. 66, 2018, pp. 18-31.
[25]. J. Reinhardt, A. Pereira, D. Beckert, K. Bengler, Dominance and movement cues of robot motion: A user study on trust and predictability, in *Proceedings of the IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, 2017, pp. 1493-1498.
[26]. R. Campagna, A. Mislin, K. Dirks, H. Elfenbein, The (mostly) robust influence of initial trustworthiness beliefs on subsequent behaviors and perceptions, *Human Relations*, Vol. 75, Issue 7, 2021, pp. 1383-1411.
[27]. J. E. Hermosilla-Diaz, E. J. Rechy-Ramirez, A. Marin-Hernandez, Robotic motion techniques for socially aware navigation: A scoping review, *Future Internet*, Vol. 17, Issue 12, 2025, 552.
[28]. T. Kruse, A. K. Pandey, R. Alami, A. Kirsch, Human-aware robot navigation: A survey, *Robotics and Autonomous Systems*, Vol. 61, Issue 12, 2013, pp. 1726-1743.
[29]. E. T. Hall, The Hidden Dimension, *Doubleday*, New York, 1966.
[30]. H. Q. Ngo, A. Steinfeld, Joint potential-vector fields for obstacle-aware legible motion planning, in *Proceedings of the 33rd IEEE International Conference on Robot and Human Interactive Communication (ROMAN)*, 2024, pp. 1856-1863.
[31]. C. Bodden, D. Rakita, B. Mutlu, M. Gleicher, A flexible optimization-based method for synthesizing intent-expressive robot arm motion, *The International Journal of Robotics Research*, Vol. 37, Issue 11, 2018, pp. 1376-1394.
[32]. S. Nikolaidis, A. Dragan, S. Srinivasa, Viewpoint based legibility optimization, in *Proceedings of the 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 2016, pp. 271-278.
[33]. A. Dragan, Legible robot motion planning, PhD Thesis, *Carnegie Mellon University*, Pittsburgh, 2015.
[34]. A. Eirale, M. Leonetti, M. Chiaberge, Learning social cost functions for human-aware path planning, in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2024, pp. 5364-5371.
[35]. P. T. Singamaneni, P. Bachiller-Burgos, L. J. Manso, A. Garrell, et al., A survey on socially aware robot navigation: Taxonomy and future challenges, *The International Journal of Robotics Research*, Vol. 43, Issue 10, 2024, pp. 1533-1572.
[36]. A. Orthey, C. Chamzas, L. E. Kavraki, Sampling-based motion planning: A comparative review, *Annual Review of Control, Robotics, and Autonomous Systems*, Vol. 7, 2024, pp. 285-310.
[37]. A. Weiner, J. Geise, Model-based deep reinforcement learning for accelerated learning from flow simulations, *arXiv*, 2024, arXiv:2402.16543.

[38]. D. Hafner, T. Lillicrap, M. Norouzi, J. Ba, Mastering atari with discrete world models, *arXiv*, 2022, arXiv:2010.02193.
[39]. P. Wu, A. Escontrela, D. Hafner, K. Goldberg, et al., DayDreamer: World models for physical robot learning, *arXiv*, 2022, arXiv:2206.14176.
[40]. W. Huang, J. Ji, C. Xia, B. Zhang, et al., SafeDreamer: Safe reinforcement learning with world models, *arXiv*, 2024, arXiv:2307.07176.
[41]. J. Nocedal, S. J. Wright, Numerical Optimization (2nd ed.), *Springer*, New York, 2006.
[42]. J. Dai, J. Ji, L. Yang, Q. Zheng, et al., Augmented proximal policy optimization for safe reinforcement learning, in *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 37, 2023, pp. 7288-7295.
[43]. G. Luo, H. Luo, J. Zhang, Y. Luo, Real-time data predictive analytics by fast sparse Gaussian networks for optimising power efficiency, in *Proceedings of the 14th International Conference on Power Electronics Machines and Drives (PEMD)*, 2025, pp. 60-67.
[44]. M. P. Deisenroth, D. Fox, C. E. Rasmussen, Gaussian processes for data-efficient learning in robotics and control, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 37, Issue 2, 2015, pp. 408-423.
[45]. J. Lee, J. Feng, M. Humt, M. G. Müller, et al., Trust your robots! Predictive uncertainty estimation of neural networks with sparse Gaussian processes, *arXiv*, 2021, arXiv:2109.09690.
[46]. S. Curi, F. Berkenkamp, A. Krause, Efficient model-based reinforcement learning through optimistic policy search and planning, in Advances in Neural Information Processing Systems 33 (H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, et al., Eds.), *Curran Associates, Inc.*, Red Hook, 2020, pp. 14156-14170.

IFSA

Advances in Robotics and Automatic Control Volume 4

Sergey Y. Yurish, Editor

Advances in Robotics and Automatic Control, Volume 4 explores important developments in modern robotics, with particular emphasis on human–robot interaction, industrial automation, and the safe integration of autonomous systems into human environments. As robots increasingly share workplaces and everyday spaces with people, ensuring their safety, reliability, predictability, and social acceptance has become a major challenge for researchers and engineers.

This volume presents two complementary studies addressing these challenges. The first examines the transformation of industrial workplaces through automated guided vehicles (AGVs) and collaborative robots (cobots), focusing on their benefits and emerging occupational health and safety risks. The second investigates sociability modelling in robot motion, introducing approaches to predictable, understandable, and trustworthy robot behaviour through motion planning, experimental evaluation, and artificial intelligence techniques.

Written by seven authors from Canada and the United Kingdom, this book provides valuable insights for researchers, engineers, and professionals working in robotics, automatic control, industrial automation, artificial intelligence, and human–robot interaction. It will also serve as a useful reference for postgraduate students, university lecturers, occupational safety specialists, and anyone interested in the development of safer, more intelligent, and human-centred robotic systems.

ISBN 978-84-09-95166-6

9788409951666