

# **Creating a Haptic Empathetic Robot Animal That Feels Touch and Emotion**

Dissertation

der Mathematisch-Naturwissenschaftlichen Fakultät

der Eberhard Karls Universität Tübingen

zur Erlangung des Grades eines

Doktors der Naturwissenschaften

(Dr. rer. nat.)

vorgelegt von

**Rachael Christine Bevill**

aus Missouri City/Vereinigte Staaten von Amerika

Tübingen  
2023

Gedruckt mit Genehmigung der Mathematisch-Naturwissenschaftlichen Fakultät der  
Eberhard Karls Universität Tübingen.

Tag der mündlichen Qualifikation:

23.02.2024

Dekan:

Prof. Dr. Thilo Stehle

1. Berichterstatterin:

Prof. Dr. Katherine J. Kuchenbecker

2. Berichterstatter:

Prof. Dr. Martin A. Giese

# Abstract

Touch is essential to everyday interaction. Humans commonly use social touch to communicate needs, express emotions, and request attention; however, utilizing touch in these ways can be challenging for children with autism. Furthermore, while socially assistive robots are generally viewed as a promising means of helping children with autism via robot-assisted therapy, existing robots possess almost no touch-perception capabilities. We propose that socially assistive robots could better understand user intentions and needs, and provide increased support opportunities, if they could perceive and intelligently react to touch.

This thesis presents the design, creation, and testing of a touch-perceptive and emotionally responsive robot, which we refer to as the Haptic Empathetic Robot Animal, or HERA. HERA is intended to demonstrate new technical capabilities that therapists could use to teach children with autism about safe and appropriate touch. We divide HERA's development into four principal stages: establishing touch-sensing guidelines, building touch-perceiving sensors, creating a long-term emotion model that responds to touch inputs, and integrating the subsystems for real-time performance.

To start, we wanted to determine whether a touch-perceiving robot would be useful to the autism support community. We therefore begin this dissertation by establishing the first-ever touch-perception guidelines for therapy robots. We examined the existing relevant literature to create an initial set of six tactile-perception requirements, and we then evaluated these requirements through interviews with eleven experienced autism specialists from a variety of backgrounds. Thematic analysis of the specialists' feedback revealed three overarching themes: the touch-seeking and touch-avoiding behavior of autistic children, their individual differences and customization needs, and the roles that a touch-perceiving robot could play in such interactions. Using feedback collected from the specialists, we refined our initial list into finalized qualitative requirements based on seven attributes: robustness and maintainability, sensing range, feel, gesture identification, spatial, temporal, and adaptation. Finally, by utilizing current best practices in human-robot interaction, tactile sensor development, and signal processing, we transformed these qualitative requirements into quantitative specifications, which future roboticists and engineers can use in the design process of social robots.

We then proceeded to use these touch-perception guidelines to endow HERA with a sense of touch. In the second part of this dissertation, we introduce a low-cost, easy-to-build, soft tactile-perception system that we created for the NAO robot, an existing rigid-bodied robot upon which HERA was built. We also share participants' feedback on touching this system. We installed four custom fabric-and-foam-based resistive sen-

sors on the curved surfaces of a NAO's left upper limb, including its hand, lower arm, upper arm, and shoulder. We then investigated how different users perform common social touches. In our user study, fifteen adults performed five types of affective touch communication gestures (hitting, poking, squeezing, stroking, and tickling) at two force intensities (gentle and energetic) on the four sensor locations. After training on these touches in a 70%-30% train-test split, a gesture-classification algorithm based on a random forest identified the correct combined touch gesture and force intensity on windows of held-out test data with an average accuracy of 74.1%, which is more than eight times better than chance. Participants rated the sensor-equipped arm as pleasant to touch and liked the robot's presence significantly more after touch interactions. Our results showed that this type of tactile-perception system can detect necessary social-touch communication cues from users, can be tailored to a variety of robot body parts, and can provide HRI researchers with the tools needed to implement social touch in their own systems.

Next, we needed HERA to react to the social touches that it detected in an emotionally intelligent way, so that children could see the impact of their touches on the robot. However, in the field of human-robot interaction, autonomous physical robots often lack a dynamic internal emotional state, instead displaying brief, fixed emotion routines. These short-term-only responses are intended to promote specific user interactions. We hypothesized that users' perceptions of a social robotic system would improve when the robot provides emotional responses on both shorter and longer time scales (reactions and moods), based on touch inputs from the user. In the third part of the dissertation, we evaluated this proposal through an online study in which 51 diverse participants watched, rated, and commented on nine randomly ordered videos (a three-by-three full-factorial design) of HERA being touched by a human. Users provided the highest ratings in terms of agency, ambient activity, enjoyability, and touch perceptivity for scenarios in which HERA showed emotional reactions and either neutral or emotional moods in response to social touch gestures. Finally, we summarized key qualitative findings about users' preferences for reaction timing, the ability of robot mood to show persisting memory, and perception of neutral behaviors as a curious or self-aware robot.

In the final portion of this dissertation, we utilize our findings to improve and integrate HERA's subsystems. We introduced a mathematical emotion model that can easily be implemented in a social robot to enable it to react intelligently to external stimuli. The robot's affective state is modeled as a second-order dynamic system analogous to a mass connected to ground by a parallel spring and damper. By adjusting the parameters of this emotion model, the three main aspects of the robot's personality can be modified. We termed these parameters as disposition, stoicism, and calmness. We also introduce an improved version of our tactile-perception system, with sixteen sensors covering HERA's full body, and upgrades to the sensor design, microcontroller, and gesture classification approach. We explain the connections between the various subsystems and demonstrate their ability to create a robotic animal that feels both touch and emotion.

The primary contribution of this thesis is to present practical methods for enabling robots to perceive and respond to dynamic social touch. Additionally, we present touch-

perception guidelines that translate the needs of autism therapists into specifications for future roboticists and engineers, easy-to-follow DIY instructions on how to build our fabric-based tactile sensors, and a mathematical representation of our emotion response algorithm. With these systems, we can create more socially intelligent robots that can use their sense of touch to better interact with people in care settings and while completing daily tasks.



# Zusammenfassung

Berührung ist ein wesentlicher Bestandteil der täglichen Interaktion. Menschen nutzen Berührungen, um Bedürfnisse mitzuteilen, Emotionen auszudrücken und um Aufmerksamkeit zu erlangen; für Kinder mit Autismus kann es jedoch schwierig sein, Berührungen auf diese Weise zu nutzen. Während sozial unterstützende Roboter allgemein als vielversprechendes Mittel angesehen werden, um Kindern mit Autismus durch roboter-gestützte Therapie zu helfen, besitzen bestehende Roboter fast keine Fähigkeiten zur Wahrnehmung von Berührungen. Wir schlagen vor, dass sozial unterstützende Roboter die Absichten und Bedürfnisse der Benutzer besser verstehen und mehr Möglichkeiten zur Unterstützung bieten könnten, wenn sie Berührungen wahrnehmen und intelligent darauf reagieren könnten.

Diese Arbeit stellt das Design, die Entwicklung und die Verifizierung eines emotional reagierenden Roboters mit Berührungswahrnehmung vor, den wir als Haptic Empathetic Robot Animal (HERA) bezeichnen. HERA soll neue technische Fähigkeiten demonstrieren, die Therapeuten nutzen könnten, um Kindern mit Autismus sichere und angemessene Berührungen beizubringen. Wir unterteilen die Entwicklung von HERA in vier Hauptphasen: Festlegung von Richtlinien für die Berührungserkennung, Bau von Sensoren, die Berührungen wahrnehmen, Erstellung eines langfristigen Emotionsmodells, das auf Berührungseingaben reagiert, und Integration der Teilsysteme für die Anwendung in Echtzeit.

Zu Beginn wollten wir herausfinden, ob ein Roboter, der Berührungen wahrnimmt, für die Gemeinschaft zur Unterstützung von Autisten nützlich wäre. Daher beginnen wir diese Dissertation mit der Erstellung der allerersten Richtlinien zur Berührungswahrnehmung für Therapieroboter. Wir untersuchten die vorhandene einschlägige Literatur, um einen ersten Satz von sechs Anforderungen an die taktile Wahrnehmung zu erstellen, und bewerteten diese Anforderungen dann durch Interviews mit elf erfahrenen Autismus-Spezialisten aus unterschiedlichen Spezialgebieten. Die thematische Analyse des Feedbacks der Spezialisten ergab drei übergreifende Themen: das berührungssuchende und berührungsvermeidende Verhalten autistischer Kinder, ihre individuellen Unterschiede und Anpassungsbedürfnisse sowie die Rolle, die ein Roboter mit Berührungswahrnehmung in solchen Interaktionen spielen könnte. Anhand des von den Spezialisten gesammelten Feedbacks verfeinerten wir unsere ursprüngliche Liste zu endgültigen qualitativen Anforderungen, die auf sieben Attributen basieren: Robustheit und Wartungsfreundlichkeit, Erfassungsbereich, Empfindung, Gestenerkennung, räumliche Auflösung, zeitliche Auflösung, sowie Anpassbarkeit. Schließlich wandelten wir diese qualitativen Anforderungen unter Verwendung aktueller bewährter Verfahren im Bere-

ich Mensch-Roboter-Interaktion, taktilem Sensorentwicklung und Signalverarbeitung in quantitative Spezifikationen um, die künftige Robotiker und Ingenieure bei der Entwicklung von sozialen Robotern verwenden können.

Anschließend nutzten wir diese Richtlinien für die Berührungswahrnehmung, um HERA mit einem Tastsinn auszustatten. Im zweiten Teil dieser Dissertation stellen wir ein kostengünstiges, einfach zu bauendes, weiches Tastwahrnehmungssystem vor, das wir für den NAO-Roboter entwickelt haben. NAO ist ein bestehender Roboter mit starrem Körper, auf dem HERA aufgebaut wurde. Wir berichten auch über die Rückmeldungen der Teilnehmer, die mit dem System interagiert haben. Wir installierten vier speziell angefertigte, auf Stoff und Schaumstoff basierende Widerstandssensoren auf den gekrümmten Oberflächen der linken oberen Extremität des NAO, einschließlich seiner Hand, seines Unterarms, seines Oberarms und seiner Schulter. Anschließend untersuchten wir, wie verschiedene Benutzer übliche Berührungen aus dem sozialen Kontext ausführen. In unserer Benutzerstudie führten fünfzehn Erwachsene fünf Arten von affektiven Berührungsgesten (Schlagen, Stoßen, Drücken, Streicheln und Kitzeln) mit zwei Kraftintensitäten (sanft und energisch) an den vier Sensorpositionen aus. Nach dem Training mit diesen Berührungen in einem 70%-30%-Train-Test-Split identifizierte ein Gesten-Klassifizierungsalgorithmus auf der Grundlage eines Random Forest die korrekte Kombination aus Berührungsgeste und Kraftintensität in Fenstern von ungesehenen Testdaten mit einer durchschnittlichen Genauigkeit von 74,1%, was mehr als achtmal besser als Zufall ist. Die Teilnehmer bewerteten den mit Sensoren ausgestatteten Arm als angenehm zu berühren und mochten die Anwesenheit des Roboters nach den Berührungsinteraktionen deutlich mehr. Unsere Ergebnisse zeigten, dass diese Art von taktilem Wahrnehmungssystem die notwendigen Kommunikationssignale durch soziale Berührungen von Nutzern erkennen kann, dass es auf eine Vielzahl von Roboter Körperteilen zugeschnitten werden kann und dass es HRI-Forschern die nötigen Werkzeuge an die Hand geben kann, um soziale Berührungen in ihren eigenen Systemen zu implementieren.

Als Nächstes musste HERA auf die erkannten sozialen Berührungen emotional intelligent reagieren, damit die Kinder die Auswirkungen ihrer Berührungen auf den Roboter erkennen konnten. Im Bereich der Mensch-Roboter-Interaktion fehlt autonomen physischen Robotern jedoch oft ein dynamischer interner emotionaler Zustand, stattdessen zeigen sie kurze, festgelegte Emotionsroutinen. Diese kurzfristigen Reaktionen sollen bestimmte Benutzerinteraktionen fördern. Wir stellten die Hypothese auf, dass sich die Wahrnehmung eines sozialen Robotersystems durch den Benutzer verbessern würde, wenn der Roboter auf der Grundlage von Berührungseingaben des Benutzers emotionale Reaktionen auf kürzeren und längeren Zeitskalen (Reaktionen und Stimmungen) zeigt. Im dritten Teil der Dissertation haben wir diesen Vorschlag durch eine Online-Studie evaluiert, in der 51 unterschiedliche Teilnehmer neun zufällig angeordnete Videos, in denen HERA von einem Menschen berührt wird, angesehen, bewertet und kommentiert haben. Die Benutzer gaben die höchsten Bewertungen in Bezug auf Handlungsfähigkeit, Umgebungsaktivität, Vergnügen und Wahrnehmbarkeit der Berührung für Szenarien ab,

in denen HERA emotionale Reaktionen und entweder neutrale oder emotionale Stimmungen als Reaktion auf soziale Berührungsgesten zeigte. Abschließend fassten wir die wichtigsten qualitativen Erkenntnisse über die Präferenzen der Benutzer hinsichtlich des Reaktionszeitpunkts, der emotionalen Erinnerungsfähigkeit des Roboters, und der Wahrnehmung von neutralem Verhalten als neugieriger oder selbstbewusster Roboter zusammen.

Im letzten Teil dieser Dissertation nutzen wir unsere Erkenntnisse, um die Teilsysteme von HERA zu verbessern und zu integrieren. Wir führten ein mathematisches Emotionsmodell ein, das leicht in einen sozialen Roboter implementiert werden kann, damit dieser intelligent auf externe Stimuli reagieren kann. Der affektive Zustand des Roboters wird als dynamisches System zweiter Ordnung modelliert, analog zu einer Masse, die über eine parallele Feder und einen Dämpfer mit dem Boden verbunden ist. Durch die Anpassung der Parameter dieses Emotionsmodells können die drei Hauptaspekte der Roboterpersönlichkeit verändert werden. Wir bezeichneten diese Parameter als Disposition, Stoizismus und Gelassenheit. Außerdem stellen wir eine verbesserte Version unseres Systems zur taktilen Wahrnehmung vor, mit sechzehn Sensoren, die den gesamten Körper von HERA abdecken, sowie Verbesserungen des Sensordesigns, des Mikrocontrollers und des Ansatzes zur Gestenklassifizierung. Wir erklären die Verbindungen zwischen den verschiedenen Teilsystemen und demonstrieren ihre Fähigkeit, einen Roboter zu schaffen, das sowohl Berührungen als auch Emotionen wahrnimmt.

Der Hauptbeitrag dieser Arbeit besteht darin, praktische Methoden vorzustellen, die es Robotern ermöglichen, dynamische soziale Berührungen wahrzunehmen und darauf zu reagieren. Darüber hinaus stellen wir Richtlinien für die Berührungswahrnehmung vor, die die Bedürfnisse von Autismus-Therapeuten in Spezifikationen für künftige Robotiker und Ingenieure umsetzen, sowie leicht verständliche Selbermach-Anleitungen für den Bau unserer stoffbasierten Tastsensoren und eine mathematische Darstellung unseres Algorithmus für die Reaktion auf Emotionen. Mit diesen Systemen können wir sozial intelligentere Roboter entwickeln, die ihren Tastsinn nutzen können, um mit Menschen in Pflegeeinrichtungen und bei der Erledigung täglicher Aufgaben besser zu interagieren.



# Acknowledgments

It takes a village to raise a Ph.D. student. Over my six years at the Max Planck Institute for Intelligent Systems, I have been blessed to get to know so many amazing people, and I have been touched by their friendship, wisdom, and support. While I cannot name everyone here, there are those that I most certainly wish to acknowledge:

- First, to my husband, Joey Burns: Thank you for moving across the world with me and moving mountains for me. The past six years have been a crazier adventure together than I could have ever imagined. Thank you for your love, encouragement, enthusiasm, patience, and support. Thank you also for getting me out of the lab on the weekends and making sure we make the most out of life! I'm so grateful for all the things we've gotten to do and the places we gotten to explore, from the nearby Christmas markets to the faraway island of Jeju. Thank you for all your technical support, too. My research went much more smoothly because of you.
- Next, I would like to thank my advisor, Dr. Katherine J. Kuchenbecker, for taking a chance on me all those years ago. Katherine was giving a talk at my university, and I was her student tour guide for the day. After spending all day together, Katherine agreed to serve as my international host if I received the Whitaker fellowship. We got the fellowship, and the rest is history. Katherine, I am so grateful for everything you have taught me. I truly have become a better researcher because of you. Thank you for your mentorship, guidance, friendship, and for this amazing opportunity. Thank you for everything.
- I would like to thank Dr. Martin Giese for serving as my University of Tübingen advisor and on my Thesis Advisory Committee. I learned a lot at our TAC meetings, and I appreciate the time and care you put into mentoring your students.
- I also had lots of support from across the pond. Mom (Jan Bevill) and Dad (Bob Bevill), thank you for instilling a love of learning in me and for always encouraging me to do my best. Christina, Laura, Robert, and Victoria, thank you for cheering me on! I also want to thank my husband's family, who welcomed me in with open arms and endless support. I love you all.
- I would also like to acknowledge the Haptic Intelligence Department, otherwise known as the Sharks. Throughout the years, I have been proud to call them my friends. I hope you all continue to support and take care of one another, and that you continue to explore the beautiful culture and land of Germany together, too!

## *Acknowledgments*

---

- To Hasti Seifi and Hyosang Lee: Thank you for being my mentors over the years, for helping me push through the challenging early years of working on HERA, and for your insights as I start looking toward life after the Ph.D. It has been wonderful to watch you advance in your careers and your lives. I have learned so much from you.
- To the wonderful women with whom I have gotten to share an office over the years – Paola Forte, Yasemin Vardar, and Giulia Ballardini: thank you for making it such a joy to come to work in room 5N14 each day.
- To Andrew Schultz: Thank you for proofreading this thesis and being my accountability buddy as we both tackled the world of academia and job applications. More importantly, thank you for the enthusiasm and kindness you bring to our department. You are a much-appreciated postdoc!
- I would also like to thank the several interns who worked on HERA with me over the years – Sophia Hass, Ruby Rosenthal, Fayo Ojo, Keshav Garg, and Jack Klingenberg. They taught me just as much as I taught them, if not more. It was a pleasure to do science and grow with you all.
- I would also like to thank Dr. Shari Young Kuchenbecker, who helped me tailor HERA’s koala costume to be even cuter and to better fit over the sensors. It is not often one gets to learn sewing as part of their Ph.D., especially from their advisor’s mom!
- I would like to thank the Max Planck Society and the International Max Planck Research School for Intelligent Systems (IMPRS-IS) for supporting me. I would also like to acknowledge the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation), which partially funded the research discussed in Chapter 4 under Germany’s Excellence Strategy – EXC 2120/1 – 390831618.
- Finally, I would like to thank Baby Burns, who is not yet here but has been a great motivator in finishing this thesis. You are already so loved.

# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
1.1	Outline . . . . .	2
1.2	Relevant Publications and Honors . . . . .	3
<b>2</b>	<b>Background</b>	<b>7</b>
2.1	Establishing Guidelines for a Touch-Perceiving Robot for Children with Autism . . . . .	7
2.1.1	Sensory processing in children with autism . . . . .	7
2.1.2	Existing intervention methods . . . . .	8
2.1.3	Socially assistive robotics for children with autism . . . . .	9
2.2	Endowing a Robot Companion with Practical Social-Touch Perception . . . . .	10
2.2.1	Touch-perception guidelines for social robots . . . . .	10
2.2.2	Tactile sensing in human-robot interaction . . . . .	11
2.2.3	Touch perception in socially assistive robotics . . . . .	12
2.2.4	Gesture classification methods for social touch . . . . .	13
2.3	Investigating User Preferences for Robot Emotions in Response to Touch . . . . .	14
2.3.1	Emotion and social touch in social robots . . . . .	14
2.3.2	Existing computational emotion models . . . . .	14
2.4	Creating a Robot That Feels Both Touch and Emotion . . . . .	15
2.4.1	Existing computational emotion models for social robots . . . . .	15
2.4.2	Approach-avoidance theory . . . . .	16
<b>3</b>	<b>Establishing guidelines for a touch-perceiving robot for children with autism</b>	<b>17</b>
3.1	Methods . . . . .	18
3.1.1	Initial touch-sensing requirements . . . . .	18
3.1.2	Participants . . . . .	20
3.1.3	Procedure . . . . .	20
3.1.4	Qualitative analysis of interviews . . . . .	25
3.2	Results: Overarching Themes . . . . .	26
3.2.1	Touch in autistic children . . . . .	26
3.2.2	Understanding individual needs and differences . . . . .	28
3.2.3	Supporting the individual and promoting independence . . . . .	32
3.3	Results: Qualitative and Quantitative Requirements . . . . .	35
3.3.1	Qualitative tactile-perception requirements . . . . .	35
3.3.2	Quantitative specifications . . . . .	36

3.4	Discussion . . . . .	40
3.4.1	Reflecting on our methods . . . . .	40
3.4.2	Implications for future research . . . . .	42
<b>4</b>	<b>Endowing a robot companion with practical social-touch perception</b>	<b>45</b>
4.1	Tactile Sensor Design and Fabrication . . . . .	46
4.1.1	Sensor overview . . . . .	46
4.1.2	Fabrication process for curved surfaces . . . . .	48
4.1.3	Tailoring sensors across NAO’s arm . . . . .	49
4.1.4	Estimated costs and fabrication time . . . . .	49
4.2	User Study Testing . . . . .	51
4.2.1	Introducing HERA – robot form factor . . . . .	51
4.2.2	Participants . . . . .	52
4.2.3	Experimental setup . . . . .	52
4.2.4	Procedure . . . . .	53
4.2.5	Gesture classification . . . . .	55
4.3	User Study Results . . . . .	56
4.3.1	Classification results: location, force, and gesture . . . . .	56
4.3.2	Classifying force and gesture together . . . . .	57
4.3.3	Preference on the feel of the tactile sensors . . . . .	58
4.3.4	Perceptions before and after touch interaction . . . . .	59
4.3.5	Expected robot reactions . . . . .	59
4.3.6	Notable observations . . . . .	60
4.4	Physical Sensor Testing and Results . . . . .	62
4.4.1	Physical experiment procedure . . . . .	62
4.4.2	Physical experiment results . . . . .	63
4.5	Discussion . . . . .	65
4.5.1	The shoulder sensor . . . . .	65
4.5.2	Physical design implications . . . . .	66
4.5.3	Taxel size . . . . .	67
4.5.4	User study limitations . . . . .	67
4.5.5	Visual appearance of the robot . . . . .	68
4.5.6	Gesture-classification approach . . . . .	68
4.5.7	Implications for future research . . . . .	68
<b>5</b>	<b>Investigating user preferences for robot emotions in response to touch</b>	<b>71</b>
5.1	User Study . . . . .	72
5.1.1	Hypotheses . . . . .	73
5.1.2	Study design . . . . .	73
5.1.3	Participants . . . . .	75
5.1.4	Procedure . . . . .	75

5.2	Quantitative Results . . . . .	77
5.2.1	Thematically combining survey items . . . . .	77
5.2.2	Data processing . . . . .	77
5.2.3	RANOVA analysis . . . . .	78
5.2.4	Posthoc testing . . . . .	79
5.3	Qualitative Results . . . . .	79
5.3.1	Lack of reaction to touch is unsettling and unintelligent . . . . .	80
5.3.2	Delayed reaction (i.e., mood) is better than no reaction . . . . .	80
5.3.3	Neutral behaviors perceived as curiosity and awareness . . . . .	80
5.3.4	Mood amplifies reaction and shows persisting memory . . . . .	80
5.3.5	Emotional reaction sounds were polarizing . . . . .	81
5.4	Discussion . . . . .	81
5.4.1	Evaluating our hypotheses . . . . .	81
5.4.2	Limitations . . . . .	82
5.4.3	Implications for future research . . . . .	83
<b>6</b>	<b>Creating a robot that feels both touch and emotion</b>	<b>85</b>
6.1	Mathematical Emotion Model . . . . .	85
6.1.1	Appropriate model order . . . . .	85
6.1.2	Second-order dynamic system . . . . .	87
6.1.3	Customizing the robot’s personality . . . . .	88
6.2	Upgrades to the Tactile-Perception System . . . . .	91
6.2.1	Sensor improvements . . . . .	91
6.2.2	Microcontroller improvements . . . . .	93
6.2.3	Collecting a new training dataset . . . . .	93
6.3	Upgrades to the Gesture Classification Algorithm . . . . .	95
6.3.1	New gesture classifier approach . . . . .	95
6.3.2	Mechanical coupling vs. electrical crosstalk . . . . .	96
6.3.3	Creating different models . . . . .	98
6.3.4	Classification results: offline evaluation . . . . .	98
6.3.5	Classification results: real-time evaluation . . . . .	100
6.3.6	Integrating the emotion model and gesture classifier . . . . .	106
6.4	Discussion . . . . .	106
6.4.1	Limitations . . . . .	106
6.4.2	Implications for future research . . . . .	107
<b>7</b>	<b>Conclusion</b>	<b>109</b>
7.1	Summary . . . . .	109
7.1.1	Establishing guidelines for a touch-perceiving robot for children with autism . . . . .	109
7.1.2	Endowing a robot companion with practical social-touch perception . . . . .	109

*Contents*

---

7.1.3	Investigating user preferences for robot emotions in response to touch . . . . .	110
7.1.4	Creating a robot that feels both touch and emotion . . . . .	110
7.2	Future Work . . . . .	110
	<b>Symbols</b>	<b>113</b>
	<b>Abbreviations</b>	<b>115</b>
	<b>Bibliography</b>	<b>125</b>

# Chapter 1

## Introduction

Social touch is an integral aspect of our daily interactions with colleagues, friends, and family. We touch other people to gain attention, communicate needs, and build empathy and attachment [31, 131]. Touch promotes social bonding and cognitive development in children [33] and is essential for emotional well-being [60]. Affective touch, or touch with an emotional component, is used to convey our emotions to others through nonverbal communication, such as a hug [17].

Separately, the field of socially assistive robotics is rapidly growing, with robots serving as therapy assistants [55, 71], caregivers [15, 51, 82], and companions [116, 65, 114]. While present social robots may feature other advanced capabilities, they are simply not equipped with adequate touch perception. If these robots have any touch sensing, it is typically limited to detecting either binary contact or a simple force threshold, and only at key locations on the robot's body [107, 132]. There is a serious need in the field of human-robot interaction (HRI) for such robots to increase in awareness and intelligence by developing an understanding of social touch.

This thesis presents the creation and refinement of a socially assistive robot that perceives and emotionally responds to social touch, which we refer to as the Haptic Empathetic Robot Animal, or HERA. HERA's development produced three further contributions. First, we established the first set of touch-perception guidelines for therapy robots by interviewing eleven autism specialists. We assessed their feedback through thematic analysis and translated the resulting qualitative requirements into quantitative specifications, which are suitable for HRI researchers and roboticists to implement in their own work. Second, based on these touch-perception guidelines, we iteratively designed fabric-based resistive tactile sensors that are robust, easy to make, pleasant to touch, and can be fitted to existing rigid-bodied robots. Our touch-perception system can detect the type of social touch gesture, force intensity, and general body location contacted, using the time-series resistance data from the sixteen sensors that cover HERA's whole body. We also provided instructions on how to build these sensors, including a step-by-step instructional guide and video, and an open-access database of sensor templates. These open-access materials and this frugal sensor design can lower the barrier of entry for equipping robots with social touch capabilities. Third, we introduce a mathematical emotion model that HERA uses to respond to the social touches it detects. By adjust-

ing the model's mathematical parameters, one can shape different aspects of HERA's personality, as well as its emotional reactions and mood levels. This approach can enable researchers to endow their robots with a dynamic personality without the need for computationally complex emotion simulations.

## 1.1 Outline

This dissertation contains four research topics. First, we established touch-perception guidelines for a therapy robot for children with autism. Next, we created a tactile-perception system that can be externally fitted to existing rigid-bodied robots in order to detect various social touches. Third, we investigated how users preferred a robot to emotionally react in response to social touch stimuli. Finally, we combined our findings and technology to create a Haptic Empathetic Robot Animal that can perceive and emotionally react to touch.

The first two chapters of this dissertation are this introduction and the background, and the final chapter is the conclusion. The detailed outline of the remaining chapters is as follows:

- **Chapter 3: Establishing guidelines for a touch-perceiving robot for children with autism.** This chapter establishes touch-perception requirements that a social robot should meet in order to be a useful therapy tool for children with autism. It begins by describing the preliminary touch-sensing requirements we derived from the literature, followed by our procedure for interviewing eleven autism specialists and the process for thematic analysis of the resulting interviews. Next, we present the three overarching themes that emerged from our interview analysis. From these themes, we derive key qualitative requirements for a touch-perceiving robot and translate them into quantitative specifications. In the discussion section, we reflect on our methods and discuss the implications of our guidelines on future socially assistive robotics research.
- **Chapter 4: Endowing a robot companion with practical social-touch perception.** This chapter describes the fabrication, user study, and physical testing of our custom fabric-based tactile sensors affixed to the arm of our social robot prototype, HERA. It begins with an overview of the sensor's working principles and do-it-yourself (DIY) instructions on how the sensors can be recreated for future HRI studies. The section on User Study Testing explains our study procedure and our gesture classification approach. The User Study Results section presents the accuracy of our gesture classification algorithm across sensor location, level of force intensity, type of gesture, and force and gesture combined, along with feedback from participants. In the section on Physical Sensor Testing and Results, we report additional details on the sensor's performance. Finally, the Discussion sec-

tion explores the effects of various design choices on our sensors' performance and offers opportunities for improvement and exploration in future research.

- **Chapter 5: Investigating user preferences for robot emotions in response to touch.** This chapter investigates viewers' preferences on HERA's short-term and long-term emotional responses to positive and negative social touch. First, we state our hypotheses, explain the study procedure, and introduce the nine video conditions that participants observed. Next, we analyze participants' quantitative ratings of the videos, followed by a qualitative analysis of their written feedback. Finally, in the Discussion section, we evaluate our hypotheses compared to the data and propose further steps with which emotional responses can be implemented in social robots.
- **Chapter 6: Creating a robot that feels both touch and emotion.** This chapter describes the steps taken to combine all the previous findings and subsystems developed while creating HERA. First, we introduce a mathematical emotion model which can be used to generate short-term and long-term emotional responses. Next, we describe improvements made to both hardware and software aspects of HERA's tactile-perception system, including upgraded tactile sensors, a new touch dataset, and an updated real-time classification algorithm. We evaluate the real-time performance of two classifier models and explain how the emotion model and classifier subsystems are connected. Finally, the Discussion section describes potential next steps, including the outline of a user study that could be conducted with HERA.

## 1.2 Relevant Publications and Honors

This dissertation contains research based on the following publications:

1. **Rachael Burns** and Katherine J. Kuchenbecker. Designing a haptic empathetic robot animal for children with autism. Workshop paper (4 pages) presented at the workshop, "Robot-Mediated Autism Intervention: Hardware, Software and Curriculum", held at *Robotics: Science and Systems (RSS)*, Pittsburgh, USA, June 2018.
2. **Rachael Bevill Burns**, Hyosang Lee, Hasti Seifi, and Katherine J. Kuchenbecker. A fabric-based sensing system for recognizing social touch. Work-in-progress paper (3 pages) presented at *the IEEE Haptics Symposium*, Washington, D.C., USA (virtual), March 2020.
3. **Rachael Bevill Burns\***, Neha Thomas\*, Hyosang Lee\*, Robert Faulkner, and Katherine J. Kuchenbecker. Tactile textiles: An assortment of fabric-based tactile sensors for contact force and contact location. Hands-on demonstration presented

at *EuroHaptics*, Leiden, Netherlands, September 2020. \*These authors contributed equally to this publication.

4. **Rachael Bevill Burns**, Hasti Seifi, Hyosang Lee, and Katherine J. Kuchenbecker. Utilizing interviews and thematic analysis to uncover specifications for a companion robot. Workshop paper (2 pages) presented at the workshop, “Enriching HRI Research with Qualitative Methods”, held at *the International Conference on Social Robotics (ICSR)*, Golden, USA (virtual), November 2020.
5. **Rachael Bevill Burns**, Hasti Seifi, Hyosang Lee, and Katherine J. Kuchenbecker. Getting in touch with children with autism: Specialist guidelines for a touch-perceiving robot. *Paladyn, Journal of Behavioral Robotics*, 12(1):115-135, 2021.
6. **Rachael Bevill Burns**, Hasti Seifi, Hyosang Lee, and Katherine J. Kuchenbecker. A haptic empathetic robot animal for children with autism. *Companion of the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 583–585, Boulder, USA, March 2021.
7. **Rachael Bevill Burns**, Hasti Seifi, and Katherine J. Kuchenbecker. Evaluation of a touch-perceiving, responsive robot koala for children with autism. Workshop paper (4 pages) presented at the workshop, “Workshop YOUR study design! Participatory Critique and Refinement of Participants’ Studies” held at *the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, Boulder, USA (virtual), March 2021.
8. **Rachael Bevill Burns**. Teaching safe social touch interactions using a robot koala. Workshop paper (1 page) presented at the workshop, “Proximity Perception in Robotics: Increasing Safety for Human-Robot Interaction Using Tactile and Proximity Perception” held at *the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. Prague, Czech Republic (virtual), September 2021.
9. **Rachael Bevill Burns**, Hyosang Lee, Hasti Seifi, Robert Faulkner, and Katherine J. Kuchenbecker. Endowing a NAO robot with practical social-touch perception. *Frontiers in Robotics and AI*, 9(840335):1-17, 2022.
10. **Rachael Bevill Burns**, Hyosang Lee, Hasti Seifi, Robert Faulkner, and Katherine J. Kuchenbecker. User study dataset for endowing a NAO robot with practical social-touch perception. Edmond, Dataset, V1. 2022.
11. **Rachael Bevill Burns**, Hyosang Lee, Hasti Seifi, Robert Faulkner, and Katherine J. Kuchenbecker. Sensor patterns dataset for endowing a NAO robot with practical social-touch perception. Edmond, Dataset, V2. 2022.

12. **Rachael Bevill Burns**, Ruby Rosenthal, Keshav Garg, and Katherine J. Kuchenbecker. Do-it-yourself whole-body social-touch perception for a NAO robot. Workshop paper (1 page) presented at the workshop, “Large-scale Robotic Skin: Perception, Interaction, and Control” held at *the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. Kyoto, Japan, October 2022.
13. **Rachael Bevill Burns** and Katherine J. Kuchenbecker. A lasting impact: Using second-order dynamics to customize the continuous emotional expression of a social robot. Workshop paper (5 pages) presented at the workshop, “Lifelong Learning and Personalization in Long-Term Human-Robot Interaction (LEAP-HRI)” held at *the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, Stockholm, Sweden, March 2023.
14. **Rachael Bevill Burns**. Creating a haptic empathetic robot animal for children with autism. Workshop paper (4 pages) presented at the workshop, “RSS Pioneers”, held at *Robotics: Science and Systems (RSS)*, Daegu, Republic of Korea, July 2023.
15. **Rachael Bevill Burns**, Fayo Ojo, and Katherine J. Kuchenbecker. Wear your heart on your sleeve: Users prefer robots with emotional reactions to touch and ambient moods. *Proceedings of the IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*, pages 1914–1921, Busan, Republic of Korea, August 2023.

The research presented in some of these publications is associated with the following honors:

1. 2019 International Max Planck Research School for Intelligent Systems (IMPRS-IS) Boot Camp Award – 2nd place for Best Research Lightning Talk
2. Selected for the 2021 cohort of HRI Pioneers at *the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*.
3. 2021 IMPRS-IS Boot Camp Award – 2nd place for Best Research Lightning Talk
4. Received first place at the Student Elevator Pitch Competition at *the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*.
5. Selected for the 2023 cohort of RSS Pioneers at *Robotics: Science and Systems (RSS)*.
6. 2023 IMPRS-IS Boot Camp Award – Honorable Mention for Best Open House Demonstration in the Stuttgart site of MPI-IS



# Chapter 2

## Background

### 2.1 Establishing Guidelines for a Touch-Perceiving Robot for Children with Autism

#### 2.1.1 Sensory processing in children with autism

Autism spectrum disorder (ASD) is a complex condition that affects many systems in the human body, from neurological aspects to physical comorbidities. Children with autism often endure sensory overload from everyday stimuli [7, 31]. They may be nonverbal and may have difficulty understanding and relaying emotions. These combined experiences can cause the child to engage in repetitive or self-injuring behavior, as well as self-isolation and heightened stress during social interactions [86].

Over 96% of children with autism have disordered sensory processing [86]. They may experience sensory overload, where one or more senses overreact to a stimulus. Likewise, they could also experience an undersaturated response to stimuli. The senses affected, as well as the over- or under-responsiveness of each sensory system, will depend on the individual child. These sensory imbalances can be distracting, frustrating, and even painful. Nonverbal children with autism may be particularly affected; unable to communicate their needs or sensory pains, they may inflict self-injury or act aggressively. Traditional techniques to help these children cope with stress and overstimulation must be taught and administered by a trained adult [13]. As the rate of autism diagnosis continues to rise [35], and as the relative supply of caregivers, therapists and paraeducators dwindles, there is an urgent need for new mechanisms to help children with autism learn to cope with stressful or unfamiliar situations.

Touch is an essential component of early childhood development [33]. Affective touch, or touch with an emotional component, promotes social bonding, secure attachment, and social communication skills. Touch is also key for environment exploration [31]. While touch is one of the most commonly affected sensory systems in children with autism, tactile processing issues in autistic children are less studied than issues with visual or auditory processing [31, 86]; this asymmetric focus within autism research mirrors the broader trend in our current scientific understanding of different senses [76].

Cascio et al. observe that individuals with autism may generally be under-responsive

to pleasant tactile sensations and over-responsive to unpleasant sensations [31]. This combination leads to an overall tactile defensiveness, causing children with autism to gravitate toward controlled, predictable, and repetitive situations and stimuli. Therefore, children with autism may show a blend of input-seeking and input-avoiding qualities. For example, a child may regularly chew on the sleeve of their sweater but avoid unfamiliar foods with unknown textures [7].

### 2.1.2 Existing intervention methods

A diverse set of autism specialists focus on different aspects of the development of children with ASD, as well as their care. Some of the most common services children with autism receive include **occupational therapy**, which focuses on cultivating sensorimotor skills and skills needed for daily life; **physical therapy**, which develops mobility and the muscular system and improves gross motor skills; and **speech therapy**, which aims to remedy language and communication impairments [95, 102]. A child with autism may receive other services besides these, such as behavioral services, orientation and mobility services, and medical services, along with any other related services as deemed necessary by their individualized education program (IEP) [78]. One important issue all therapists must face is calming down a child who becomes stressed while in their care.

Deep-touch pressure (DTP) therapy is a tactile-based intervention method for helping children with autism reduce their stress. It is the current gold standard of treatment and utilizes a wide range of therapy tools, such as weighted blankets, weighted vests, and Wilbarger brushes [83]. However, the large range of therapeutic tools available within the classification of DTP means that customization to find the method most beneficial for an individual child may take many attempts [13]. Furthermore, DTP must be administered and monitored by a trained adult.

Animal-assisted intervention is a promising alternative to DTP that encourages independent social and physical interaction between autistic children and an animal companion [99]. Guinea pigs [97], dogs [89], and horses [9] have been found to improve the stress levels, mood, and behavior of children with ASD. Additional benefits included increased independent sensory seeking and social contact with peers. However, a live animal companion is not always feasible for many reasons, including cost, allergen or hygiene concerns, and availability. Additionally, even with robust training, an animal's behavior cannot be predicted entirely. Furthermore, as young children with autism may not fully understand the effect of their touch on others, they may accidentally touch the animal in an unsafe way (e.g., touching the eyes, nose, or teeth of a dog). Hurting or startling even a well-trained animal could cause it to react badly to a child.

To further support families, there is growing interest in incorporating socially assistive robots into the therapy, routines, and care of children with autism. Robots are ideal assistants in autism intervention because they have simplified features, can utilize a variety of sensory outputs to reinforce communication (such as colored lights and sound effects to signify different emotions), and are less intimidating to the children than a human

interaction partner [29].

### **2.1.3 Socially assistive robotics for children with autism**

The field of socially assistive robotics (SAR) aims to use interaction with robots to improve daily life, often for those with impairments [49]. For example, PARO, a robot baby seal, reduces stress and improves various symptoms in elderly patients with dementia [116]. Huggable, a blue teddy-bear-inspired robot, was designed to comfort and entertain children during long-term hospitalization [65]. The Haptic Creature, a furry, lap-sized robot animal, uses calm physical breathing patterns to improve mood and decrease stress; thus far, its benefits have been validated in neurotypical participants [114]. Therabot, a robot dog inspired by animal-assisted therapy, is designed to help survivors of trauma reduce the stress they experience [14].

The use of SAR is especially being investigated for children with autism. While preferences may vary greatly between individuals, children with autism generally seem to prefer child-sized robots [29] with simple, often exaggerated, features [101, 106]. To reward further interactions, the robot should produce behavior that is somewhat predictable and that can be controlled by the child [101, 29, 106]. Robot appearances and/or behaviors are often customized to align with these findings. For example, the IROMEC robotic toy encourages children with autism to engage in robot-assisted play and promotes growth in several developmental categories, including cognitive development, through cause-and-effect activities [50]. In another research effort, five children with autism interacted with the humanoid robot NAO through two different play modes – dancing and an interactive touching game [124]. The play modes encourage prosocial behaviors and promote a sense of bodily awareness in the child.

Socially assistive robots can positively impact children with autism through interaction in many roles: as social mediators to encourage communication [71], as therapy assistants to elicit positive behaviors such as joint attention [45] and emotion recognition and expression [63], and as playmates [29, 120, 50]. In these roles, timely responses from the robot can promote the child’s continued learning and interest [71]. It is also important to assess how children interact with a robot in a long-term setting. Three different longitudinal studies were recently conducted with different robot form factors and children between ages 3.5 to 12 years old; two of these studies focused on autistic children, and the third did not report the neurodevelopmental status of the participants. Beneficially, they found improved social skills [111], increased communication attempts [100], and a maintained interest in the robot at the end of the two-month-long interaction period [134].

As physical contact and tactile exploration are key tools for child development [31], endowing a robot with robust touch-sensing capabilities could greatly enhance its interaction possibilities, and by extension, its potential benefit for an autistic child. Tactile sensing varies greatly across SAR technology, ranging from no sensing to sensing at a large number of discrete points on the robot’s body. Even when the robot is equipped with some tactile sensors, the majority of existing robot behavior modes rely almost

exclusively on voice recognition, visual cues, and/or remote control of the robot by a trained operator [65, 71, 45, 64, 115]. These approaches either miss important touch input or include it at the cost of a human operator. Other robots, such as PARO, the robot dog AIBO, and the humanoids NAO and Pepper, limit physical tactile perception to a few key areas, such as the head and limb extremities, often with only binary output at each location. The childlike robot KASPAR utilizes force-sensing resistors (FSRs) at several discrete points on its body [105]. Finally, some research prototypes in SAR include a large number of contact-sensing points. An initial model of the Haptic Creature detected touch using an internal accelerometer and 56 FSRs distributed across its body [36]. A later model utilized pressure-sensing piezoresistive fabric [30]. The mobile ball robot CARBO, intended to provide an automated therapy experience similar to DTP, collects tactile input from 67 miniature trackballs spread across its shell [75]. While these studies show growing interest in touch sensing for SAR, the variety of designs suggests a lack of design guidelines in this area.

Existing observation studies with children with autism provide preliminary guidelines for touch-sensing in autism therapy. Notably, in a study by Robins et al., three child participants with autism were given unconstrained interaction time with the KASPAR robot for a preliminary case study [107]. Instances of physical contact ranged from less than one second to more than two seconds. Intensities of forces applied ranged from a very gentle touch to forceful squeezing. Furthermore, several instances of affective touch were observed, such as kisses, gentle hand-holding, and stroking the robot's cheek. In a three-year-long longitudinal study of over 100 sessions, a small snowman-shaped robot called Keepon was placed in a daycare center and teleoperated by a remote person. More than 30 preschool-aged children with autism ultimately interacted with the robot. Their results highlight a variety of touch interactions among a select set of participants over time. One participant initially guided her therapist's hand to put a hat on Keepon's head (her fifth session with the robot – S5), but she became bolder in her own physical interaction with Keepon, first by poking it with a xylophone stick, later touching its belly, and finally giving Keepon a kiss (S14). Another participant showed initial disinterest in Keepon but slowly gained willingness to physically interact, touching it for the first time during S10 and poking the robot frequently to prompt responses by S17. A third participant initially kicked Keepon over, but he became much gentler and friendlier toward the robot in later sessions [74].

## **2.2 Endowing a Robot Companion with Practical Social-Touch Perception**

### **2.2.1 Touch-perception guidelines for social robots**

While current social and socially assistive robots are generally limited in their touch-sensing capabilities, several teams have investigated whether touch perception is a worth-

while endeavour to pursue. They have also established which aspects of touch are most important for a robot to understand social touch communication from a human user and which features are less needed.

Tanaka *et al.* [126] found that robots would benefit from having a system that can detect touch contacts across their whole body, and that contact locations can be grouped into general anatomical regions without fine spatial sensing resolution within each region. Experimenters coded videos from 45 field sessions of children interacting with the small humanoid robot QRIO in a classroom. The coders marked binary inputs indicating whether touch occurred across any of eight general body regions of the robot. By analyzing these data using a regression model, the researchers found that touch anywhere on the robot's body was a strong predictor of the quality of social interaction, both between the human and the robot and between the human and their peers.

Andreasson *et al.* [3] used video annotation to observe 32 adult participants performing 23 possible social-touch gestures on a NAO robot to convey nine different emotions. In order to replicate a human-human interaction study by Hertenstein *et al.* [61], they visually divided the NAO robot's body into 16 general body regions, and they hand-coded the gesture intensity, duration, general location, and touch type. The highest percentage of touches occurred on the robot's arms and hands. In another part of the same user study, they covered the NAO in removable fabric garments intended to represent and eventually contain tactile sensors [84]. Their results showed that participants used stronger intensities, interacted for a longer duration, and touched more locations on the robot when it wore the garments compared to the uncovered NAO robot.

### 2.2.2 Tactile sensing in human-robot interaction

While tactile sensors have been studied for several decades [42], most development has been focused on robot fingertip sensors, which support object manipulation by providing high spatial resolution and accuracy across a small surface area [5]. On the other hand, large-scale tactile sensors, which can provide sensing across the broad surfaces of a robot's whole body, have received far less attention due to their high cost and high system complexity. For example, one existing approach is to utilize small discrete sensors that are spaced intermittently across the desired sensing area, such as the FSRs on the original Haptic Creature [36] or miniature capacitive force sensors like the SingleTact [103]. However, due to their small size (usually between 8 mm and 15 mm diameter), one must apply a great number of these sensors to cover the surface area of a robot. Such sensors are thus typically more costly per unit area than other methods, such as fabric-based tactile sensing, and they cannot sense contacts that occur in the regions between the sensors. Furthermore, this type of thin laminated sensor often experiences delamination, which harms the sensor's performance [125]; laminated flexible sensors can also be difficult to apply on a robot's curved surfaces as they require a flat surface for good adhesion. However, recent advances in sensing materials and computation technologies have enabled the creation of some simple and affordable alternative large-scale

sensing solutions [85, 129].

Fabric-based tactile sensors are large-scale tactile sensors that are typically low cost and easy to manufacture. They are also flexible [43], which makes them a better fit than rigid tactile sensors for curved surfaces like a robot's arm. The simplest fabric-based design uses one layer of low-conductivity fabric sandwiched between two layers of high-conductivity fabric. When a user compresses the sensor, the two highly conductive layers come closer together, causing the sensor's electrical resistance to decrease. However, due to the small distance (only one piece of fabric) between the two conductive layers, this design is not very soft and exhibits only a low range of resistances in response to different force inputs. In the design of Day *et al.* [43], an individual taxel demonstrates a dynamic force-sensing range from roughly 0.5 N to 5.5 N. One way to increase the sensing range is to add a middle layer of plastic mesh, which serves as a spacer between the two conductive layers. However, an initial amount of force is required before contact can be detected through the plastic mesh. When indented with a 5 cm<sup>2</sup> probe tip, the fabric sensor of Büscher *et al.* [28] does not begin changing resistance until a threshold of roughly 2.5 kPa, or 1.25 N; gentle forces below this level are critical for perceiving affective touch [127]. While fabric-based sensors are cost effective and easy to make, this field lacks a sensor design that has both a *wide sensing range* and *good low-force touch detection*.

### 2.2.3 Touch perception in socially assistive robotics

Most robots in SAR are very limited in terms of touch perception. The humanoid robots NAO and Pepper detect binary touch inputs – pressed or not pressed – through capacitive sensors at a small number of discrete key body locations. The baby seal robot PARO has touch sensors in its head and flippers, and it reacts positively to gentle touches and negatively to forceful touches [132]. While most of these systems are fairly simple, two examples of the state of the art for SAR touch perception are Kaspar [135] and the Haptic Creature [114].

Kaspar is a child-like humanoid robot designed as an assistive tool for children with autism [135]. The current iteration of Kaspar, K5, features fifteen FSRs positioned at discrete points across the hands, arms, feet, legs, chest and face of the robot. These FSRs are used to identify whether firm or gentle force is used on the robot (separating the levels at 0.6 N), but not what kind of gesture has been performed [107]. In contrast to our goals, Kaspar's FSR sensing areas are small, and its surfaces are rigid.

Envisioned as a therapeutic companion, the Haptic Creature is a furry lap-sized robot animal that calms participants by imitating slow breathing patterns [36, 114]. The initial model of the Haptic Creature detected touch through 56 FSRs distributed across its body [36]. In Sefidgar *et al.* [114]'s later version, the FSRs were replaced with a large-scale fabric-based piezoresistive sensor array. The robot's tactile sensing system was used to identify what gesture had been performed, but not to try to deduce the gesture's force intensity. A random forest classifier with 20-fold cross validation was used to iden-

tify six gestures performed on the robot by ten participants with 88.6% accuracy [30]; this performance is about five times higher than random chance. In contrast to our hardware goals, the Haptic Creature is a fully custom robot with a simple body and few degrees of freedom.

### 2.2.4 Gesture classification methods for social touch

To perceive the information gathered by tactile sensors as specific touch gestures, a robot needs a mental framework for understanding what the possible gestures are. This skill of tactile gesture recognition is often implemented by processing acquired tactile data through a machine-learning classifier that has been trained on labeled examples. These classification algorithms typically analyze the time-series touch data using data sampling windows with a fixed time duration [30].

To learn more about various classifier techniques, we refer the reader to the following papers, which detect social touch gestures using a variety of classification models: Bayesian classifiers [67], decision trees [72], LogitBoost [118], neural networks [121], random forests [2], and support vector machine (SVM) variations [68, 40]. Additionally, Jung *et al.* [68] provide a comprehensive list of research on the use of tactile sensing to detect and identify social touch, including a table identifying the surface touched (e.g., a mannequin arm, a stationary robot), the gestures performed, the classifier method, and the system's classification accuracy.

The random forest classifier approach has been successfully utilized in several works studying touch gesture recognition [70, 2], including the second version of the Haptic Creature [30]. In a comparative study by Keshmiri *et al.* [70], participants conducted gentle and strong versions of three touches on a mannequin wearing a vest with a  $32 \times 32$  grid of pressure sensors: hitting on the chest, hitting on the shoulder, and hugging. The random forest classifier performed significantly better at identifying the different touches than five other classification approaches, achieving an average accuracy of 85% (about five times higher than the random chance level of 16.67%).

While many systems have classified social touch gestures, very few have sought to identify gesture and force intensity together. *Gesture type* and *force intensity* convey two distinct and important features of social touch [22]. By detecting both, a robot could potentially obtain a better understanding of the user's intent and how to respond. The gesture-recognition algorithm we use to evaluate our tactile sensors is based on a random forest and classifies both gesture and force intensity.

## 2.3 Investigating User Preferences for Robot Emotions in Response to Touch

### 2.3.1 Emotion and social touch in social robots

A robot can display emotions to make the user aware of its internal state, to reinforce a user's action, or even to guide the user to a goal behavior [20]. Emotional display directly improves user interaction: children reacted more expressively and positively toward a NAO with a teleoperated emotion model than toward its non-emotive counterpart [130]. The autonomous Roboceptionist by Kirby et al. used emotions, moods, and long-term attitudes toward repeat visitors (identified by a swipe of their university ID cards) during chat-based interactions [73]. Users significantly changed their interactions with the robot based on its mood, hinting at the power of this capability for social robots.

Additionally, while existing commercially available robots have been limited to no touch perception, researchers are investigating how upcoming robots should react to social touch. Fitter and Kuchenbecker found that users' perceptions of a hand-clapping robot's pleasantness and energeticness significantly increased when the robot used facial reactivity to acknowledge users' touch contacts [54]. Lehmann et al. investigated what type of directional movement was considered an appropriate response to touch contact on the hand for Pepper and NAO [79]. Participants perceived the android robot ERICA to be more human-like when it gave an immediate subtle reaction to touch compared to responding after a two-second delay [117]. HuggieBot 2.0 haptically detects when a user initiates and ends a hug [17]. HuggieBot 3.0 also identifies the type of intra-hug gesture the user performs (e.g., rubbing or patting the robot's back) and reciprocates with an intra-hug gesture of its own [18]. The robot seal PARO coos pleasantly in response to gentle touches and cries if it is handled with a high level of force [5]. While there have been several instances of robots providing an immediate acknowledgement of social touch contacts, we want to investigate how both shorter- and longer-term emotional responses that are customized to the user's touch input affect their perceptions of the system.

### 2.3.2 Existing computational emotion models

While complex emotion models exist for chatbots and virtual agents, within the study of emotion usage for autonomous social robots, there has been a strong focus on human recognition of and responses to static robot emotions, rather than building dynamic emotion systems that change over time [122]. However, some models exist that track the robot's internal state, such as the TAME framework by Moshkina et al. [94]. This framework presents an emotion model composed of traits (T), attitudes (A), moods (M), and emotions (E). Each of these four categories is affected by both internal and external factors over varying time scales, with fundamental traits remaining constant over time,

attitudes changing very slowly, moods changing over the course of a day, and emotions changing quickly in reaction to immediate stimuli. Two proof-of-concept user studies with partial implementations of the framework showed preliminary success, with the traits and emotion components enabled in the dog-like robot AIBO [93] and a simple demonstration of mood and emotion enabled in the humanoid robot NAO [92]. This second example involved a mock-search-and-rescue scenario: after the overhead lights dimmed, the NAO used different approaches to tell the participant that the area was unsafe and needed to be evacuated. Participants exited the study area further and faster in the conditions that used a negative mood through voice and actions compared to the neutral control condition [92], indicating the power of affective communication in HRI.

Notably, nearly all robot systems that emote based on user input focus on vision- or audio-based sensing modalities [34]. We did not find any models that suggested tactile input, such as social touch from a user, as suitable stimuli. Therefore, there is a need to systematically investigate the relationship between social touch as an input and what users perceive to be appropriate affective robot responses as an output.

## 2.4 Creating a Robot That Feels Both Touch and Emotion

### 2.4.1 Existing computational emotion models for social robots

Ojha et al. provide a thorough review of existing computational models of emotion from the last two decades [98]. However, several of these models have been implemented with virtual agents [113] rather than on physical robots, and they typically focus on providing lifelike visual responses to text- or chat-based input [8, 46], which are complex and highly cognitive. Some models utilize machine-learning approaches, such as a Gaussian mixture model [81] or a series of neural networks [38], to recognize and reproduce the user's emotional cues or otherwise adjust the robot's affective behavior.

Within the study of social robots displaying emotions, there has been a strong focus on human recognition of and responses to static robot emotions, rather than methods for building dynamic emotion systems that change over time [123, 21]. This focus on static emotions may stem from the fact that existing computational emotion models are multi-faceted and difficult to translate to a different robotic platform. The TAME Framework by Moshkina et al. [94] showed promising initial results at creating an affect model for physical robots; however, the existing equations for the individual affect categories are complex, and no code is provided for replication or adaptation by other researchers. There is a need to develop emotion systems that allow for the easy adjustment of personality parameters, demonstration of the internal affective state over increasingly long time scales, and usability across a variety of robotic systems.

## 2.4.2 Approach-avoidance theory

To facilitate high-quality HRI, we aimed to develop an emotion model that is rooted in behaviors that can be observed in nature. Approach-avoidance theory is a well-validated theory in psychology which states that when an organism receives a positive stimulus (i.e., something which supports survival), it feels a positive emotion, which motivates approach behaviors, such as lifting arms and cooing for a baby. A negative stimulus (i.e., something which hinders survival) elicits a negative emotion and motivates the organism to avoid and withdraw [47].

Furthermore, approach-avoidance theory explains that intense emotional responses tend to be bi-phasic: if a very negative stimulus is presented and then removed, the emotional state not only returns to neutral, but also temporarily springs up to a positive state [44]. Conversely, the removal of a very positive stimulus can result in a negative mood. For example, a person might feel negative emotions while they are sick or injured. When they are healthy again, they not only no longer feel negative, but they also feel even more positive and more appreciative of their good health than they did before the illness began. This pattern suggests that a robot's emotional state should have this property of bouncing back beyond neutral after spending a longer duration of time either positive or negative.

By using an emotion model consistent with approach-avoidance theory, robots can produce naturalistic responses during social interaction. This behavior is especially important for robots that primarily function in social settings, such as a peer-like tutor, an empathetic caretaker, or an attentive animal companion.

## Chapter 3

# Establishing guidelines for a touch-perceiving robot for children with autism

**Note:** This chapter is based on Burns, Seifi, Lee, and Kuchenbecker’s article “Getting in Touch with Children with Autism: Specialist Guidelines for a Touch-Perceiving Robot”, which was published in *Paladyn, Journal of Behavioral Robotics* [22] and is licensed under a Creative Commons license (CC BY). Some paragraphs in this dissertation’s Abstract, Introduction, Background, and Conclusion are also adapted from this publication.

While much work is already being done to bring robots into autism therapy and care, the use of tactile sensing is severely lacking in comparison to other sensing modalities. Children with ASD often struggle with speech and visual emotion cues, so touch is too crucial a communication channel to ignore [31].

Taking inspiration from existing methods such as deep-touch pressure therapy [83] and animal-assisted intervention [99], we aim to broaden the tactile interaction capabilities of socially assistive robots in autism intervention. We thus investigated guidelines for the touch-sensing capabilities of a robot companion for autistic children. Our approach combines an initial literature review, an in-depth interview study, and current best practices in tactile sensor development and signal processing. We translated the literature into a set of initial requirements for touch sensing, and we then explored these requirements through hour-long interviews with eleven autism specialists from a variety of backgrounds. We systematically examined these interviews using the method of thematic analysis [19].

Our results highlight three overarching themes: the touch tendencies of children with autism, the importance of their individual differences, and the role of therapists in each child’s development. Participants recommended inclusive design strategies that would enable the companion robot and its tactile-sensing system to work for both touch-seeking and touch-averse children with autism. Specialists frequently alluded to or specifically requested customizable features that could be adjusted by the child’s caregivers and/or therapists. They also described a variety of roles that a tactile-sensing robot could perform to support a child with autism, including a teacher, a companion, a tool for regulat-

ing emotional state, and a potential tool for communication. Based on these findings, we provide the following four contributions:

- Informed by our in-depth interviews with eleven autism specialists, we recommend **seven key touch-sensing requirements** that a robot should meet in order to be a good companion for children with autism.
- We translate these qualitative requirements into **quantitative touch-sensing specifications** for human-robot interaction (HRI) researchers and future robot creators.
- We present **guidelines for the robot itself** to promote successful interaction, including its role, responses, and form.
- Additionally, we identify **areas** of the robot a child with ASD is most likely to touch, and the type of **gestures** they are likely to utilize.

The chapter is organized into the following sections: Section 3.1 describes our methods for deriving a set of preliminary touch-sensing requirements from the literature, followed by detailed procedures of our interview study and analysis. Section 3.2 presents the three overarching themes that emerged from our analysis of the interviews. Section 3.3 synthesizes these themes into seven qualitative requirements and translates them into quantitative specifications for touch-perceiving robots. Section 3.4 discusses all of our results and their implications on the future of socially assistive robotics in the autism community.

## **3.1 Methods**

We derived an initial list of requirements for a tactile-sensing robot companion by thoroughly reviewing the relevant literature that was summarized in Section 2.1. In order to validate and fine-tune these initial requirements, we then conducted an interview study with autism professionals from a wide range of specialities and with several years of field experience with children across the spectrum. After conducting the interviews, we transcribed, coded, and analyzed the data using an approach called thematic analysis. We then augmented our initial requirements based on the experts' input. Finally, we further utilized the literature and established methods in the fields of robotics, sensors, signal processing, and machine learning to translate the resulting qualitative requirements into a set of proposed quantitative specifications for tactile-perceiving systems in companion robots.

### **3.1.1 Initial touch-sensing requirements**

We explored a breadth of papers related to robot-mediated autism intervention and studied the variety of physical child-robot interactions that occurred in each study, whether

Table 3.1: Initial list of six touch-sensing requirements that we derived from the literature. Specific sources that inspired each requirement are cited in the “Motivation” column. For each requirement, we also show the follow-up question that the experimenter asked to spark additional feedback after completion of the ranking task.

<b>Requirement</b>	<b>Motivation</b>	<b>Qualitative Description</b>	<b>Follow-up Question</b>
<b>Gesture Identification</b>	Touch applied to an interaction partner can be used to relay immediate physical needs or requests (deictic or instrumental gestural communication) [74, 87].	The robot should be able to feel physical communication gestures, such as a poke.	What kinds of physical communication gestures should the robot be able to feel? Participants could select gestures from the touch dictionary by Yohanan et al. [136]. We omitted one gesture (“finger idly”) from the 30-item list, to avoid participant confusion.
<b>Sensitivity</b>	The child may use soft, gentle touches to communicate positive feelings [107, 74].	The robot should be sensitive enough to detect affective touch communication, such as a gentle hand rest.	What kinds of affective touch communication should the robot be able to feel?
<b>Spatial</b>	Children with autism respond best to direct rewards. When a child touches the robot’s tactile sensor, the robot’s response serves as a reward [106, 29, 50, 105].	The robot should detect touch across a high proportion of its body, so that it may react and therefore encourage future interaction attempts.	How much of its body should be touch sensitive?
<b>Temporal</b>	If time elapses between the child’s touch and the robot’s response, the child may not form a meaningful connection between their touch action and the robot’s resulting behavior [50, 71, 105].	The robot should provide a near-immediate response to touch interaction, similar to a human’s responses.	What kinds of responses should a touch-sensitive robot provide?
<b>Adaptation</b>	A variety of assistive robots are being studied for use in robot-assisted therapy [107, 74, 111, 29, 64, 124].	A general tactile sensing system should be easily adaptable to fit robots of different shapes and sizes.	What kinds of robots do you think are most important to be able to use? (Adjectives, such as animal, humanoid, big, small, etc.)
<b>Robustness</b>	Children with autism may be rough during an interaction [74, 29].	The tactile sensing system should be robust and keep working properly even after rough treatment.	What kind of rough treatment would you expect?

they were planned or spontaneous. We started this process by searching for relevant works in review papers, such as those by Begum et al. [11] and Cabibihan et al. [29]. Once we found a relevant publication that discussed the concept or instance of physical interaction, we then searched for additional sources in both its forward references (new publications that cite the article) and its backward references (older studies cited in the article). Table 3.1 shows the initial requirements that we derived from this examination of the literature. We used this requirement list as a major discussion point for our interviews with autism specialists. As detailed later in this section, we asked each participant to both rank the requirements we proposed and share suggestions for their improvement. For each requirement, we also identified the gaps in the literature and included them as follow-up questions for the specialists. The follow-up questions can also be seen in Table 3.1 but were not shown to the participant during the ranking task.

### **3.1.2 Participants**

We focused on accessing a breadth of autism expertise by recruiting specialists with different focuses through email and snowball sampling. We started by recruiting professionals in applied behavior analysis and speech-language pathology. After multiple participants stressed that their occupational therapist (OT) colleagues would be very well suited to answer the study questions, we decided to recruit OTs as well. We ultimately interviewed eleven participants (4 male, 7 female). Four were occupational therapists, and seven were in other occupations.

All eleven participants were located in the United States or Canada. Their occupations included adaptive physical education teacher, board-certified behavior analyst, neurology professor, occupational therapist (two with focuses on pediatrics, one focusing on early childhood special education, and one on sensory processing disorders), paraprofessional/paraeducator, physical therapist, relationship development intervention consultant, and speech-language pathologist. Participant age ranged from 26 to 55 years old, with a mean of 41, a median of 46, and a standard deviation (SD) of 12. Their years of experience working with children with autism ranged from 2 to 30 (mean: 13, median: 10, SD: 10). Table 3.2 summarizes their backgrounds, as self-reported during the interviews.

### **3.1.3 Procedure**

This research study was approved by the Ethics Council of the Max Planck Society under the Haptic Intelligence Department’s framework agreement as protocol number F003A. Participants gave informed consent and did not receive compensation for their participation. Each interview was conducted in English over WebEx, which is a secure video-conferencing system [39]. Interviews generally lasted about 60 minutes. As an exception, the interview with P10 had a duration of about 100 minutes, as this participant was excited to share many examples from their own experiences. The audio and video of

Table 3.2: Summary of participant experience and backgrounds.

**P1** was a board-certified behavior analyst (BCBA) who also had experience as a paraprofessional and as an applied behavior analyst (ABA). P1 worked in several settings over 7 years including a residential facility, in home services, in various school districts, and in a school specifically for students with an autism diagnosis.

**P2** was a Ph.D. student in Human Development and a relationship development intervention consultant at a private practice. P2 studied inclusive work for autistic individuals entering the workforce and coached parents who are adjusting to their child's autism diagnosis. P2 had 10 years of experience with children with autism.

**P3** was a speech-language pathologist (SLP) in a public school who also previously worked as a paraeducator (also referred to as a paraprofessional). P3 had about 8 years of experience with autistic children ages 5 to 10.

**P4** was a paraprofessional who assisted children with special needs with navigating their day at school. P4 worked in a public school district, in an elementary school in a typical classroom setting for 5 years and in a self-contained autism room for 2 years, and later in a middle school classroom setting for 8 years. P4 has 5-7 years experience specifically with children with autism.

**P5** was a neurology professor who conducts interdisciplinary research on various aspects of autism, including genes, brain development, social interaction, and treatment studies. P5 has roughly 20 years of experience working with children with autism.

**P6** was a pediatric occupational therapist (OT) who worked in a pediatric clinic and a day treatment center specifically for students with an autism diagnosis, both in one-on-one and group therapy settings. P6 had 3 years of experience working with special needs individuals from ages 0 to 22.

**P7** was an OT in an early childhood education setting in a public school district. P7 has 15 years of experience with children with special needs, with and without autism diagnoses.

**P8** was an adaptive physical education teacher who teaches physical education class for children with special needs. P8 has 17 years of experience with children with autism, working both in a public school setting and at an alternative behavior school.

**P9** was a physical therapist (PT) in a public school district with two years of experience working with children with autism. P9 uses physical therapy to help students access their educational environment and curriculum.

**P10** was an OT at a school specifically for individuals with an autism diagnosis between ages 4 and 21. P10 has 30 years of experience, worked with children with autism in 4 different countries, and specializes in sensory processing disorders.

**P11** was an OT in a public school system that serves individuals with autism ages 3 to 22. P11 has 30 years of experience working with children with autism in a variety of settings, such as a residential summer camp, an early intervention program, and home-based instruction.

these interviews were recorded with each participant's explicit consent. The interviews were then transcribed for coding and analysis.

The scripted interview questions were designed to gather detailed feedback in a structured manner. The experimenter guided each autism specialist to provide information on their education background and work experience, describe the touch behavior of autistic children, and rank and comment on touch-sensing and response requirements for a companion robot. A description of the six stages of the interview, as well as the general time allotment for each stage, is provided as follows:

1. **Study setup and demographic questionnaire (10 minutes)** – The experimenter explained the motivation behind this research, the goals we hoped to accomplish through the study, and the general setup for the interview session. We then collected demographic data from participants. Next, we asked participants about their experiences in relation to children with autism, including current and previous related jobs, years of experience, and countries in which they had worked with children with autism.
2. **Children with ASD and general touch interactions (15 minutes)** – We asked participants to comment on how children with autism utilize touch, including whether there are differences compared to neurotypical peers, as well as similarities and differences within the autistic population. We also asked how participants respond to children's requests that are presented through touch, and what methods they have used to calm an autistic child who was stressed. We then showed them two examples of a NAO robot (Figure 3.1), with one wearing a thin shirt and the other wearing a koala suit, and we asked what types of physical touches the robot should be able to detect. To show the robots, the interviewer either redirected her camera or shared a picture of the robots in the WebEx browser. In both scenarios, the robots were powered off and in the same pose.
3. **Ranking task (15 minutes)** – We sent the participant a link to the requirement-ranking interface, where the initial requirements were presented in an order that was randomized for each participant. We asked the participant to rearrange the requirements, ranking them from most to least important, and to explain their thought process out loud. The ranking task interface was a shared online presentation application similar to Microsoft PowerPoint (Figure 3.2). Each initial requirement from Table 3.1 was displayed on its own slide, along with a brief explanation of the concepts that motivated the requirement. Participants could rearrange the requirements into their desired order by clicking and dragging the slide navigation thumbnails on the left side. If a participant was not able to navigate the ranking interface, they could also tell the interviewer the sequence in which they wanted the requirements rearranged. Participants were encouraged to tell the interviewer if they felt a requirement should be edited, removed, or added.



Figure 3.1: The prototype robot companion for children with autism that was presented to the therapists in our interview study. It is a NAO robot (SoftBank Robotics); the NAO robot shown on the left is wearing a thin fabric shirt, and the one on the right is enclosed in a soft koala suit. Photo redistributed from [22] under a Creative Commons license – Attribution 4.0 International (CC BY 4.0).

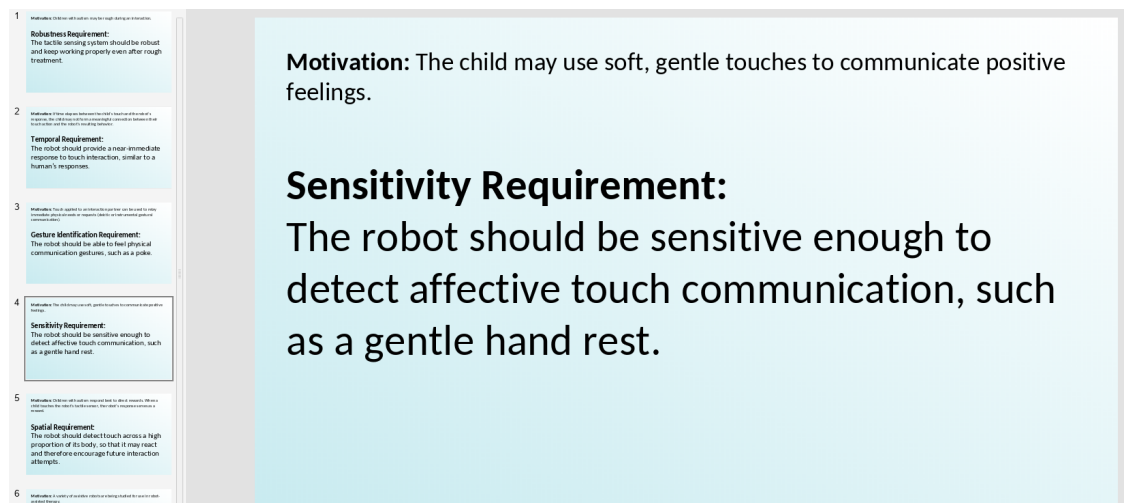


Figure 3.2: The requirement-ranking interface enabled participants to prioritize tactile-sensing requirements for the robot as they preferred. Participants could reorder the six provided requirements by clicking on and dragging the respective slides in the navigation pane on the left. The main view enabled the participant to clearly see the individual requirements and the motivations that inspired them. Illustration redistributed from [22], CC BY.

4. **Follow-up questions (10 minutes)** – Next, we asked the participant a specific follow-up question for each requirement (Table 3.1), in order to initiate a more detailed dialogue and derive further specifications. While the initial requirements were worded generally, the follow-up questions prompted the participant to give specific details about their ideal robot companion. The question order followed the participant’s ranking, from most important to least important requirement.
5. **Comments on tactile sensor design (5 minutes)** – To elicit further comments on existing sensing technology, we showed the participant a prototype fabric-based tactile sensor (Figure 3.3) and explained how it worked in layman’s terms. We asked for their thoughts on the design, as well as related questions and concerns.
6. **Closing questions and recommendations (5 minutes)** – Finally, we asked the participant what movements and sounds the robot should provide in response to a child’s input. We also asked what other factors we should consider in creating a touch-sensing robot companion for children with ASD, and what closing comments or questions the participant had for us.

Similar to how we used the initial touch-sensing requirement list in stage 3, the robots presented in stage 2 and the tactile sensor presented in stage 5 were meant to provide tangible starting points for the discussion. Given available resources and time limitations, we presented only one type of robot and one sensor design. However, we encouraged




Figure 3.3: A photograph of the prototype fabric-based tactile sensor for a robot companion that was shown to participants. Participants were told that several such sensors could be placed across the robot's body, and that they could be used to detect the intensity and/or type of touch being enacted on the sensor. Participants were asked to give feedback and share any concerns they might have with such a design. Photo from [22], CC BY.

the participant to think about a variety of form factors and approaches in both cases, not limiting their thinking to the presented examples.

### 3.1.4 Qualitative analysis of interviews

We utilized thematic analysis to analyze the data from the interviews. Thematic analysis is a method for interpreting and organizing qualitative data, such as a series of interviews, into meaningful patterns [19]. A block diagram illustrating our workflow can be seen in Figure 3.4. After transcribing each interview, two authors analyzed each complete transcript line-by-line and labeled the participant's responses with descriptive codes, using an open, iterative coding approach. The two coders met frequently to discuss findings and compare codes. We started to notice repetitions across participant responses and data saturation around the seventh participant. At this point, the coders began to aggregate all the codes into related groupings in order to identify the overarching themes. The coders separately searched for themes, and they then discussed findings and merged theme results, applying focused coding methods to finalize the analysis. We converged on three themes that captured the vast majority of the comments shared by our eleven participants.




Figure 3.4: A diagram illustrating the flow of our process and the links between our literature review, interview study, and resulting themes. The thematic analysis portion of our study is highlighted in green. The identified themes were used to refine our initial requirement list into finalized qualitative requirements, which we then translated into quantitative specifications. Figure adapted from [22], CC BY.

## 3.2 Results: Overarching Themes

We present the three overarching themes that we identified in our data: the touch tendencies of children with autism, the importance of individual differences, and the role of therapists in each child’s development. Within each theme we report the relevant requirements that the participants described for both touch sensing and child-robot interactions. Quotes taken directly from a participant’s transcribed interview are marked in italics and enclosed in double quotation marks. We use bold to highlight terminology from the autism specialists that may be new to the reader.

### 3.2.1 Touch in autistic children

Children with autism frequently experience sensory stimuli differently than neurotypical children. Several of our participants explained that children with autism are often on one of the two tail ends of the touch-sensitivity spectrum. As P5 described, “...if there is a distribution [of touch preference], the kids with autism will be out here on the two tails, and typically developing kids will fill that average ... you would hardly ever meet a typically developing child who just seemed to really crave being touched, or hating it, but with autism you generally won’t find them to be in the middle.” Participants used several different words to describe these two ends of the distribution, including under- or over-responsive to touch, hypo- or hyper-sensitive to touch, loving or hating touch, and touch seeking and touch averse. We refer to them herein as “touch seekers” and “touch avoiders”.

**Touch seekers** are hypo-sensitive to touch. They crave contact and use it to investigate their surroundings and function in the world. They enjoy deep pressure touches, like squeezes and firm hugs. **Touch avoiders** are hyper-sensitive to touch. For them, touch can be upsetting, unpleasant, and even painful at times. Touch avoiders may use a very light touch, or get very close to the person or object of interest without touching. If there is touching, a touch avoider would prefer to control and initiate the interaction, such as guiding another person’s hand to request assistance. While a child will generally display

traits for one of these tail ends (a touch seeker or avoider), the child's receptiveness to touch also fluctuates based on external factors, such as events in their day and their familiarity with the interaction partner or object. P6 explained this topic particularly clearly, stating, "*...just like the autism spectrum, touch and all of our sensory systems are also on a spectrum... Some kids [with autism] seek out touch as their way of functioning in life. Some kids completely avoid it... a lot of our kids who are over-responsive to touch kind of hold back. They don't use it as much. They have a really low threshold, so it can be upsetting to them. They don't like to be touched in certain places. They don't like to touch things with their hands and they could be super over-responsive to heat, pain, cold, stuff like that. And then on the other hand, people who are under-responsive, as you probably know, they kind of seek that out... they'll touch literally anything to kinda get that input to their brain of what they're doing, how they're doing it, and kind of function throughout their day.*"

Touch can help children with autism explore the world, self-regulate, and communicate their needs. To investigate objects, autistic children frequently use their hands, their mouth or lips (P8: "*...There is a lot of mouth in my class.*"), and even their entire body (P2: "*I have other clients that like to take fuzzy things and roll in them.*"). While touch seekers typically need firm squeezes to calm down, touch avoiders may also use touch for self-regulation, utilizing predictable and controlled interactions like hand holding. Children with autism, regardless of whether they are a touch seeker or avoider, often have favorite toys or sensory items. Schools and therapy clinics usually also have a variety of objects that can create different tactile sensations for the children to choose from. Finally, children with ASD may be non-verbal or have limited speaking skills. Therefore, touch becomes especially important for communication. They may use touch to communicate needs (e.g., guide the caregiver to an object), to express feelings (e.g., demonstrate their distress), and to socialize with others (e.g., seek attention or space).

Importantly, children with ASD often do not understand social conventions. As such, they may use socially inappropriate touch. When interacting with someone, a child may get very close to the person – in their personal "*bubble*" (P2). If they find a feature or object on the person's body interesting, such as a mole or an ID badge, they may touch or grab it without asking. They may unknowingly touch private body parts or use inappropriate touch gestures (e.g., a slap or pinch), typically with the wrong pressure level or a high frequency. P6 explained, "*...a lot of students will go up to another person and when they're trying to be really nice and you know, to tap them on the shoulder, they're actually going to give them a good slap because they're not fully understanding their touch, their proprioception.*"

#### **Sensor and robot recommendations**

The participants recommended requirements that would enable a tactile-sensing robot companion to work for touch seekers and touch avoiders alike. The tactile sensors should be able to detect a wide range of contact intensities, from the light touches of touch

avoiders to the deep squeezes of touch seekers. Sensitivity should be consistent across a single sensor's area and also across different sensors. Additionally, depending on the child or children using the robot, the same type of gesture might occur with different intensities.

The robot should detect touch across a wide range of its body to encourage future touch interactions, especially for touch avoiders. 10 out of 11 participants requested at least 80% of the robot's body be touch-sensitive. They said that the location of the touch, not just its intensity, should be detected. They would want the robot to know the general region of its body that had been touched, such as the arms or face, or whether the child had touched a region that was appropriate or inappropriate for social interaction. Notably, the participants did not ask for exact contact location discrimination within a general body region. Often, while participants preferred the idea of whole-body sensing, they also specified which body regions they would prioritize to equip with touch sensing, in case whole-body sensing was not possible. Figure 3.5 presents a summary of the locations suggested by the participants.

The robot and its sensors must be very robust. They must survive rough exploration, excited interactions, and tantrums. The sensor material should be durable and not tear easily. Sensors should be adhered securely to the robot to withstand rough squeezing, picking and pulling. The robot and its sensors as a whole system must also be safe for oral exploration, without sharp edges or risk of electric shock.

The tactile feel of everyday items or the robot itself became an important topic of discussion during many of our interviews. Four participants (P1, P3, P5, P8) noted that children with autism may be sensitive to the feel of their clothes or objects (e.g., the tag on their shirt, the feel of jeans, or the texture of a stuffed animal). Five other participants (P2, P4, P6, P7, P10) emphasized that much attention should be paid to the passive feel of the robot. They recommended the robot provide a range of tactile sensations, such as a firm base and a soft and squishy outer covering. Visibly and tactilely noticeable seams and wiring should be avoided. The outer texture of the robot should have a soft, neutral sensation – not extremely fluffy or rough.

### **3.2.2 Understanding individual needs and differences**

All of the participants referred to the wide range of characteristics that can be found in children on the autism spectrum. These individual needs and differences need to be taken into consideration when designing a robot companion.

The specialists noted that children on the spectrum may differ in their preference for touch sensations, their ability to utilize touch gestures, and their processing of sensory input. Specifically, ASD children may fixate on a favorite body part or tactile sensation and use that repeatedly. They may have other diagnoses and symptoms comorbid with their autism that affect how they are able to interact with the robot companion. For example, a child may have low muscle mass, have poor motor coordination, fatigue easily, or have a vision impairment. Thus, the child may be able to apply only a light

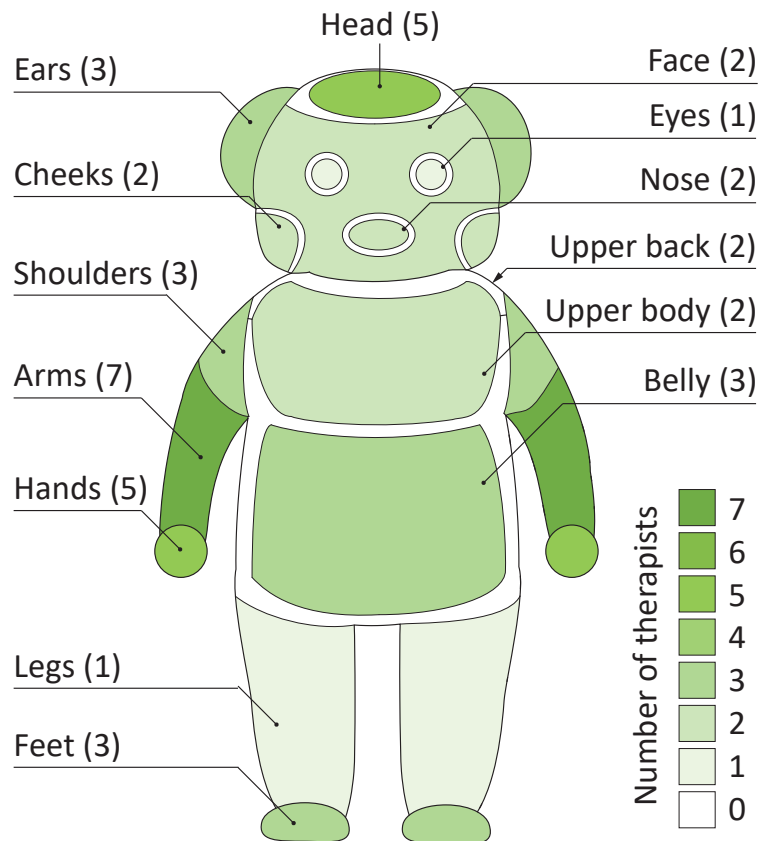


Figure 3.5: Locations that participants specifically prioritized to equip with touch sensing, displayed on a diagram of a generalized robot companion. The number of therapists who suggested each region is shown in parentheses. Most of our eleven participants requested whole-body touch sensing, but they also provided these specific locations in case whole-body sensing is not possible. Illustration from [22], CC BY.

touch, use their whole hand instead of a single finger, or touch a general region on the robot rather than a specific area. Finally, children on the spectrum may need additional time to process information or sensory input. As explained by P5, *“If you could code social interactions between kids with autism and adults, it’s almost as though the person with autism is one step behind, even though they are responding, and they get kind of out of phase.”*

Children with autism communicate in a variety of ways. Their level of verbal communication can range across non-verbal, non-verbal using augmentative and alternative communication, and verbal. They also may have developed custom sign language or gestures to communicate with their caregivers. They will have different cues to express distress or alert the start of a *“meltdown”* (P1) – when the child is very upset, their senses are overstimulated, and further information cannot easily be processed. Therapists work closely with each child to understand their needs and communication methods. When a therapist first begins working with a new client, much of their communication with the child may be guesswork.

Like any child, a child with autism will have their own interests – a favorite color, toy, cartoon character, etc. They will have their own motivators, skills, and fears. The child’s education program is customized by their individual therapists and education team. The therapist encourages the child to complete their therapy session by using **motivators** – activities or items the child likes, which can therefore be used as a reward. Motivators are personalized for each child. Therapists constantly monitor and update motivators to reflect the child’s current interests.

### **Sensor and robot recommendations**

Most participants either alluded to or explicitly described the need for customization, agreeing that customizability will be key for creating a robot companion that is adaptable to children across the autism spectrum. They requested the option for the parent or therapists to customize several features, such as the force levels and types of gestures the robot can detect, the timing delay of the robot’s responses, and the robot’s physical appearance.

Along with touch avoiders, some children may also have physical conditions that restrict how much they can control their physical interactions with the robot. As such, the participants reiterated the importance of the tactile sensors reliably detecting touch, whether it is a very light touch or a tight squeeze, so that the child can be rewarded for their interaction attempt in either scenario. Additionally, some children cannot apply fine-grained gestures, and the implied meanings of their gestures may vary. Consequently, the participants suggested that the parents or therapists be able to direct the robot about their child’s actions. They could specify what touch gestures the child commonly uses or avoids. In addition, they could specify the intent of a gesture – a hard slap might be a friendly interaction from some children and a negative interaction coming from others.

The timing of the robot's reactions was our most controversial point of discussion. Participants were divided on how quickly the robot should respond, wanting either a near-immediate response or a time delay. Several participants suggested a compromise – equipping the response setting with a customizable time-delay window. It is important to note that the delay was requested only for the robot's response to the touch. No delay was requested for the speed at which the tactile sensors detect the child's touch.

One group of participants strongly agreed with our initial temporal requirement that the robot should provide a near-immediate response to touch (P1, P4, P6, P10, P11). They explained that this near-immediate response would help the child form a direct correlation between their actions and the robot's responses. A near-immediate response would also provide tangible positive reinforcement for the child's touch and encourage additional interaction. Finally, P1 said that delaying the response could reinforce the wrong interaction, especially if the child was frustrated by the robot's silence and began using more force.

Other participants disliked the “near-immediate” wording of our initial temporal requirement (P3, P5, P9). They recommended having the robot respond after a time delay. As the child might have a processing delay, a near-immediate response might be confusing to the child, happening too early or too fast for the child to understand. Participants with this opinion often recommended that the robot respond with a customizable time delay, as set by the parent or therapist. Finally, P3 was concerned that the child could become dependent on the fast response and noted that the caregiver could gradually increase the time delay to help the child build tolerance. The remaining participants (P2, P7, and P8) did not provide specific comments for or against the “near-immediate” timing of the temporal requirement.

Customization of the robot's appearance is desired, but it is less important than other requirements. Participants differed in opinion on what form factor is most desirable for a robot companion. Some participants preferred an animal appearance (P1, P4, P6, P7, P8, P11). They suggested interacting with an animal would be more inviting and calming than interacting with a humanoid partner. Others further specified that a toy-like, stuffed animal form would be ideal, citing that children with autism often already have a treasured particular toy or stuffed animal. As P7 explained, “*Well, just make it something that they really want to hold onto... Something that they want; they want it to be their special animal, their special fuzzy, you know? Kids cling onto an animal for forever.*”

Other participants suggested a humanoid appearance (P5, P11). While P1 felt that a humanoid robot would be perceived as creepy, P5 suggested that interacting with a humanoid companion robot would provide good social interaction practice and enable the child to transfer skills to human interactions with greater ease. P5 stated this preference as follows: “*Humanoid for sure, because ...you know, I think what you are trying to train up is using touch in a social or communicative sense.*” Finally, P11 noted that an animal form would be more comforting for children, while a humanoid form might be more appropriate for children at the high-functioning end of the spectrum, as well as when a

child's treatment plan progresses.

Other participants suggested giving the robot a form factor similar to objects that specifically interested that child, such as a toy truck or a cartoon character (P9, P10). Many participants suggested having several interchangeable options for an outside skin on the robot. Such a design would enable therapists to easily customize one robot's appearance for several different children. A washable outer skin was also recommended to keep the robot sanitary. Several participants noted that the robot should be a portable size (P3, P4, P6, P7, P10), and one emphasized that the tactile sensors should be scaled to match the robot's size (P5).

### 3.2.3 Supporting the individual and promoting independence

Therapists help children with autism learn socially appropriate behaviors, navigate meltdowns, build tolerance to uncomfortable stimuli, and eventually cope with these stimuli through self-regulation skills. Regardless of the exact details of their occupational title, their underlying goal is always the same: to support the child and to promote the child's independence. As P1 described, *"We'll talk to the parents, we'll talk to pertinent people in their life, to try and determine what are the skills that are lacking...because our goal is always to have them achieve their highest functioning level, whatever that might be at that moment in time."*

Therapists help children with autism improve their communication skills by using a variety of methods such as gestures, picture cards, sign language, alternative and augmentative communication (AAC) devices, and speech modeling. Therapists also often teach children with autism proprioceptive skills, such as understanding their own body parts and the strength of their touch. Therapists improve these skills through a process called shaping.

**Shaping** refers to the therapist guiding the child toward an ideal behavior, often through several incremental steps of accepted behavior. For example, a therapist may encourage a child to change their mode of communication over time, gradually migrating from using a guiding touch on their caregiver's arm, to selecting a picture card symbolizing their request, with the end goal of asking verbally, if possible. The therapist may also help shape the child's touch behavior, guiding them to utilize socially appropriate intensity, duration, and location when touching others. The therapists may use a *"token system"* (P1, P3) to help shape behavior: every time the child successfully completes a therapy task, they are awarded a **token** – a positive visual marker such as a sticker or a check mark – on a token board. When the child finishes the therapy activity and fills the token board, they are rewarded with one of their motivators.

Therapists utilize various methods to safely navigate a child with ASD out of a meltdown. To start with, the therapist will often reduce sensory input if possible, such as dimming the lights or moving to a quiet room. As the child may have a difficult time processing additional input during a meltdown, the therapist will limit their talking, using language that is as simple and direct as possible. They may provide the child with

picture cues, so the child can express what is wrong or what they need. The therapist may provide positive distractions in the form of calming physical sensations (e.g., deep hugs or squeezing, access to toys with different sensations, using a swing, riding a wagon, or taking a walk). A therapist can also try to remind the child of the awaiting motivator to help them get through the situation.

Ultimately, the therapist aims to help the child build their independence. They seek this goal by giving control and choices to the child where possible. They help build tolerance to unpleasant stimuli. They provide tools and teach strategies that the child can later use on their own. The goal is to help the child thrive and become as self-sufficient as possible. As articulated by P3, *“You’re trying to build up that tolerance ability to the unpreferred or less preferred situation. So I think, as you know, in the education setting, we’re doing as much as we can to help them gain independence, or gain the strategies that they will need to be functional. Just that little bit at a time.”* In order to reach this goal, the therapist must first build rapport, establish a relationship, and gain the child’s trust. Trust is key for building the child’s tolerance to activities or objects they dislike, and to help a child when they are particularly distressed.

#### **Sensor and robot recommendations**

The participants discussed the various roles the robot companion could play in order to support the child and promote the child’s independence. Participants seemed to see the robot serving as a teacher, a companion, a tool for regulating the child’s emotional state, or a tool for communication. These roles are not necessarily exclusive from one another: some participants thought the robot could fulfill multiple roles.

Several participants saw an opportunity for **the robot to act as a teacher** (P5, P7, P10). P10 praised the idea of using a robot to reinforce concepts from educators through play, saying, *“You have to give [children with autism] time to play and process if they’re behind... you gotta give them time to process it, integrate it, to feel comfortable, to repeat it... the summary is very important. So, I think tools like a robot could be that, that moment where they get to practice the goals that are coming from the educators, teachers, and therapists. So I think it’s very, very important. So they can repeat and do it in a more calm and enjoyable way.”* The robot could be used to teach about body parts and promote understanding of the child’s own body (P7). The robot could also teach about socially acceptable locations to touch others. The spatial settings could be adapted and reduced over time to promote touching the robot in areas that are socially acceptable for human interactions.

Additionally, the robot could teach what kinds of gestures are socially acceptable to use when touching others. The robot could slowly shape the appropriate touch type, location, intensity, and frequency during interactions with the child. Certain reactions, such as clapping and positive verbal responses, could be used to reinforce positive behavior. Participants gave a variety of suggestions for how the robot should react to less desired touches, such as giving a firm verbal response for negative behavior, or turning off and




Figure 3.6: A visualization of the touch gestures that specialists recommended (or did not recommend) a touch-perceiving robot companion should be able to detect from children with autism. An asterisk above a gesture indicates that it was selected as a “top five” gesture by more than three participants. Figure from [22], CC BY.

giving no response at all. The robot’s gesture identification feature could help reward the appropriate type of touch at the start of the behavior shaping, even if the intensity is wrong (e.g., a tight hug is okay at the start, because a hug is a socially acceptable gesture). P5 suggested matching the force detected by the tactile sensors to the human perceptions of pleasure and pain. They suspected that matching the robot’s responses to expected human responses at the same force levels could help the child transfer their practice with the robot to human interaction. Figure 3.6 presents the touch gestures that the therapists observed in autistic children and the top five gestures that they recommended the robot to detect. While most of the gestures are from the touch dictionary by Yohanan et al. [136], some participants also suggested gestures of their own: “light touch”, “squish”, “heavy catch”, and “gentle hand rest”.

Multiple participants requested that the robot detect if a gesture is repeatedly occurring, in order to prevent the robot from quickly repeating a response over and over. A repeated gesture could be an indicator that the child is stimming. Even if it is not a stim behavior, repeating a poke over and over, for example, is not a desirable social interaction.

Some participants saw **the robot as a companion** for the child (P3, P4, P5, P6, P11). They felt it was well suited to act as a friend and source of comfort. P4 extended this idea to helping the child befriend other children, saying “*But that would be awesome because, you know... it’s hard for those kids to make friends, and this could be a friend, you know, an extra friend for them... and I can see it drawing attention too, to get other children interested in communicating more with that child also.*” Therapists highlighted that the robot should behave in a manner that was calm, reassuring, and gentle. In particular,

the child's initial interaction with the robot is very important to gaining trust. P3 and P4 suggested that the robot should start with very small predictable responses or no movement at all, so as not to frighten the child. P4, P5, and P11 suggested the robot build rapport by reciprocating and mirroring the child's communicative actions, such as giving and receiving hugs, holding the child's hand, and playfully poking back. They also suggested other actions the robot should perform, such as greeting the child and looking in the direction of interest (either at the child or where the child touched).

Participants also suggested that **the robot could serve as a tool for regulating emotional state** (P1, P7, P9, P10, P11). Three participants (P7, P10, P11) felt that the robot could calm the child down if they were overexcited, perhaps by giving hugs, playing calm music, or playing white noise. They also suggested having the robot play comforting custom audio messages recorded by the child's family members. Conversely, the robot could help energize the child if they were feeling lethargic, using colorful lights (P10) and singing and dancing (P7, P9). Two participants (P1, P7) also suggested the robot could replicate some of the cool-down methods therapists use to navigate the child out of a meltdown.

Finally, participants saw **the robot as a potential tool for communication** (P6, P8, P9, P11). The robot could encourage communication, perhaps acting as an AAC device. The child could use the robot as a safe companion for practicing communication requests. The robot could use its cameras and/or touch sensors to identify a child's touch requests, and it could then verbalize those requests out loud (P6). The robot could also use its sensors to detect how the child is feeling and then verbalize this observation, to help give the child vocabulary for what they are feeling (P11).

## 3.3 Results: Qualitative and Quantitative Requirements

Building on the aforementioned themes, we present seven qualitative requirements and further translate them into quantitative specifications for a touch-perceiving robot.

### 3.3.1 Qualitative tactile-perception requirements

The results of the requirement-ranking task can be seen in Figure 3.7. Five participants explicitly stressed the importance of reviewing and carefully selecting the tactile properties of the robot companion and its tactile-sensing system, and an additional four participants separately mentioned sensitivity in autistic children towards certain tactile textures and sensations. We therefore added the *feel* requirement to our initial list. Based on the participants' input, we provide a finalized version of these seven qualitative tactile-sensing requirements in Table 3.3. We revised the requirements' initial descriptions from Table 3.1 to better reflect the recommendations gathered from participants. Additionally, the robustness requirement was renamed to ***robustness and maintainability***, and the sensitivity requirement was renamed to ***sensing range***, in order to better match their revised

descriptions. We bold and italicize the finalized requirements in Table 3 and 4 and in the text to allow readers to distinguish them from our initial requirement list. The order of the final requirements corresponds to the median ranking across all participants. As the *feel* requirement was consistently requested or implied by almost all participants, we placed it third in the requirement priority order.

### 3.3.2 Quantitative specifications

Next, we returned to the literature and established methods for implementing touch perception to translate our qualitative requirements from Table 3.3 into the quantitative specifications shown in Table 3.4. Rather than limit the translation of qualitative requirements into specifications that work for only one particular tactile-sensing technology, we have attempted to provide quantitative specifications that are as “approach agnostic” as possible. It is our hope that engineers developing many different sensing technologies and processing methods could utilize these specifications to create successful robot companions. Below, we describe the rationale and terminology associated with each requirement in further detail.

In terms of *robustness and maintainability*, we propose that the tactile-sensing system must remain fully functional, at minimum, for at least the total duration of a child’s therapy session. This specification ensures the robot does not stop working during a child’s session, which could distress them and discourage them and the therapist from utilizing the robot in the future. The typical duration of an autism therapy session was found to be between 30 and 45 minutes, with some sessions lasting even longer [133]. The sensing system and robot itself must also be able to withstand the maximum pressing force of a human finger (30 N) without breaking [91]. Furthermore, it should continue to operate normally even after being subjected to such treatment. Lastly, all devices used with children must follow applicable local safety standards.

The *sensing range* requirement indicates that the robot should be sensitive enough to detect a wide range of contact intensities, from light touches to deep squeezes. We propose minimum force detection for a light touch based on the force reported to deliver a light, pleasant touch during affective touch studies, 0.4 N [127]. For the high end of our force sensing range, we suggest that the sensor should be able to detect at least the force at which the human arm is in pain, 25 N [88]. These values help align the robot’s force sensing capabilities with those of a human, which can in turn help a child learn socially acceptable touch through interaction with the robot. The standard minimum acceptable signal-to-noise ratio in sensor development is set to be 3.3, following the limit of detection [6]; we thus propose this value as a reasonable guideline for ensuring good performance across the sensing range.

The sensing system should enable reliable *gesture identification* despite different intensity levels, execution rates, and locations. As such, we believe there should be less than a 5% change in gesture characterization if the same gesture is applied repeatedly over time, or if the same gesture is applied at a different position on the same sensor.

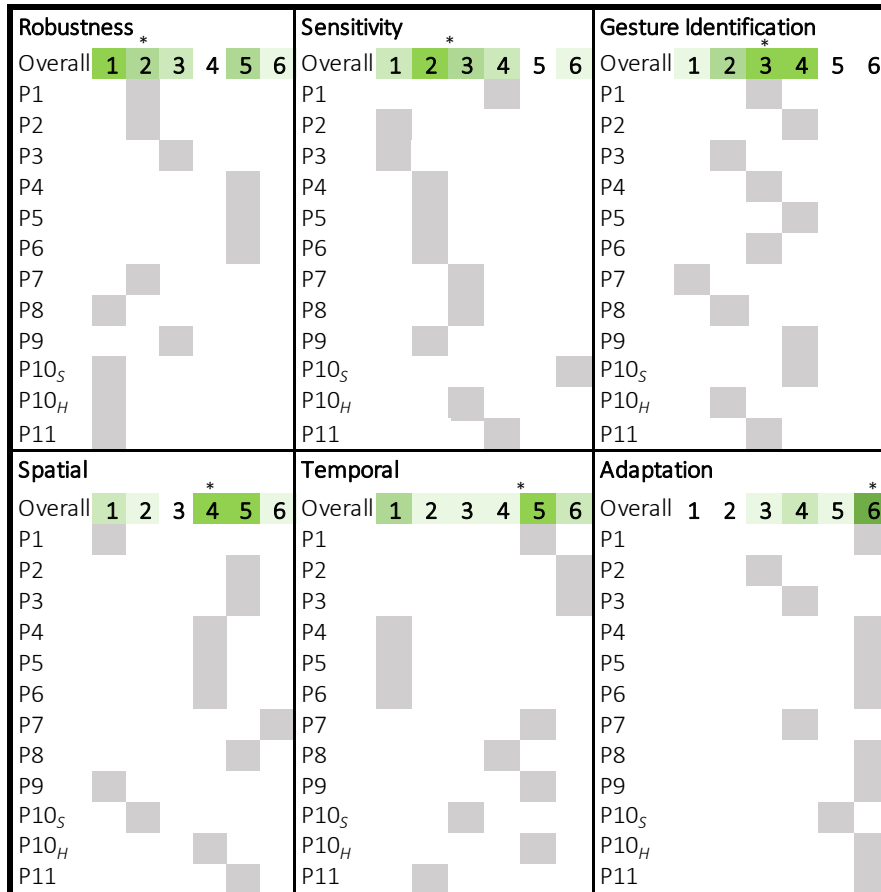


Figure 3.7: The participants’ individual responses for the requirement-ranking task. A ranking of 1 indicates the most important requirement, and 6 indicates the least important. A summary of each requirement’s rankings is also displayed at the top of each panel using a green opacity overlay – the darkness of a number’s green background indicates the number of participants who selected that rank for that requirement. White means that no participants chose that ranking. The asterisk above each overall summary indicates that requirement’s median ranking. Participant 10 completed two different versions of the ranking – one for using the robot in a school setting (P10<sub>S</sub>), and the second for using it in a home (P10<sub>H</sub>). Figure from [22], CC BY.

Table 3.3: Our final qualitative guidelines for a touch-sensing robot companion for children with autism, as derived from the recommendations of 11 autism specialists. Requirements are listed in descending order of importance, with 1 indicating most important.

Requirement	Qualitative Description
<b>1. Robustness and maintainability</b>	The tactile sensing system and robot should be robust and keep working properly even after rough treatment. The sensor material should be robust to vigorous mechanical interactions as well as oral exploration. The sensor attachment and wires should be robust to pulling and rubbing gestures. The sensor and/or its outer cover should be easy to wash and repair by caregivers.
<b>2. Sensing range</b>	The robot should be sensitive enough to detect a wide range of contact intensities, from light touches to deep squeezes, similar to humans.
<b>3. Feel</b>	The sensing system should be pleasant to touch (e.g., soft and squishy). The sensors and wires should be minimally detectable and seamless to touch.
<b>4. Gesture identification</b>	The robot should be able to differentiate physical communication gestures. The five most recommended gestures to detect include hug, poke, squeeze, hold, and tickle. Gestures should be identifiable at different intensity levels, at different execution rates and speeds (e.g., to detect stimming behavior), and at different locations.
<b>5. Spatial</b>	The robot should detect touch across all (or a high proportion) of its body. The robot should discriminate which body part or region was touched but does not need exact contact localization within each region.
<b>6. Temporal</b>	The robot should respond to touch interaction with timing that can be customized to the child's processing needs. The sensor's measuring capabilities should be fast enough to capture all human contacts.
<b>7. Adaptation</b>	The sensors should be easy to scale to different sizes and adapt to different curvatures present on a robot body. The entire robot should also be portable. Adapting to different robot types would be nice but not essential. The appearance of the sensing system should be customizable to the child's preference.

Furthermore, many participants expressed interest in using the robot to teach socially appropriate and inappropriate behaviors. It is thus crucial that the robot can differentiate appropriate gesture types (e.g., hug, stroke) from inappropriate ones (e.g., pinch, slap). Therefore, we recommend a confusion rate lower than or equal to 5% when discriminating between these two categories of gestures; confusion rates between gestures in the same category may be higher. For all three of these quantitative specifications, we have recommended the value of 5% to match the accepted level of error that scientists most commonly use when performing statistical analyses (i.e.,  $p \leq 0.05$  in significance testing).

For the *spatial* requirement, we based our quantitative specifications on the over 80% tactile-sensing coverage of the robot's body and the 14 body regions that the participants requested during the interviews. The robot's arms, head, and hands were the most requested body regions (Figure 3.5).

According to the *temporal* requirement, the robot should respond to touch interaction

### 3.3 Results: Qualitative and Quantitative Requirements

Table 3.4: The minimum quantitative specifications we propose for a touch-perceiving companion robot for children with autism. These quantitative specifications were translated from our final qualitative requirements (Table 3.3).

Requirement	Quantitative specifications
<b>1. Robustness and maintainability</b>	<ul style="list-style-type: none"> <li>• Minimum duration of consistent operation: <math>\geq 45</math> minutes</li> <li>• Minimum force to withstand without malfunction: <math>\geq 30</math> N</li> <li>• Compliant with all relevant safety standards in the country of use</li> </ul>
<b>2. Sensing range</b>	<ul style="list-style-type: none"> <li>• Minimum detectable force: <math>\leq 0.4</math> N</li> <li>• Maximum detectable force: <math>\geq 25</math> N</li> <li>• Minimum signal-to-noise ratio (SNR): <math>\geq 3.3</math></li> </ul>
<b>3. Feel</b>	<ul style="list-style-type: none"> <li>• We did not find any references that provide quantitative specifications for tactile pleasantness.</li> </ul>
<b>4. Gesture identification</b>	<ul style="list-style-type: none"> <li>• Gesture recognition: <math>\leq 5\%</math> confusion between socially appropriate and socially inappropriate gestures</li> <li>• Intensity perception: able to detect each gesture at <math>\geq 2</math> intensity levels</li> <li>• Spatial consistency: <math>\leq 5\%</math> change in recognized gesture for application of the same gesture in different locations on the same sensor</li> <li>• Temporal consistency: <math>\leq 5\%</math> change in recognized gesture for repeated application of the same gesture over time</li> </ul>
<b>5. Spatial</b>	<ul style="list-style-type: none"> <li>• Surface area capable of touch detection: <math>\geq 80\%</math></li> <li>• Number of contact-sensing regions: <math>\geq 14</math> different areas distributed across the robot's body, without needing to localize contact within each area</li> </ul>
<b>6. Temporal</b>	<ul style="list-style-type: none"> <li>• Minimum cut-off frequency for detecting dynamic contacts: <math>\geq 20</math> Hz</li> <li>• Maximum delay between tactile interaction and recognition: <math>\leq 155</math> ms</li> </ul>
<b>7. Adaptation</b>	<ul style="list-style-type: none"> <li>• Minimum sensor size: <math>\leq 9</math> cm<sup>2</sup></li> <li>• Maximum sensor size: <math>\geq 100</math> cm<sup>2</sup></li> <li>• Maximum convex curvature: <math>\geq 0.4</math> cm<sup>-1</sup></li> <li>• Maximum concave curvature: <math>\geq 0.05</math> cm<sup>-1</sup></li> </ul>

with timing that can be customized to the child's processing needs. To accommodate variation across the autism spectrum, we recommend that the tactile-sensing system's bandwidth (low-pass cut-off frequency) should be at least twice as fast as that of the fastest human movement frequency of 10 Hz [96, 53, 56] to adequately capture even energetic or violent touch actions that a child performs. Since textural contact generates large high-frequency vibrations [37], even higher bandwidth may benefit perception of gestures that involve motion across the robot's surfaces, such as stroking or tickling. While the participants stressed the importance of providing a customizable time delay, some children with autism may be capable of interacting at the same rate as neurotypical individuals. In this case, the robot should detect the child's touch with near-immediate recognition. We propose that recognition rate match that of a human's average touch reaction time, around 155 milliseconds, to promote the child's skill transfer from robot interaction to human interaction [90].

For the *feel* requirement, several studies describe tactile pleasantness in qualitative

terms and give examples of surfaces that people commonly perceive as pleasant to touch [32, 127, 31, 48]; however, we were unable to locate any work that provides quantitative specifications for tactile pleasantness.

The *adaptation* requirement states that the tactile sensing system should be scalable for robots of different shapes and sizes. We translated this qualitative goal to quantitative specifications by referencing the measurements of the NAO, a child-sized robot by SoftBank Robotics that has been used in autism research [128, 124, 59] and was shown to participants during the study. We estimate that a minimum sensor size of  $9 \text{ cm}^2$  (3 cm by 3 cm) is small enough to highlight the smallest body regions requested by several participants, such as the hands of the robot. Beyond a maximum sensor size of  $100 \text{ cm}^2$  (10 cm by 10 cm, approximating a robot's belly or upper back), researchers may find that they lose the ability to discriminate between robot body regions, which was heavily prioritized by the participants. The most highly curved robot body part that several participants requested is the arms; each arm of the NAO has an approximate radius of 2.5 cm, which corresponds to a maximum convex curvature of  $0.4 \text{ cm}^{-1}$ . Many robot surfaces have lower curvature than the arms, including flat surfaces with zero curvature. Some robot surfaces that need tactile sensing may also be gently concave. The most highly curved concave surfaces on the NAO robot are at the front of each leg, below the knee; these surfaces have an approximate radius of 20 cm, which corresponds to a maximum concave curvature of  $0.05 \text{ cm}^{-1}$ .

## 3.4 Discussion

This chapter has presented guidelines for tactile perception in a robot companion, considering both the qualitative requirements and the quantitative specifications that resulted from our study. This section reflects on the scientific methods we employed and discusses how our results contribute to the fields of socially assistive robotics and robot-mediated autism intervention. With our key tactile-perception guidelines, we attempt to form a bridge between tactile sensor developers, HRI researchers, and the target populations – children with autism, their families and their therapist teams. Although many socially assistive robots currently exist to help children with autism, they rarely incorporate touch perception and thus typically cannot feel or react to contacts that children apply. The added feature of rich touch sensing, designed specifically with autistic children in mind, could make their interactions even more meaningful.

### 3.4.1 Reflecting on our methods

We explored the existing literature to develop an initial set of six tactile-sensing requirements for a robot companion for children with autism (Table 3.1). Interviewing eleven autism specialists enabled us to verify the importance of these requirements and expand them into a set of seven richer and more complex descriptions than could be derived

from the literature alone (Table 3.3). The autism specialists were eager to work with us to design a robot specifically catered to the needs of autistic children. Participants actively encouraged us to take into account the tactile properties of the robot (new *feel* requirement), which will heavily influence the child’s interest in interacting with it, no matter whether they are touch seekers or touch avoiders. The participants alerted us if a requirement needed a shift in thinking, such as changing the approach for the *temporal* requirement to feature a customizable reaction delay. They also told us if they considered a requirement to be only a low priority. This open dialogue allowed us to form a ranked list of requirements that are focused on the therapeutic needs and capabilities of children with autism; however, as in any human-subject study, our chosen sampling method, interview format, and materials contributed to and influenced our results. Here, we reflect on the rationale for and the limitations caused by these choices.

Our participant pool consisted of eleven participants from the United States and Canada. As children across the autism spectrum behave differently and have different needs, one could argue that a larger pool of specialists could have been interviewed. However, as each participant adds about one hour of interview footage and several hours of transcription and data analysis, large-scale recruitment in a study of this format is not feasible. More importantly, we found we reached a saturation in input from these eleven specialists, most likely because most of them have interacted with dozens or hundreds of children with autism over their career. Therefore, we believe our results effectively reflect autism care in the United States and Canada, where similar therapy methods are practiced [69]. Future work could verify and revise our recommendations for autism care practices in other countries.

The remote format of the interviews allowed us to collect data from a variety of autism specialists who are geographically distributed, but it also prevented participants from physically interacting with the study materials. If the interviews had been conducted in person, we could have asked the specialists to demonstrate preferred gestures and locations directly on a robot or a sensor. Furthermore, the participants could have felt the presented robot and sensor prototype. However, given that recruiting specialists is challenging, we opted for the remote format to be able to recruit widely. In retrospect, we believe that maintaining some distance from the specific robot and sensor we showed may have enabled the participants to think more freely about the technology when answering our questions.

Seeing the example robots and tactile sensor seemed to help the participants better understand our research goals. On the other hand, showing only one robot model and one sensor design may have limited what the participants perceived to be their options and thereby impacted their responses. However, we did alter NAO’s appearance to present both a humanoid and animal form. Furthermore, when asked to describe an ideal appearance for the robot companion, participants gave a variety of unique answers beyond NAO’s traditional humanoid form, suggesting additional forms such as cartoon characters, stuffed toys, trucks, and other objects unique to the child’s specific interests. Participants also answered all of the questions regarding their recommendations for the robot’s

tactile perception before seeing the prototype tactile sensor.

Presenting the therapists with an initial set of touch-sensing requirements supported their thought process, but it did not seem to bias their responses. For example, the *adaptation* requirement was not perceived as important by the majority of the participants, and they thus suggested significant revisions to its qualitative definition. New requirements about the sensor's *robustness and maintainability* and *feel* were also suggested during many interviews.

We used the final list of seven qualitative requirements to propose corresponding quantitative specifications, but further studies would be needed to validate the proposed specifications. We hope that our methods, findings, and recommendations can guide the design of such studies in the future.

### 3.4.2 Implications for future research

The necessity of touch perception depends on the robot's role. Therapist responses in our study suggest that all four recommended roles – teacher, companion, emotion regulation, and communication – would benefit from tactile interaction. However, the benefits accrued from adding touch sensing may vary widely. For some roles, this capability may fall into the category of “nice to have” or “not necessary”, such as a robot teaching math skills. For others, such as emotion regulation and social companionship, touch sensing brings additional information that can help the robot better judge the child's needs and respond to them more appropriately (e.g., reciprocating a hug). Finally, for teaching acceptable touch interactions and facilitating nonverbal communication, touch sensing is a crucial requirement that should take the primary focus of child-robot interaction design.

Our touch-sensing guidelines span the hardware, software, and user-interaction components of a robotic system that must be tightly integrated. For example, the *sensing range* and *spatial* requirements reference hardware specifications, while other recommendations, such as having a customizable *temporal* delay for robot response, lean more on software and interaction-design approaches. Building a touch-perceiving robot for autism intervention poses different challenges for the sensing, computing, and HRI research communities. Here, we present the main challenges that we foresee in these areas, and we suggest future work to address them.

The biggest challenge for the sensing hardware is ensuring robust and reliable measurements in dynamic, uncontrolled environments, such as home or school settings. High-quality force sensors are typically rigid and fragile. In contrast, recently developed stretchable fabric-based sensors (e.g., [28, 77, 137, 104, 119]) can provide a promising solution, as they tend to be robust to high-force contacts and impacts and can cover non-planar surfaces. However, their soft materials often result in nonlinear sensing performance, which may make *gesture recognition* more difficult. HRI researchers should stress-test soft sensors with the touch gestures that are common among children on the autism spectrum to identify their potential failure points before field deployment. If using an existing commercial social robot, researchers should also test whether the robot

demonstrates suitable *robustness and maintainability* to withstand being touched by energetic users. Unfortunately, most existing commercial robots do not have enough tactile-sensing capabilities onboard to meet any of our quantitative specifications. However, one does not necessarily need to design and build a whole new robot. An external tactile-sensing system could be developed, mechanically fitted to the robot, and also intelligently integrated into the robot's existing processing technology.

Importantly, the tactile-perception guidelines we propose here are in reference to passive touch, meaning the robot is being touched by the child. Defining the requirements for active touch, where the robot touches the child, is beyond the scope of our current goals and would require a future study that considers robot mechanics, motion, and control alongside sensing capabilities.

On the software level, identifying touch gestures applied by children with autism is not trivial. The robot will need to process simultaneous tactile sensor inputs from regions all across its body. Some of these tactile sensors will be activated by the robot's own movement or pose, rather than the child's contact; developing strategies to screen out these self-caused tactile sensations is still a nascent research topic. Also, large individual differences among children with ASD (e.g., physical abilities, communication abilities, and intent) further add to the challenge of recognizing diverse gestures as they are applied in a realistic setting. Here, HRI research can build on the ongoing work in the artificial intelligence (AI) and personalization research communities. The solutions can range from designing user interfaces for manual calibration and customization of the robot by caregivers to developing algorithms that learn and adapt to a child's actions over time.

Further exploration is needed to collect quantitative tactile interaction information between robots and children with autism. We were unable to identify any literature that provided quantified measurements of physical contact between these interaction partners, so it seems that new physical experiments are needed. Touch locations and touch gesture data could be further characterized through an observation study with autistic children and a large online survey with their caregivers.

Conducting this study has convinced us that touch is a worthwhile sense to pursue when creating a robot companion for children with autism. Although such an endeavor poses major engineering challenges, we believe that a touch-perceiving robot that follows the qualitative requirements and quantitative specifications we have formulated will be able to engage in meaningful passive-touch interactions with children with ASD. In turn, such progress will hopefully increase the level of education and companionship that socially assistive robots are able to provide for this population and more broadly.



# Chapter 4

## Endowing a robot companion with practical social-touch perception

**Note:** This chapter is based on Burns, Lee, Seifi, Faulkner, and Kuchenbecker’s article “Endowing a NAO Robot with Practical Social-Touch Perception”, which was published in *Frontiers in Robotics and AI* [24] and is licensed under a Creative Commons license (CC BY 4.0). Some paragraphs in this dissertation’s Abstract, Introduction, Background, and Conclusion are also adapted from this publication.

Endowing socially assistive robots with tactile perception has long been a challenging task due to the complex and expensive nature of whole-body tactile sensing [42, 112, 132, 135]. Unlike the soft and robust skin of mammals, traditional tactile sensors have been rigid and fragile [42, 5]. Recently introduced fabric-based tactile sensor designs are promising for social touch due to their soft texture, simple design, and robustness [85, 43, 77].

Even though existing social robots have limited tactile sensing, one does not need to wait and purchase an entirely new robot with built-in touch sensors to obtain whole-body tactile sensing that is robust and informative. Rather, an existing robot can be externally fitted with a tactile-perception system. This proposed approach can enable researchers, medical staff, teachers, and caregivers to continue to use technology they have already purchased and learned about, while adding a whole new channel of possible interaction methodologies alongside existing audio and visual sensing methods. We have created a robust tactile-perception system that is easy to manufacture, is pleasant to touch, and can be applied to existing rigid-bodied robots, as seen in Figure 4.1. It consists of fabric-based resistive tactile sensors whose outputs are processed using a touch-gesture classification algorithm. However, our goals go beyond providing a practical touch-perception solution; we also wanted to inform the creation of a whole touch-perceptive robot system and investigate users’ reactions to such a robot. We therefore added this tactile-perception system to a NAO robot to create our social robot prototype, the Haptic Empathetic Robot Animal (HERA). HERA is intended to serve as a tool to help therapists teach children with autism about safe and appropriate touch [23]. While the NAO robot is frequently used in robot-assisted therapy studies to help children with

autism [128, 59, 124], we intend to introduce practical touch-perception to add a new set of teaching opportunities. We tested the performance of this tactile-perception system on HERA in a user study with healthy adults.

Based on our findings, we present the following contributions:

- We created a design framework for **fabric-based resistive tactile sensors** that can be mounted on the curved surfaces of existing hard-bodied robots, with easy-to-follow step-by-step building instructions and an open-source database of four pre-made patterns.
- We share a **labeled dataset** of sensor readings and classifier source code from a user study in which fifteen adults touched the robot’s arm at each of the four sensor locations using five affective touch-communication gestures (hitting, poking, squeezing, stroking, and tickling) at two force intensities (gentle, energetic). After training on sensor data from the study participants, our touch-perception system perceives combined gesture and force intensity with an accuracy of 74.1% on held-out test data.
- We provide analysis of **participant impressions** of the overall robot system, including their reported feedback, their suggestions on how the robot should respond to the ten tested social-touch gestures, and our observations of their touch interactions with the system.
- We outline the **physical performance characteristics** of our tactile sensor design, including a robust sensing range from 0 to 30 N, gathered from sensors mounted to the hand, lower arm, upper arm, and shoulder of a NAO robot.

## 4.1 Tactile Sensor Design and Fabrication

### 4.1.1 Sensor overview

Our fabric-based tactile sensor design consists of a layer of low-conductivity foam sandwiched between two outer layers of high-conductivity fabric, which serve as electrodes. The two conductive-fabric electrodes are connected to a microcontroller (Uno, Arduino, Italy) via metal clothing snaps and wires. The sensor is powered by enabling the internal pull-up resistor of an analog input pin on the Arduino. One electrode layer is connected to this analog input pin, and the other electrode is connected to ground. This simple circuit creates a voltage divider formed by the internal pull-up resistor and our sensor; we measure the voltage drop across the fabric sensor to calculate its instantaneous resistance.

An Arduino’s internal pull-up resistors have resistances between 20–50 k $\Omega$  [4]. We tested the pins of our microcontroller and found them to have internal pull-up resistance



Figure 4.1: Left: Custom tactile sensors were built from fabric and foam to wrap around the four rigid segments of NAO’s arm. Right: As seen in this photo-merge, the sensors are secured on top of the robot’s plastic exterior and are hidden underneath a soft koala suit to create the robot companion HERA. Photos redistributed from [24] under a Creative Commons License, CC BY.

values of approximately  $37 \text{ k}\Omega$ . The chosen high-conductivity fabric (Shieldex Combi-Tex, Statex, Germany) has a sheet resistance of  $<1 \text{ }\Omega/\text{sq.}$ , meaning that a square piece of material of any size would have an edge-to-edge resistance of  $1 \text{ }\Omega$ , which is negligible. The chosen low-conductivity foam (6.0 mm RS Pro Low Density ESD Foam, RS Components, United Kingdom) has a sheet resistance that depends on the force exerted upon it; when undeformed, it is approximately  $100 \text{ k}\Omega/\text{sq.}$ , and it can decrease down to  $0.1 \text{ k}\Omega/\text{sq.}$  when heavily compressed, giving the sensor a correspondingly wide sensing range. The electrical layers are placed on top of a base layer of nonconductive neoprene foam (1.5 mm, Neopren Solution GmbH, Germany) to disperse touch across a wider surface area and increase the softness of the robot form. Each layer is adhered to its adjacent layers at the edges using heat-activated tape (Thermal Bonding Film 583, 3M, United States).

Because of its uniform structure, applying the same contact at different places on this sensor produces approximately the same signal. Therefore, a single sensor equates to one tactile pixel (“taxel”) and does not provide location within the sensor or have multi-point touch discrimination. Identifying a general contact region of the robot’s body meets the guidelines for detecting social touch discussed in Section 2.2. Thus, this sensor design provides touch data that is both physically and computationally simple yet effective for social interaction. If an application requires more precise location information, such as splitting the upper arm into front and back halves, then multiple smaller sensors can be built and mounted next to one another.

### 4.1.2 Fabrication process for curved surfaces

While a fabric sensor manufactured on a flat surface has good electrical performance on a flat surface, its performance significantly degrades when it is bent around a curved surface. This degradation occurs because all the layers of a flat-manufactured sensor are cut to the same size. When the sensor is applied to a curved surface, the outer layers of the flat-manufactured sensor must stretch to reach around the larger perimeter. Additionally, the innermost layers may also bunch up. These stretching and compressing deformations cause the sensor to perform as though it is already being touched, reducing its sensitivity and sensing range. This practical challenge is important to address because most rigid robot body parts are composed of complex shapes rather than flat surfaces; curved tactile sensors are thus required. After trying various mitigation techniques, we discovered that one can best solve this problem by *creating a curved sensor from the beginning*, using a curved building surface that closely matches the curvature of the sensor's final intended location. We design the sensor's dimensions to increase with each layer so that the material does not need to stretch.

We provide a full step-by-step visual guide of how curved sensors are constructed in Figure 4.2. We also provide a video showing the entire fabrication process, which can be found in the supplementary material of our associated article [24] and is summarized as follows. In **step 1**, we identify the shape and curvature of the sensor's final location. We created the base layer of each sensor using tailoring: measuring the area by hand and drawing the resulting shape. One could also generate the base layer using methods such as three-dimensional (3D) scanning or molding. In **step 2**, we create templates for the sensor's subsequent layers, which can be determined either by sequentially placing and tracing each new layer or by using drafting software. In **step 3**, we assemble all of the necessary raw materials: high-conductivity fabric, low-conductivity foam, neoprene foam, heat-activated tape, rubber silicone elastic bands, and plastic clips. **Step 4** and **step 5** illustrate the phenomenon explained previously; for curved surfaces, the sensor layers must increase in size to compensate for the increased circumference that comes with increased layer height. In **step 6**, the materials are cut to the desired shape using scissors or a laser cutter. A tab-shaped protrusion is left on the end of each conductive outer layer. The tape is cut to 10 mm wide and lengths suitable for placement along all edges. **Step 7** shows the order of all material layers as a reference. Additionally, **step 8** indicates that the two tabs should be placed in different locations to prevent them from contacting, which would short the sensor. For **step 9**, we secure the sensor around a cylindrical form with a radius that closely matches that of the final sensor location on the robot. A plastic clasp is sewn onto each rubber silicone elastic band. The silicone bands are secured onto the Neoprene base layer using either sewing or heat-activated tape. The silicone coating increases the friction between the sensor and the outer surface of the robot, to keep the sensors stationary during interactions. In **step 10**, heat and pressure are applied to fuse the heat-activated tape, using methods such as a clothing iron or a mug press. For this study, approximately 149°C heat was applied using a clothing iron for

roughly 3 seconds in each area (as recommended by the tape’s bonding guidelines [1]), using a protective cloth covering and small circular movements to slowly work across the sensor. Each layer is added and heated incrementally to ensure the layers adhere properly. Finally, in **step 11**, the supply voltage and ground wires are connected to their respective electrode tabs via clothing snaps, and the sensor is ready to be mounted on the robot. The snaps also allow the sensor to be easily removed if needed, e.g., for repair or cleaning the setup. The final sensor is 10 mm thick.

### 4.1.3 Tailoring sensors across NAO’s arm

As a proof of concept for our system, we built four sensors to cover the left arm of a NAO robot. The arm has four rigid segments, as seen in Figure 4.1 and Step 1 of Figure 4.2: the hand, the lower arm (between the wrist and elbow), the upper arm (between the elbow and shoulder joints), and the shoulder. Each sensor was fabricated around a glass or metal cylinder of a curvature matching its corresponding arm segment. The cylindrical radii for the hand, lower arm, upper arm, and shoulder were measured at roughly 3.5 cm, 2.5 cm, 3.0 cm, and 4.0 cm, respectively. Rather than being a perfect cylinder, each arm segment has unique geometry, such as a bulge on the upper arm representing a bicep muscle, and a deep cavity in the hand to leave room for fingers. Therefore, each sensor was fabricated with a custom pattern for an optimal fit. We have publicly shared the patterns for creating these four sensors in an online repository [25].

### 4.1.4 Estimated costs and fabrication time

We purchased the high-conductivity fabric, low-conductivity foam, and neoprene foam in pieces that were 100 cm by 130 cm, 100 cm, and 55 cm wide, respectively. With these fabric quantities and the heat-activated tape, we estimate the price of the sensor’s raw materials to be 1.88€ / 10 cm<sup>2</sup>. As the surface area of the sensor increases, the percentage of surface area covered by tape will decrease, and so this price ratio will also decrease. The plastic clasps, silicone bands, and metal clothing snaps add a nominal cost. A sensor for the NAO’s upper arm, like that in Figure 4.2, with a surface area of 162 cm<sup>2</sup> and a perimeter of 55 cm, would cost approximately 7.78€. The Arduino Uno microcontroller we used retails at roughly 20€.

We estimate that a novice maker using our sensor patterns would need approximately 2 hours to create a fully operational sensor using scissors and a clothing iron. Using tools such as a laser cutter and a mug press should streamline this process to take closer to 45 minutes. Regardless of the tools used, the time needed per sensor would decrease if the maker prepares layers for multiple sensors in bulk, and as they gain experience with the techniques and materials. Beneficially, using the fabrication process and materials we propose, no waiting time is required between steps.




Figure 4.2: A visual guide highlighting the eleven key steps required for building our fabric-based tactile sensors. Illustration from [24], CC BY.

## 4.2 User Study Testing

After designing and constructing sensors for the NAO, we conducted a user study to learn about the performance of our tactile sensors as well as how users would receive a touch-sensitive robot system of the envisioned type. During initial testing of the sensor design, we observed that different touch gestures and force intensities could often be distinguished through visual inspection of the resistance data alone. We suspected that a machine-learning algorithm could automatically capture these patterns to perform accurate identification. Therefore, we hypothesized that the simple yet practical time-series resistance data produced using our tactile sensors would be sufficient to identify which sensor was being touched, the level of force intensity a user was exerting, and which gesture was being performed. Furthermore, our sensor design was motivated by the touch-perception guidelines detailed in Section 2.2. Using conductive foam not only gives our sensors a wide sensing range but also provides a soft and squishy feel. As the autism specialists recommended the robot's exterior to have such tactile qualities in order to promote user interaction, we also hypothesized that users would find our fabric-based sensors pleasant to touch and the act of touching a robot to be engaging and appealing.

To test these hypotheses, we introduced participants to our prototype robot companion, referred to as the Haptic Empathetic Robot Animal, or HERA [23]. We asked 15 adults to conduct five prominent social-touch communication gestures (as identified by the therapists in Section 3.2) at two force intensities on each of the four tactile sensors across HERA's arm in a full-factorial design. To evaluate the usability of our sensors, we wrote a machine-learning-based gesture-classification algorithm that we trained to process the sensor data over time. The raw sensor data from our user study is publicly available online, along with start and end indices for all gesture performances, a MATLAB script for generating the indices, and the source code for our gesture-classification algorithm [26]. To evaluate the personal experiences of the users, participants answered several questions about their interaction with the robot through opening and closing evaluations with sliding-scale ratings and open-ended questions.

### 4.2.1 Introducing HERA – robot form factor

As we were inspired by the positive impact of animals and AAI on children with autism [99, 97], we sought to create an animal-themed robot to use as our initial robot companion. The majority of robot animals available on the market are toys with little to no sensing or reprogramming functionality. Rather than design a new robot, we elected to use a NAO by SoftBank Robotics for our prototype. NAO is an ideal robotic agent for our purposes due to its small stature, processing power, and ease of use. Furthermore, NAO is often used in robotic therapy studies for children with ASD [124, 128, 59], giving us confidence in its acceptability for our target audience.

Robots modeled after non-familiar animals, such as PARO the baby seal, are more accepted as credible social agents than familiar-animal robots, such as a cat robot [116].

This preference is due to users' accumulated knowledge about familiar-animal behaviors and appearances through interactions with pets. We therefore decided to give HERA the outward appearance of a non-familiar animal by enclosing the NAO in a soft koala suit. This koala form factor also fit the appearance recommendations from the autism specialists.

## **4.2.2 Participants**

As our country was implementing lockdown restrictions due to the coronavirus (COVID-19) pandemic at the time of this study, we focused recruitment on our own research institute, including employees not in research positions and family and friends of employees. We distributed an institute-wide email advertisement.

We recruited 15 participants (7 female, 8 male) who were all adults (mean: 32, SD: 5), spoke English capably, and came from 10 different home countries, as reported in a demographic questionnaire. Four participants were from the United States, two from South Korea, and two from Italy, with the remaining individuals coming from China, France, India, Iran, Israel, Malaysia, and Switzerland. We also ran one pilot subject whose data are not included because the study setup and procedures were slightly modified after their session; participant numbers thus range from P2 to P16. Participants rated their familiarity with robots on a five-point scale. Two participants had no prior experience with robots, one had seen some commercial robots (novice), one had interacted with some commercial robots (beginner), seven had done some designing, building, and/or programming of robots (intermediate), and four had frequently designed, built, and/or programmed robots (expert).

## **4.2.3 Experimental setup**

To start, we secured the four custom-tailored tactile sensors across the NAO's left arm. Next, a koala suit outer layer was secured on top of the sensors. The purpose of the koala suit is four-fold: to create a friendly robot animal appearance in order to invite touch, to hide the added sensors, to further soften contact with the robot, and to serve as an additional electrical insulator to keep the user physically separated from the sensor circuit. We cut a hole in the suit behind the robot's head and its internal cooling fan to prevent overheating. The wires that attach to the electrode tabs of each sensor were hidden underneath each sensor, guided up the robot's arm, down the back, and out of the koala suit through a small opening. The wires connected to the Arduino, which was hidden behind a black tri-fold poster that stood on the table and behind the robot. A webcam was placed near the setup to record both the robot and the participants' hands as they interacted with the robot. This video provided a ground-truth record of the testing to which the experimenters could later return if there were any irregularities in the data. The robot was powered on, and its arm motors were engaged to hold the same joint configuration with maximum stiffness for each test.

Table 4.1: The seven questions in the robot acceptance survey asked in both the opening and closing evaluations.

---

I like the <b>presence</b> of the robot.
I feel <b>threatened</b> by the robot.
I am <b>afraid</b> of the robot.
I am afraid to <b>break</b> something while using the robot.
I think using the robot is a <b>good</b> idea.
I feel <b>safe</b> touching the robot.
It could be <b>useful</b> to touch the robot.

---

#### 4.2.4 Procedure

This user study was approved by the Ethics Council of the Max Planck Society under the Haptic Intelligence Department’s framework agreement as protocol number F008A. Participants not employed by our institution were compensated at a nominal hourly rate. Throughout the user study, the robot did not respond to any of the performed gestures; it simply held its constant pose. The four stages of the study proceeded as follows:

1. **Demographic questionnaire and opening evaluation (10 minutes)** – After providing informed consent, the participant filled out a demographic questionnaire that included topics such as country and familiarity with robots. The participant then completed an opening evaluation wherein they rated and commented on the tactile appeal of the robot’s two arms (with and without the added sensors) and also stated their level of agreement with seven general statements about the robot; these statements were adapted from a survey deployed by [55] and can be seen in Table 4.1.
2. **Instructions for touch interactions (5 minutes)** – The experimenter showed instructional slides explaining the upcoming touch interactions. The participant was introduced to the five gestures to perform in the study – “hitting”, “poking”, “stroking”, “squeezing”, and “tickling”. We also provided printed definitions for the touch gestures as defined by [136] in their touch dictionary. Next, the participant was informed which force intensity they should use on the robot first – “gentle” or “energetic” – and practiced at least one sample trial of their first force condition.
3. **Touch interactions on the robot (2 × 10 minutes)** – The experimenter switched to an automated presentation that indicated the force intensity, sensor location, and touch gesture to use in each trial. Figure 4.3 shows sensor data being recorded as a participant interacts with the sensor; the automated instructions can be seen in the background of the inset image. The participant was instructed to use the same force intensity for all of the touch interactions in this half of the task. Additionally, gestures were performed on one sensor location at a time. At each sensor

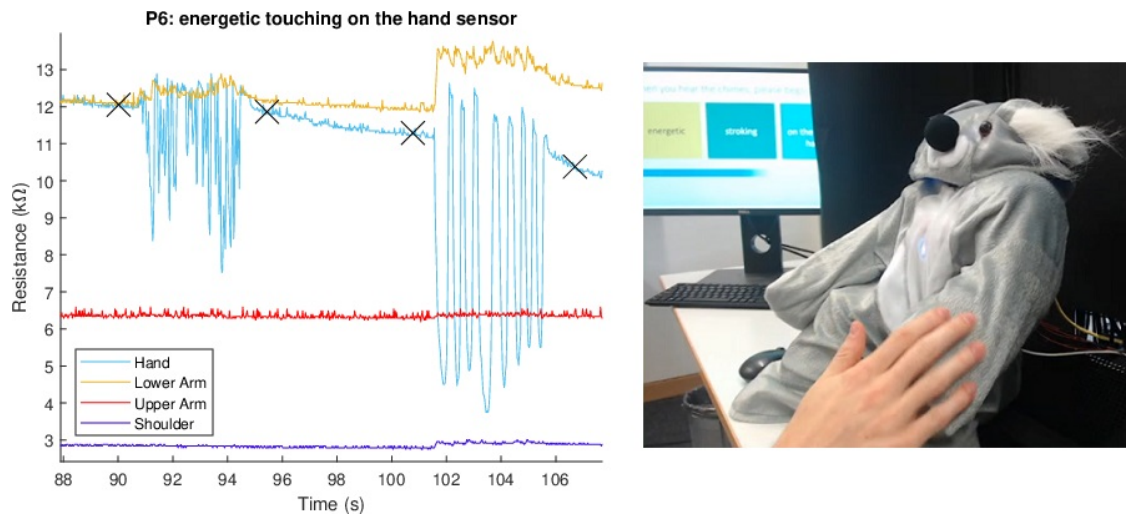


Figure 4.3: Real-time sensor readings show how the four sensor resistance values change (left) as the participant interacts with the robot (right). The four  $\times$  symbols on the plot mark the experimenter’s timestamps. In this example, the participant performed an energetic tickling gesture on the robot’s hand between 90 and 96 seconds (s) and performed energetic stroking on the same body part between 101 and 107 s. Figure from [24], CC BY.

location, the participant executed each of the five gestures in a randomly assigned order. Each gesture was performed for five seconds, with a five-second break of no touching between successive gestures. This approach is consistent with previous literature, such as [30]. Within the five seconds for touching, the users could do as many instances of the gesture as they wanted. This process was repeated twice for a total of three trials of each gesture. Once all of the gestures had been performed three times at the sensor location, the participant was instructed to move on to the next randomly ordered sensor location. This process was repeated until gesture data had been collected at all four sensor locations using the first force intensity. Finally, this entire step was repeated using the second force intensity. In total, each of the 15 users performed the 2 force intensities  $\times$  5 gestures  $\times$  4 sensors  $\times$  3 repetitions = 120 trials. A sample recording of the experimenter’s interface during this phase of the study can be seen in the second supplemental video of our associated article [24].

4. **Closing evaluation (5 minutes)** – The participant rated their level of agreement with the same seven statements from the opening evaluation and provided any final comments. They were additionally asked how they expected the robot would react to the gestures and force intensities from the touch interactions.

The resistance measurements from all four sensors were recorded at a sampling rate

of 40 Hz during all gesture performances. The experimenter held a manual button that was connected to the microcontroller. For every trial, the experimenter pressed this button just before the participant began a gesture and again immediately after the gesture's completion, which created timestamps marking the boundaries of each gesture performance.

### 4.2.5 Gesture classification

The resulting raw data from the four tactile sensors were then prepared for classification. As the randomized gesture order was generated in advance, the experimenter had a ground-truth order with which to pair the timestamps. We could then systematically label gesture data for each dataset; time spans with no gesture were labeled “none”. Every dataset was inspected in Tableau Software to ensure there were no missing or added markers, which would cause mislabeled data. Upon detailed inspection, we found that momentary electrical shorts to ground had occurred in the shoulder sensor's resistance data during some gesture performances. Most of these shorts occurred during participants' interactions with the neighboring upper arm sensor. We also discovered two trials where the hand and lower arm sensor temporarily shifted to identical resistance values, which we assume was caused by momentary electrical contact. We discarded all trials that contained one or more momentary shorts to avoid providing corrupted data to the classification algorithm; in total, 117 of the 1800 recorded gesture performances (15 participants  $\times$  120 trials per participant) were removed, leaving 1683 labeled intervals.

Next, we developed and refined the gesture classification algorithm. The classification was conducted using Python and the Scikit-learn library. We segmented the time-series data using a moving window size of  $w$  data points and an overlap size of  $0.5w$  between successive windows. If a data segment had an overlap of at least  $0.75w$  with any of the gesture time spans, the segment was labeled with that gesture. Otherwise, it was labeled as “none”. For each segment, we calculated a vector of 124 features (31 from each of the four sensors), including the entropy of each resistance signal, and the sum, max, min, average, median, standard deviation, variance, area under the curve, interquartile range, and number of peaks for the signal itself, as well as all ten of these metrics for the signal's first and second derivatives over time. These features were selected based on previous work on touch gesture classification, such as [57].

We trained and tested a random forest classifier on the raw data segments using a 70%-30% train-test scheme that was randomly sampled across all participant data. We selected a window size of  $w = 80$  samples (about 2 seconds), as it provides good test accuracy on gesture classification and a reasonably fast reaction time for future responses from the robot. This window size is also consistent with prior work on touch gesture classification, such as [30]. The parameters we used for the random forest model can be found in Table 4.2.

Table 4.2: The parameter values that were used for the gesture-classification algorithms. Unless otherwise specified, the same value was used for classifying contact location, intensity, gesture, and gesture and intensity together.

Parameter	Value
bootstrap	True
class_weight	None
criterion	“gini”
max_depth	Location: 40 Intensity: 20 Gesture: 40 Gesture and Intensity: 32
max_features	20
max_leaf_nodes	None
min_impurity_decrease	0.0
min_impurity_split	None
min_samples_leaf	1
min_samples_split	8
min_weight_fraction_leaf	0.0
n_estimators	Location: 50 Intensity: 100 Gesture: 100 Gesture and Intensity: 100
n_jobs	1
oob_score	False
random_state	None
verbose	0
warm_start	False

## 4.3 User Study Results

This section summarizes the results from our user study, including the gesture classification algorithm’s performance, the participants’ answers from both the opening and closing evaluations, and our observations of notable unexpected behaviors that users performed while interacting with the robot.

### 4.3.1 Classification results: location, force, and gesture

In all cases, the classification algorithm operated on the simultaneous data from all four sensor channels. To characterize the expected performance in unstructured interactions, we also evaluated the system’s ability to recognize the lack of contact, which was labeled “none”. Figure 4.4 shows the gesture classification algorithm’s performance as confusion matrices, classifying location, force intensity, and gesture individually. The location of each touch interaction was identified with 96% or greater accuracy for all sensor locations (chance level is  $1/5 = 20\%$ ). The system correctly identified force intensity with an

average accuracy of 89.3% (chance level is  $1/3 = 33.3\%$ ). In the case of identifying only gesture, without intensity, the classification system had an average accuracy of 74.3% (chance level is  $1/6 = 16.7\%$ ). We also looked at the gesture classification accuracy by sensor location: gestures on the hand, lower arm, upper arm, and shoulder had an average classification accuracy of 78%, 74%, 77%, and 68%, respectively.

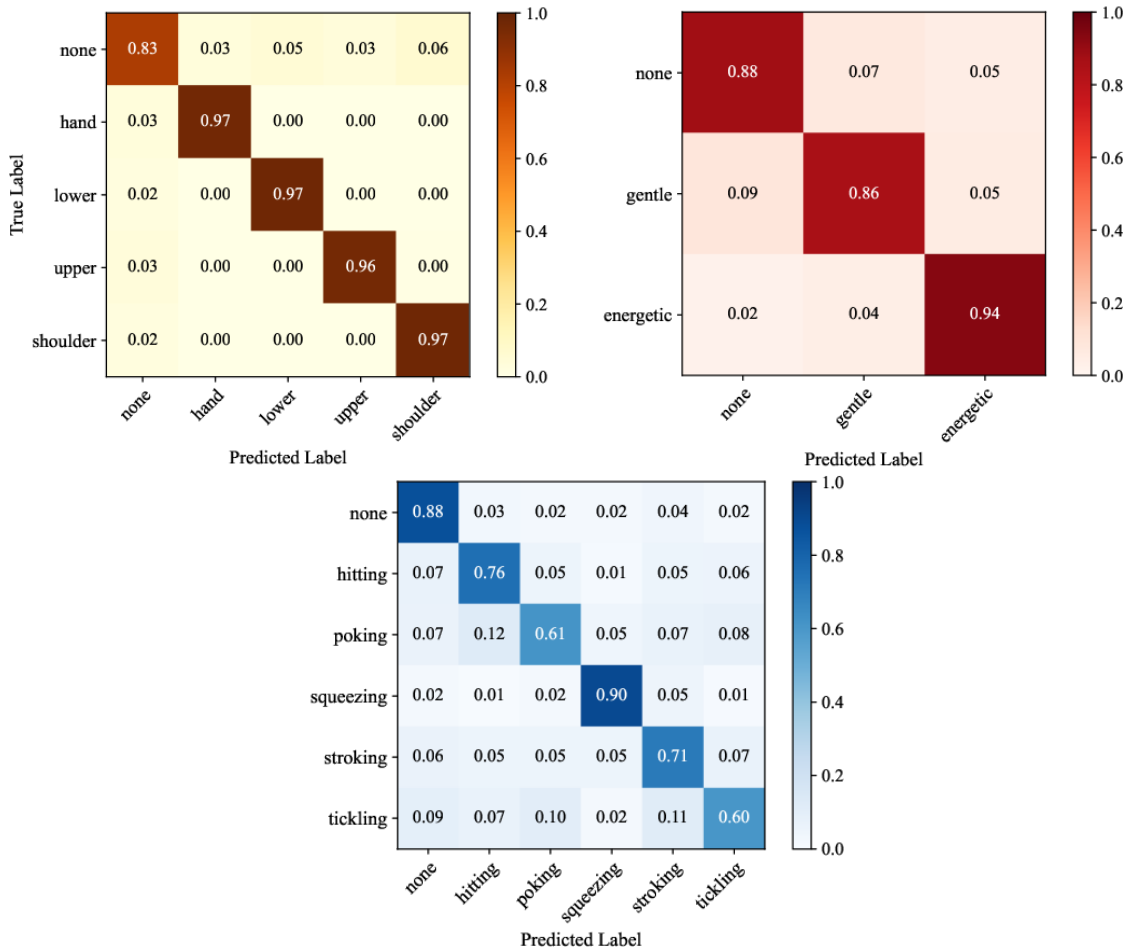


Figure 4.4: Normalized confusion matrices showing the classification results for location, force intensity, and gesture. Graphs from [24], CC BY.

### 4.3.2 Classifying force and gesture together

Next, we present the case where the gesture classification system predicted both gesture and force intensity together. The resulting confusion matrix is presented in Figure 4.5. The system had an average label accuracy of 74.1%, which is about 8.2 times higher than the chance of randomly guessing a gesture correctly ( $1/11 = 9\%$ ). Energetic squeezing

had the highest recognition rate at 88%. Gentle tickling was the hardest to identify, obtaining 64% recognition accuracy. Gentle tickling was most commonly confused with gentle poking (11% occurrence). Compared to results based on identifying gesture alone, identifying both gesture and force level improved the recognition accuracy for poking and tickling at both force levels. For the hitting and stroking gestures, the energetic force level had higher identification accuracy, and the gentle force level had lower identification accuracy. Finally, squeezing without a force level had a 90% recognition rate, whereas gentle squeezing and energetic squeezing had 80% and 88%, respectively.

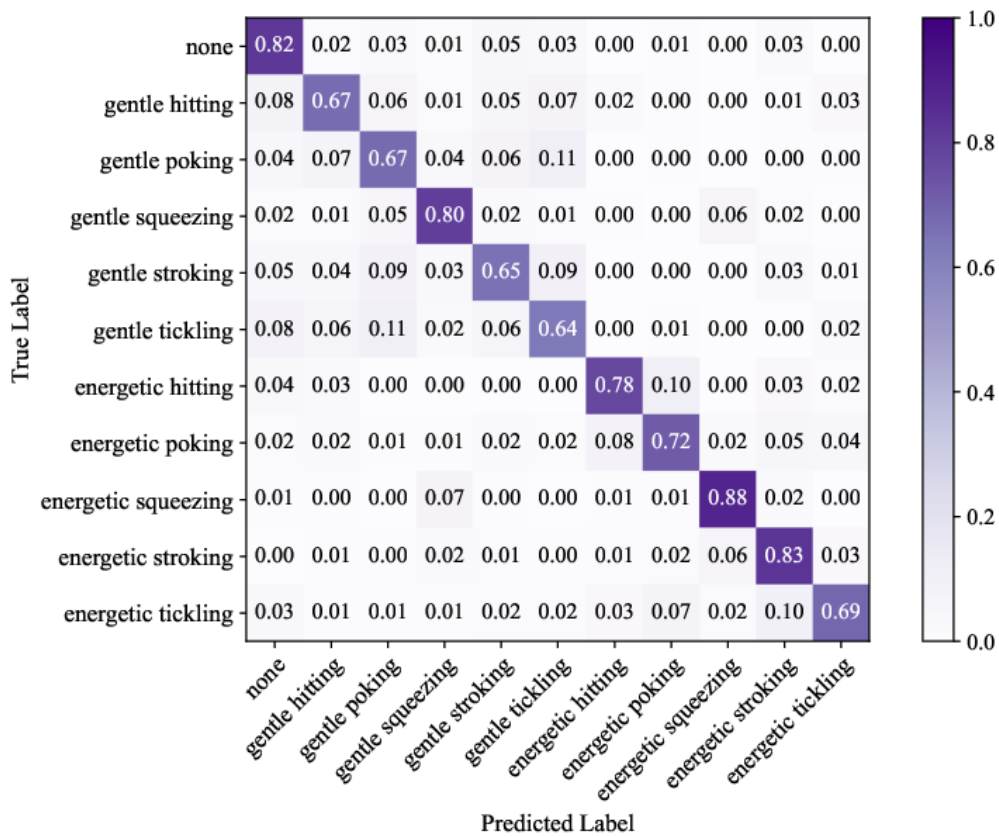


Figure 4.5: Normalized confusion matrices showing the classification results for both force intensity and gesture together. Figure from [24], CC BY.

### 4.3.3 Preference on the feel of the tactile sensors

Participants were asked to rate the feel of both the robot’s left and right arms on a scale from 0 to 10, with 0 being “Unpleasant to touch” and 10 being “Pleasant to touch”. While both arms were covered with the soft koala suit, the robot’s left arm also had the added fabric-based tactile sensors underneath. One participant ranked both arms as

equally pleasant, and all other participants ranked the arm with the added tactile sensors as more pleasant to touch. Anderson-Darling tests confirmed that participants' ratings for both arms were normally distributed. A paired-samples t-test revealed that there was a significant difference in the pleasantness scores for the feel of the arm with the added tactile sensors (mean: 8.0, SD: 1.3) and the feel of the arm without the added sensors (mean: 3.5, SD: 1.7,  $t(14) = -6.5$ ,  $p < 0.001$ ).

Participants were also asked to elaborate on their numerical ratings with an open-ended written response. Several participants referred to the sensor-covered arm with positive phrases such as “*soft and squishy*”, “*friendly*”, “*warm*”, and “*quite pleasant*”. Conversely, the arm without added sensors was referred to with phrases such as “*stiff*” and “*hard and unnatural*”. Participant P2 wrote, “*It would be even better if the whole robot was squishy.*”

#### 4.3.4 Perceptions before and after touch interaction

Participants rated their agreement with seven statements (Table 4.1) in the opening and closing evaluations on a scale from 0 (“Strongly disagree”) to 10 (“Strongly agree”). The results can be seen in Figure 4.6, labeled by the bold keyword in each statement.

Anderson-Darling tests revealed that the assumption of normality was violated for the results of four opening evaluation statements (threatened, afraid, break, and safe), and for three of the closing evaluation statements (threatened, afraid, and safe). Therefore, we evaluated each of the seven robot acceptance statements with a Wilcoxon signed ranks test. We found that only one prompt underwent a statistically significant change in rating. Participants agreed with the statement, “I like the presence of the robot,” significantly more after the touch interaction sessions (mean: 7.9, SD: 2.0) compared to before the touch interactions (mean: 6.9, SD: 2.1,  $n = 15$ ,  $Z = 2.937$ ,  $p = 0.003$ ). There were no other statistically significant changes in ratings among the robot acceptance questions. The non-significant results for the remaining comparisons are as follows: threatened ( $n = 15$ ,  $Z = 0.677$ ,  $p = 0.498$ ), afraid ( $n = 15$ ,  $Z = -0.405$ ,  $p = 0.686$ ), break ( $n = 15$ ,  $Z = -1.099$ ,  $p = 0.272$ ), good ( $n = 15$ ,  $Z = 1.482$ ,  $p = 0.138$ ), safe ( $n = 15$ ,  $Z = 1.265$ ,  $p = 0.206$ ), and useful ( $n = 15$ ,  $Z = 1.425$ ,  $p = 0.154$ ).

#### 4.3.5 Expected robot reactions

The participants rated how they expected the robot to react to the various touches, on a scale from 0 (a very negative reaction) to 10 (a very positive reaction). The resulting ratings can be seen in Figure 4.7 in the form of a box-and-whisker plot. Energetic poking, energetic hitting, and energetic squeezing all had median responses below 5, indicating that participants expected the robot to give a generally negative response for these three actions. All other gesture and force combinations had a median value higher than 5, signifying a generally positive response. However, it is interesting to note that some actions did not have a clear consensus between participant responses. For example, while

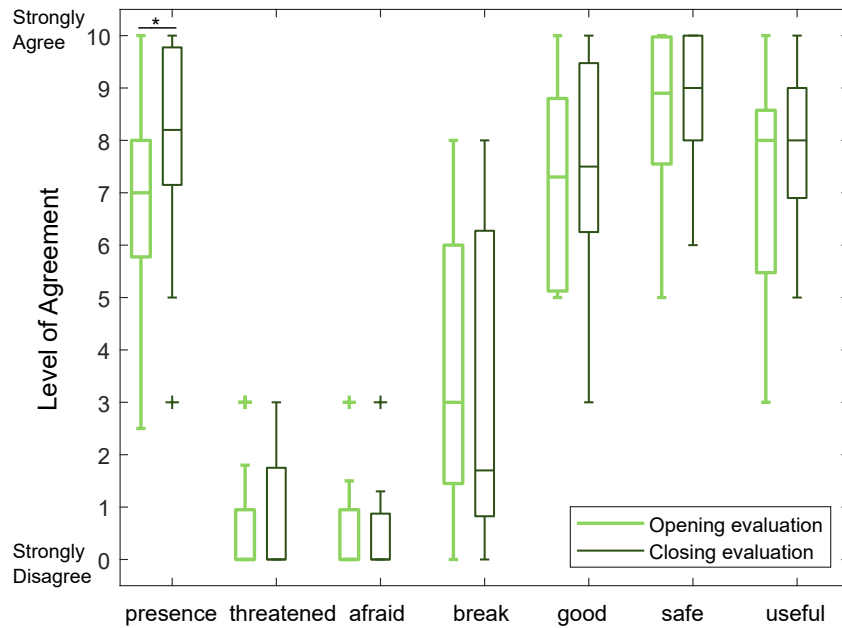


Figure 4.6: A box-and-whisker plot comparing participants’ level of agreement with seven survey statements before and after interacting with the robot. For each distribution, the central line indicates the median, the box shows the interquartile range (IQR), the whiskers show the range up to 1.5 times the IQR, and + marks indicate outliers. Plot from [24], CC BY.

energetic squeezing and energetic tickling received median ratings of 4.0 and 6.5, respectively, their relatively large interquartile ranges show that participants predicted both positive and negative responses. Some participants explained that they felt the response from the robot should be situation- or mood-dependent. For example, a deep squeeze could provide comfort and support, or it could be rough and hurtful.

### 4.3.6 Notable observations

During the course of our study, some participants exhibited two unexpected behavioral patterns. We share these observations to provide additional context for the gesture classification and survey results.

First, we observed that participants appeared to have differing perceptions of “gentle” and “energetic”. A gentle touch was typically reflected in the data as a narrow range of resistances, and an energetic contact tended to translate to a wide range. Figure 4.8 provides box-and-whisker plots grouped by sensor location to show the resistance ranges observed during each gesture. Although energetic performances typically did have higher ranges, some participants applied the same amount of force for both touch interaction

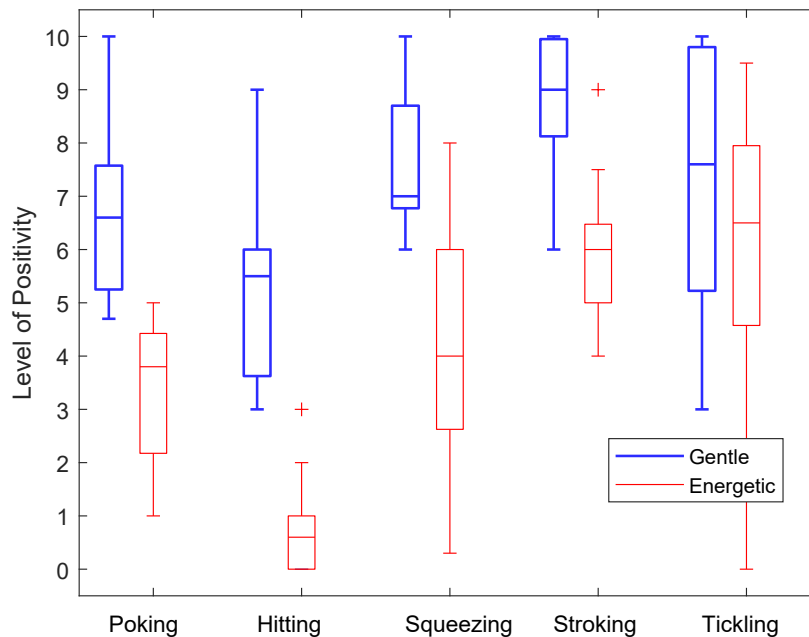


Figure 4.7: A box-and-whisker plot showcasing participants’ ratings of how they expect the robot would react in the future to the various touch gestures. A rating of 10 indicates a very positive response, and a rating of 0 indicates a very negative response. Touches with a gentle force intensity are shown in blue, and energetic gesture variations are colored red. Plot from [24], CC BY.

sessions (e.g., P3 and P9). Additionally, energetic gestures by some participants were in the range of gentle gestures performed by others, such as P7’s unusually strong “gentle” gestures compared to P2’s low-range “energetic” gestures. We believe that individual preference, inattentiveness, and fatigue may have all played roles in generating these varying performances. Some participants were also hesitant to touch the robot energetically; for example, participant P14 wrote in the closing evaluation, “*I felt uncomfortable with [performing] the energetic gestures as if it was a living creature.*”

Additionally, of all the gestures we instructed participants to perform, we observed that the hitting gesture evoked the most vocal responses. Several participants expressed reluctance or distress at being instructed to hit the robot. During the touch interactions, one participant (P4) said, “*I don’t like to hit the robot.*” Another participant, P5, asked, “*Do you want me to actually hit the robot?*” While performing a set of energetic touch gestures, P6 hit the robot with particularly high force and then quickly said, “*Oh, sorry!*” to the robot.

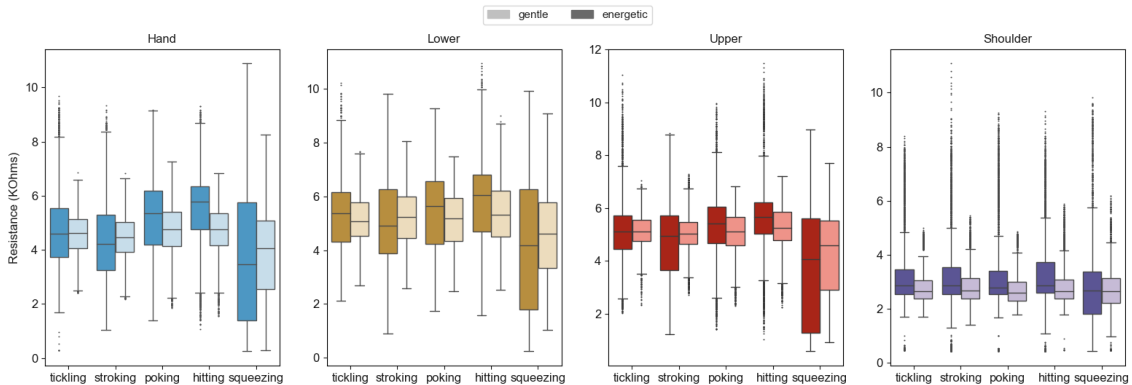


Figure 4.8: Box-and-whisker plots indicating the range of sensor resistance readings during all participants’ cumulative gesture performances. Gentle trials are shown in a pale color directly next to the more darkly colored energetic trials of the same type. Figure redistributed from [24], CC BY.

## 4.4 Physical Sensor Testing and Results

While conducting the user study, we noted that the shoulder sensor appeared to have a smaller sensing range than its counterparts. Due to all four sensors’ differences in surface area, shape, and curvature, we were curious to understand whether and how their sensing performance varied.

### 4.4.1 Physical experiment procedure

Our physical sensor experiments utilized the same setup as the user study experiment, except that rather than touching the robot directly, the experimenter used a compression load cell (FC-22, TE Connectivity, Switzerland) attached to a cylindrical indenter tip with a diameter of 13 mm (to mimic the size of a fingertip).

We wanted to evaluate the fabrication and performance consistency of our sensors by testing the sensitivity at several points across each sensor’s surface area. We also wanted to characterize the hysteresis of our sensors – the change in sensor resistance output values based on whether the external force on the sensor is increasing (loading force) or decreasing (unloading force). To accomplish both of these goals, an experimenter used the indenter tip of the load cell to press across the surface area of each tactile sensor with a force range of 0 to 30 N. To determine whether the koala suit affects the system’s performance, this experiment was performed both with and without the koala suit on top of the sensors. Data collection always started from no contact. The experimenter pressed down with the load cell until it reached 30 N and then withdrew the indentation again back down to no contact. Each touch indentation occurred at intervals spaced roughly 2 cm apart across each sensor. The experimenter simultaneously provided support to NAO’s arm at the next closest arm segment to prevent damage to the robot’s fully

engaged motors. The experimenter was careful not to touch the sensor currently being tested while supporting the robot’s arm.

#### 4.4.2 Physical experiment results

Performance markers were calculated for each of the four tactile sensors with and without the koala suit. First, it was found that the sensors have different initial resistances (i.e., the measured resistance when no external force is being applied) and different sensing ranges (i.e., the full range of resistance values a sensor may have depending on the force exerted upon it) due to their differing surface areas, manufacturing variations, and pre-compression by the koala suit. For example, with the koala suit on, the hand sensor had an average initial resistance of 11.5 k $\Omega$ , while the lower arm, upper arm, and shoulder had resistances of 18 k $\Omega$ , 9.5 k $\Omega$ , and 4.0 k $\Omega$ , respectively. Therefore, Figure 4.9 provides a normalized comparison of the measured resistance,  $R$ , by dividing each reading by that sensor’s initial resistance,  $R_0$ , in that condition. This visualization shows that the sensors on the hand, lower arm, and upper arm demonstrate very similar sensing behavior relative to their starting resistance, as does the shoulder sensor in the “without suit” condition. The shoulder sensor shows much less response when in the koala suit. Looking more closely, one sees that the lower arm and upper arm sensors also become somewhat less sensitive when covered by the suit.

Figure 4.9 is annotated to show each sensor’s initial sensitivity, final sensitivity, and hysteresis, with and without the koala suit present. The initial and final sensitivities for each sensor were calculated with and without the suit using the slope of the corresponding force-resistance loading curve; we calculated the average slope from the first and last five data samples, respectively. With the suit, the average normalized initial sensitivity of the hand, lower arm, and upper arm sensors is  $-0.181 \text{ N}^{-1}$ . The average final sensitivity was  $-0.004 \text{ N}^{-1}$ , confirming that all three of these sensors reach their saturation point at roughly 30 N. With the suit, the shoulder sensor was found to be nearly 10 times less sensitive than the other three sensors at initial contact and similarly sensitive at final contact. Returning to the other three sensors, the human sense of touch prioritizes good resolution at low force and has coarser resolution at high force. For example, on the arm, one experiences a pleasant mood from affective touches occurring around 0.6 N [127] and a maximum pain threshold at 25 N [88]. Aside from the shoulder sensor, which will be discussed below, our sensor design emulates this nonlinear response, making it suitable for detecting a broad range of social-touch gestures.

The hand, lower arm, and upper arm sensors all showed fairly consistent initial sensitivities across their surface areas. We observed that they all demonstrated a slightly diminished initial sensitivity around the perimeter, potentially due to the layers of non-conductive heat-activated tape that bind the layers at the edges. The least sensitive point on any edge of these three sensors had an initial sensitivity registered at  $-0.044 \text{ N}^{-1}$ , which is about four times worse than the average in the center and twice as good as the average in the center of the shoulder sensor.

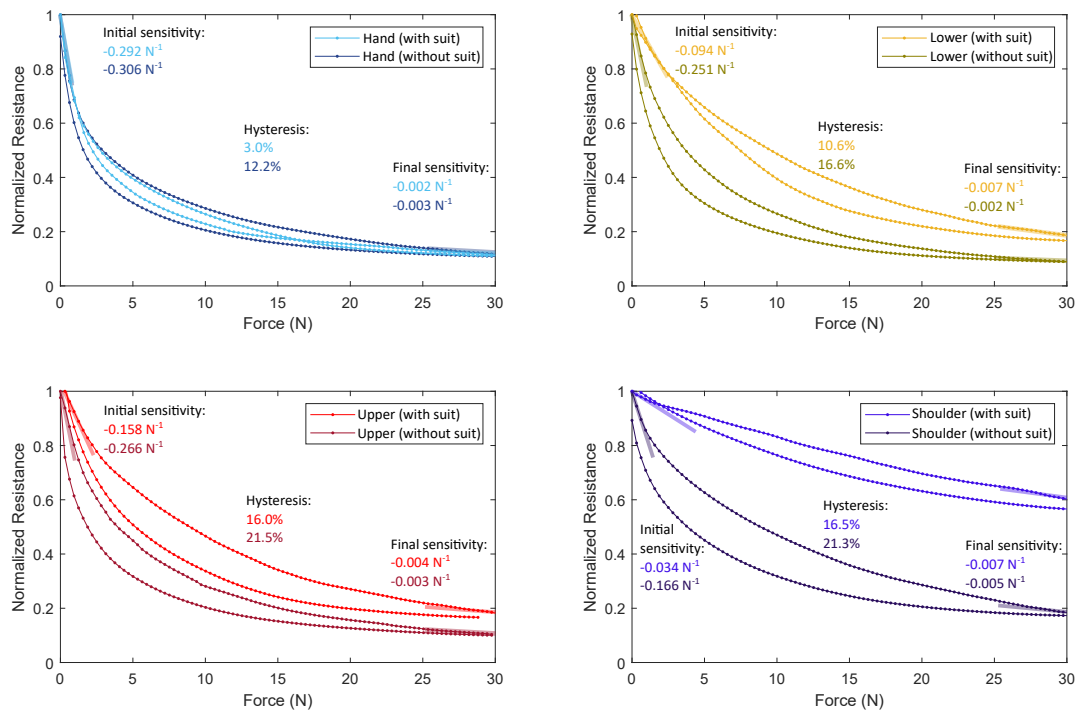


Figure 4.9: Normalized comparisons of individual sensor’s resistances across the tested force range with and without the koala suit, averaged across all trials of each physical experiment. The initial and final sensitivities are represented using semi-transparent lines oriented at the appropriate slopes. Each plot also lists the corresponding numerical values for sensitivity and hysteresis. Graphs from [24], CC BY.

The effect of hysteresis is present and consistent across indentations; it appears as distinct loading and unloading curves in Figure 4.9. We numerically quantified the average observed hysteresis for each of these eight conditions by calculating the ratio between the maximum vertical separation between the loading and unloading resistances across the tested force range divided by the total range of the minimum and maximum resistance readings. Without the suit the four sensors show between 12.2% and 21.5% hysteresis. When the suit is added, the hysteresis decreases somewhat.

## 4.5 Discussion

This chapter has presented a soft, low-cost, easy-to-manufacture tactile perception system that can be externally fitted onto existing hard-bodied robots for social-physical human-robot interaction. We sought to create a system that can work on curved surfaces and differentiate both the contact gesture and force intensity applied. Our system's hardware consists of a set of pre-curved fabric-and-foam-based tactile sensors attached to the left arm of a NAO robot. We analyzed the outputs of these four sensors by applying a gesture-classification algorithm to a dataset collected during a user study and also performing select physical experiments. Furthermore, we investigated users' reactions to interacting with a touch-perceiving robot through opening and closing surveys. This section reflects on the strengths and limitations of the scientific methods we utilized, and it discusses the contributions of our work toward future research in human-robot interaction.

### 4.5.1 The shoulder sensor

In the user study, the shoulder sensor had the lowest gesture classification accuracy of the four sensors (68% versus 78%, 74%, and 77%). The physical experiments showed that the shoulder sensor was nearly ten times less sensitive than the other three sensors. Why did this sensor perform so poorly? While we originally believed that the cause lay with our fabrication process, investigation revealed that the external koala suit was the core issue. We customized a child-size koala pajama suit that roughly fit the robot. However, unlike children, the NAO robot has large shoulder pads. Since we did not tailor the suit to provide additional room in the shoulders, the koala outfit pre-compressed the shoulder sensor. When we measured this sensor's response with the suit on, we obtained an initial resistance that was much lower than without the suit; the compression tests with the suit were thus measuring only the tail-end of the force-resistance curve, after the sensor had already been greatly compressed. When the suit was removed and the shoulder sensor was measured separately, it performed similarly to the other three sensors. Therefore, if the robot is fitted with an external suit, it is important to ensure that the suit fits consistently across the sensor system. Moving forward, we will modify

our koala suit so that it no longer exerts pressure onto the shoulder sensor in comparison to the other sensors.

## 4.5.2 Physical design implications

We discovered some aspects of our sensor design that should be addressed in future versions of similar tactile sensing systems. First, while the top and bottom faces of the sensors are covered with insulating material, the sides of the sensor are exposed. When the edges of two sensors touch, they electrically connect and interfere with each other's signals. We had to remove 6.5% of the recorded gesture trials from our dataset due to different sensors experiencing momentary electrical shorts. The electrically exposed sensor edges could be better protected by adding an insulating layer that surrounds the entire sensor, though care should be taken not to reduce the active sensing area or introduce hard edges that disrupt the feel of the system.

Secondly, we observed that the sensors display a peak of high resistance after contact and release, followed by a slow descent back to the baseline resistance, rather than an immediate return; an example of this behavior is visible in the hand sensor data in Figure 4.3. We suspect this slow return down to baseline resistance has mechanical origins. The sensor layers are secured together around the perimeter with heat-activated tape, leaving the middle of the sensor to have a small air gap between layers, which helps the sensors detect low-force touches. However, when the sensor is touched and released, the foam layer expands more slowly than the outer fabric layer. As the foam fills in to its original size, the air gap between these two layers is reduced, increasing the conductive surface area back to its original baseline. This phenomenon causes the sensor resistance to have somewhat different values depending on how recently it was touched. In the future, this variable air gap could be removed by using conductive adhesive across the whole surface area of each sensor, rather than just taping the edges. However, this solution may also lower the sensor's ability to detect low-force touches.

Incidentally, we also found that the baseline sensor resistance gradually increases over the course of months, presumably due to tarnishing of the conductive fabric and foam; normalizing by the sensor's initial resistance in a session removes this effect. Finally, even though we fixed the sensors in place with clasps and the silicone-grip bands, we found that sensor locations slightly shifted during the user study when the participant touched the robot particularly energetically. Moving to slightly different poses on the robot's body changes the baseline resistance of these sensors, as it changes their curvature; an example of this behavior is visible in the lower arm sensor data in Figure 4.3 at the start of the second gesture. We tried applying a high-pass filter at 0.5 Hz to remove the resulting baseline shifts from our user study dataset and also stabilize the slow return to baseline mentioned above. However, we ultimately omitted this processing step as we found it did not improve the classification accuracy, most likely because it also discards useful information about the intensity of contacts that occur. Instead, these changes can be reduced in the future by securing the sensors more tightly in place on the robot's

body parts, either by using higher-friction straps that can be tightened, or by securing the sensors with another mechanism such as double-sided tape.

### 4.5.3 Taxel size

As each of our sensors acts as a single taxel, we identified the general touch location based on where the sensor was fastened on the robot's body. It is possible that attaching several smaller sensors to each rigid body part could improve gesture recognition accuracy, since some gestures tend to move across the surface. However, increasing the number of taxels would also increase the hardware complexity and computational load of the touch-perception system. Low system complexity is an important design feature when one wants to add sensors to a large portion of the robot body. Furthermore, the borders between neighboring sensors might be particularly susceptible to electrical shorting and could diminish the system's pleasant feel. While we showcased the quality of gesture classification that is possible using four sensors that each cover a relatively large surface area, future researchers will need to consider the size and number of taxels best suited for their intended application.

### 4.5.4 User study limitations

The majority of our study participants had intermediate- or expert-level familiarity with robots (7 and 4 out of 15, respectively). While we had hoped to gather data from more novices, the COVID-19 pandemic limited our recruitment to participants with some connection to our research institute. Having participants who are less experienced with robots may have produced different results. However, recruiting internally enabled us to conduct this preliminary user study despite pandemic restrictions, which was essential for evaluating our tactile sensors and understanding participants' perceptions of the overall robot system. Furthermore, the users of future touch-perception systems for robots are most likely to already be familiar with robots. Additionally, there may have been some self-selection bias in our recruitment, with only those positively inclined toward touching robots answering our email. Nonetheless, our participants were diverse in many ways, such as home country, gender, and age, and they voiced a wide range of opinions about the presented robotic system.

We gave our participants limited instructions during the touch interaction sessions. By not providing specific numerical ranges for the force intensities to exert, we encountered a wide range of perceptions of what it meant to perform gestures gently or energetically. Furthermore, in the case of the "hitting" gesture, two participants actually conducted a "poking" gesture instead, including a participant who said they did not enjoy hitting the robot. Such behaviors make the classification problem nearly impossible and lead to some misclassifications. However, we believe that using these limited instructions enabled us to record participants contacting the robot very naturally, creating a more

general classification model that we believe will be useful beyond this user study during everyday interaction.

#### 4.5.5 Visual appearance of the robot

Our robot HERA has a cuddly koala exterior. Four participants referred to it as “*cute*” or “*adorable*”, and P14 said that it was like “*a living creature*”. It is possible that participants would have performed their touch gestures differently, or answered the evaluations differently, if the robot had a different appearance (e.g., a humanoid, an abstract design, or a scarier animal). Participants’ ratings on the three relevant questions of the robot acceptance survey showed they liked the robot’s presence, they found the robot to be non-threatening, and they were not afraid of it. These positive ratings could explain why participants expressed discomfort in hitting the robot. Future research could investigate how people touch and act toward touch-sensing robots with other appearances and form factors.

#### 4.5.6 Gesture-classification approach

We asked participants to perform each touch gesture for a fixed duration of five seconds per trial. Within these five seconds, the user could do the gesture as many times as they wanted. Like the variability in exerted force intensity, these variations in gesture frequency likely made the gesture classification task more difficult. During a natural interaction between a user and the robot, some gestures might be conducted over time frames that are different from our data collection window, such as a user poking or hitting only once. While windowing is a common method for classifying time-series data, another method would be to use thresholding on the signals themselves, or their derivatives, to promptly identify the start and end of contact.

Nonetheless, our preliminary results suggest that the proposed touch-perception system can accurately detect social touch including the contacted body part, force intensity, and gesture. With an average accuracy of 74.1% (about eight times better than chance), our classifier performs comparably to other social-touch setups with more complex hardware and software, such as Keshmiri *et al.* [70], even when identifying more force and gesture combinations. These results support our hypothesis that the rich time-series resistance data from our tactile sensors can capture both the gesture and force level conducted. Future researchers can further improve the presented data collection and analysis approaches for specific use cases.

#### 4.5.7 Implications for future research

People enjoyed touching the soft, sensor-covered robot. Their impression of the robot was also significantly improved after interacting with it. Furthermore, multiple participants actively voiced displeasure when asked to hit the robot, with some opting to touch

it in gentler ways. Altogether, these results lend credence to our hypothesis that users would find it appealing and engaging to interact with a touch-perceptive robot. It will be interesting to study how user impressions further evolve when a robot provides a response to their touches. Although we received initial recommendations from our participants, additional research will be needed to investigate the optimal ways for a robot to react to the various touch gestures it feels. Different responses may be more context appropriate to the input based on the robot's role. For example, robots serving as caretakers, teachers, or animal companions might all need to react differently to an "energetic poking" gesture.

For this study, we covered NAO's left arm with tactile sensors. As a next step, we want to extend sensor coverage across NAO's entire body. We are now starting to digitize the sensor fabrication process, including designing additional sensors in drafting software, utilizing a laser cutter to create the layers, and swapping out our clothing iron for a mug press, as shown in the supplementary fabrication video. These steps will streamline the sensor creation process and ensure even higher standards of reproducibility. We also would like to assist others in creating similar sensors for their own hard-bodied robot systems. Therefore, we plan to continue creating and sharing sensor patterns for the NAO robot in our database [25]. Our next goal is total system integration. We plan to move the gesture classification from offline to real-time operation by establishing direct communication between the tactile sensors and the algorithm. This improvement will let us test how well our classification approach generalizes to completely new sensor data, and it will enable the robot to immediately react to detected gestures using sounds, lights, or movements. When the robot has the ability to move its body, we will need to account for any mechanical stimulation of the sensors caused by the robot's own motion. We plan to train our system to predict self-touch artifacts generated by robot motions so that it can quickly and accurately identify external contacts that occur.

To date, there has been a glaring absence of social-touch perception in the current technology for human-robot interaction and socially assistive robotics. Conducting this study has proven to us that it is both possible and worthwhile to add tactile perception of both gesture and force intensities to social robots. Although the task of fully integrating an external system into an existing robot can be challenging, we believe that adding touch perception will enable robots to mimic the types of social touch interactions that occur so commonly between humans in everyday life, thereby providing users with more engaging and meaningful teaching, assistance, and companionship experiences.



# Chapter 5

## Investigating user preferences for robot emotions in response to touch

**Note:** This chapter is based on Burns, Ojo, and Kuchenbecker’s article, “Wear Your Heart on Your Sleeve: Users Prefer Robots with Emotional Reactions to Touch and Ambient Moods”, which was published in the *Proceedings of the IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)* [27]. Some paragraphs in the Abstract, Introduction, Background, and Conclusion are also adapted from this publication.

Robots may soon join our daily interactions as home assistants, companions, educational tutors [58], and therapy aids [100]. While meeting a new robot is often exciting, the novelty can wear off quickly [80]. It is important for such robots to maintain user interest over a sustained period of time to maximize the benefits of their use, such as completing an educational game series [111] or a physical-therapy regimen. One effective way to promote long-term HRI is to create adaptive robot behaviors that mirror aspects of human-human interactions [20]. In particular, robots can convey emotions during social interaction to increase their perceived naturalness (i.e., how similar the robot’s behaviors are to what the user expects), attentiveness (i.e., how much the robot detects its environment), and engagement (i.e., how the robot reacts to the detected input) [34].

In many user studies, the robot is controlled by a human operator to provide fast and appropriate emotional responses [66]. However, teleoperation is not a sustainable method of interaction for autonomous robots. In other research approaches, either the robot’s affective state (i.e., its simulation of emotion) is a fixed routine, regardless of user interaction, or the robot’s affective state is instantly changed by user action, usually to reward certain user behaviors, and then returns to a default setting [64]. These approaches neither demonstrate situational awareness from the robot nor adapt with the user, and therefore they may not promote long-term interaction. Furthermore, despite the importance of social touch in human communication, none of these approaches use social touch as a stimulus.

We investigated user preferences for shorter- and longer-term robot emotions through a video-watching study where 51 participants evaluated nine options for the richness



Figure 5.1: A screenshot from one of the nine videos in our study, which showcased three levels of immediate reaction to social touch and three long-term mood responses by HERA. Photo redistributed from [27], ©2023 IEEE.

of a robot’s response to social touch, as seen in Figure 5.1. These carefully crafted videos showcased robot behaviors ranging from no touch response at all to full emotional reactions and moods inspired by nonverbal cues used in human-human interaction. The nine depicted robot behaviors represented a full-factorial experimental design with two factors (reactions and moods) that each have three levels.

In terms of social touch, we found that users felt it was very important for the robot to immediately react to, or at least acknowledge, that a touch had occurred; users most preferred when the robot showed an immediate emotional reaction. Our qualitative analysis additionally revealed that if there was no immediate reaction, users still perceived an emotional mood as an appropriate, albeit delayed, response to touch. Furthermore, having the robot display a mood, either neutral or with emotion, improved users’ overall opinions.

In this chapter, we describe our user study in Section 5.1. The study’s quantitative and qualitative results are reported in Sections 5.2 and 5.3, respectively. We discuss the implications of our findings in Section 5.4.

## 5.1 User Study

Given the prevalence of robots that display only brief, static emotions, we wanted to investigate whether the combination of an immediate reaction and a visually perceivable, dynamic internal emotional state (i.e., mood) would be a feature that users would notice, appreciate, or even prefer, especially in regard to social-touch interaction. We evaluated this question through a video-watching study. In particular, we sought to understand

Table 5.1: A comparison of the nine robot behavior conditions shown in the study, each including a mood level and a reaction level.

	<i>Reaction Level: None •</i>	<i>Reaction Level: Neutral r</i>	<i>Reaction Level: Emotional R</i>
<i>Mood Level: None •</i>	Condition abbreviation: <b>••</b> No immediate reaction to touch. No ambient mood between touches.	Condition abbreviation: <b>•r</b> Looks at touched location. No ambient mood between touches.	Condition abbreviation: <b>•R</b> Emotional reaction to touch (positive or negative). No ambient mood between touches.
<i>Mood Level: Neutral m</i>	Condition abbreviation: <b>m•</b> No immediate reaction to touch. Neutral ambient mood between touches.	Condition abbreviation: <b>mr</b> Looks at touched location. Neutral ambient mood between touches.	Condition abbreviation: <b>mR</b> Emotional reaction to touch (positive or negative). Neutral ambient mood between touches.
<i>Mood Level: Emotional M</i>	Condition abbreviation: <b>M•</b> No immediate reaction to touch. Emotional ambient mood (positive or negative) based on prior touches.	Condition abbreviation: <b>Mr</b> Looks at touched location. Emotional ambient mood (positive or negative) based on prior touches.	Condition abbreviation: <b>MR</b> Emotional reaction to touch (positive or negative). Emotional ambient mood (positive or negative) based on prior touches.

the difference between emotionally neutral and emotional robot behaviors, as well as shorter-term and longer-term responses. Here, we list our hypotheses and explain the study.

### 5.1.1 Hypotheses

- H1. A robot that responds to physical contacts with **immediate positive or negative reactions** will be perceived as more intelligent than a robot that does not respond to touch. A robot that acknowledges physical contacts without emotion will be perceived as having moderate intelligence.
- H2. A robot that shows its internal emotional state through **positive and negative ambient moods** will be perceived as more intelligent than a robot with no ambient mood. A robot that displays a neutral ambient mood will be perceived as having moderate intelligence between these two extremes.
- H3. **Increasing levels** of both touch reactivity and mood will be increasingly exciting and engaging to participants.

### 5.1.2 Study design

In order to test these hypotheses, we designed a user study with two independent variables: robot reactions and robot mood. We define a reaction as the robot's immediate

response to a touch, and a mood as the internal emotional state of the robot, which changes over time based on touch stimuli. We created a  $3 \times 3$  full-factorial design with three **reaction** levels (abbreviated herein as none **•**, neutral **r**, and emotional **R**) and three **mood** levels (none **•**, neutral **m**, and emotional **M**). We created videos for all possible pairings of reaction and mood levels, with a total of nine conditions. Table 5.1 showcases the nine conditions compared in this study, including the abbreviations used to refer to each condition and a description of what generally occurred in each video.

Due to COVID-19 regulations, we conducted the study online and presented each participant with nine prerecorded videos. We utilized HERA as the sample robot for our videos, as zoomorphic robots are known to work well in therapy and care settings for users of all ages [110, 114, 62]. For this study, we used only the onboard tactile sensors built into the NAO robot’s hands and head, rather than HERA’s full tactile-perception system, as only a simple setup was needed.

We designed various physical behaviors for the robot to perform according to each reaction and mood level. For the neutral and emotional reaction levels, the robot reacted immediately after the contact occurred. For the neutral reaction, it turned its head toward the touch. For a negative emotional reaction, the robot pulled its arms inward and let out a cry. For a positive reaction, it cheered and waved its arms. The audio cues for the emotional reactions were used with permission from Javed et al. [64]. For the “none” reaction level, the robot neither moved nor made any sounds immediately following contact. For the neutral and emotional mood levels, the robot visually indicated its mood long after it had been touched and after the immediate reaction, if applicable. We modeled the mood movements on human body language and previous robot imitations of body language [41, 10]. For the neutral mood, it looked around the room. For the negative emotional mood, the robot lowered its head, drew its head and arms close to its chest, and moved slowly. For the positive emotional mood, the robot raised its head, demonstrated open body posture with open arms, and moved at a faster pace. For the “none” mood level, the robot neither moved nor made any sounds after the immediate reaction sequence was completed.

The nine videos were carefully created to differ only in the robot’s responses. Figure 5.2 provides a timeline to help visualize this timing and sequence of events. Each video begins with the robot waving its arm as a greeting at 5 seconds, but the reaction and mood performances that follow depend on the condition being shown. At 13 seconds, an experimenter hits the robot on its left arm, triggering its first reaction for the six videos belonging to reaction levels **r** and **R**. At 30 seconds, the robot displays its first ambient mood for the six videos belonging to mood levels **m** and **M**. At 43 seconds, the experimenter pets the robot on its head, triggering the second reaction. At 60 seconds, the robot displays its second ambient mood behavior. A simultaneous compilation of all nine video conditions can be seen in the supplementary video of our respective article [27].




Figure 5.2: Timeline illustrating the order of stimuli and responses that viewers saw in each of the nine video conditions. While the robot’s reaction and mood levels varied by condition, the timing and sequence of events remained the same across all nine videos. Illustration from [27], ©2023 IEEE.

### 5.1.3 Participants

Participants were adults who spoke and understood English well, had normal or corrected-to-normal vision, and had access to a device with internet and in-browser video and audio capabilities. As this survey was entirely online, we could recruit participants regardless of their location. We advertised our study via relevant email lists and on several social-media platforms. We also used snowball sampling by asking participants to share the study information with interested friends and colleagues. In total, 52 participants were recruited; one participant’s results were omitted for not following study instructions.

As we aimed to understand perceptions of our robot system across a general population, rather than a niche audience, we recruited a diverse set of participants across age, gender, home country, and experience with robots. Our final population (26 female, 24 male, 1 preferred not to answer) were all adults who ranged in age from 18 to 83 (mean: 37, SD: 17) and came from 18 different home countries (28 participants from the United States). Participants self-identified their familiarity with robots based on pre-defined levels in our demographic questionnaire: 14 had no prior experience, 13 were novices (had seen some commercial robots), 9 were beginners (had interacted with some commercial robots), 7 were intermediate (had designed, built, and/or programmed some robots), and 8 were experts (frequently design, build, and/or program robots).

### 5.1.4 Procedure

This study was approved by the Ethics Council of the Max Planck Society under the Haptic Intelligence Department’s framework agreement as protocol number F017A. Participants did not receive any compensation. After providing informed consent, they received a link to the study on an online survey platform.

After completing the demographic questionnaire, the user was asked to watch and listen to one of the nine videos corresponding to the conditions listed in Table 5.1 and described in Section 5.1.2. The videos were shown in random order to mitigate bias.

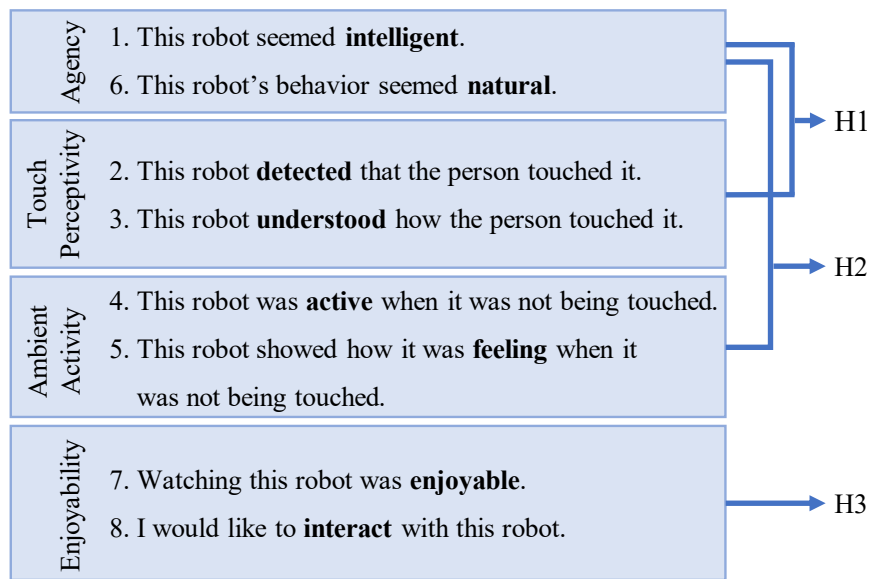


Figure 5.3: The statements that participants rated for each of the videos, numbered by presentation order. The boxes show how the eight statements were grouped into categories, and the arrows show the categories used to evaluate each hypothesis. Figure from [27], ©2023 IEEE.

After watching each video, the participant rated their agreement with each of the eight statements listed in Figure 5.3 using a sliding scale from 0 (strongly disagree) to 10 (strongly agree) with slider increments of 0.1. These statements were carefully designed to address our three hypotheses from different angles, as well as to evaluate whether there were differences between the neutral and emotional levels of reaction and mood. The connection between each statement and the hypothesis (or hypotheses) it addresses can also be seen in Figure 5.3. Participants then completed an open-ended response by writing what they liked and disliked about the robot behavior shown in the current video. This process was repeated until the participant had seen, rated, and commented on all nine videos. One final optional text box at the end of the survey gave participants the opportunity to provide any closing comments or suggestions.

The survey included approximately ten minutes of footage between all nine videos. As the survey was conducted online, there was no time limit required for completion, and participants were encouraged to take breaks from watching videos and answering questions as needed. The median survey completion time was 38 minutes (minimum: 20 minutes, mean: 50 minutes, SD: 36 minutes).

This procedure resulted in 72 sliding-scale answers and either nine or ten open-ended answers per participant. 51 completed surveys created a total dataset of 3,672 sliding-scale answers and 500 open-ended responses.

## 5.2 Quantitative Results

### 5.2.1 Thematically combining survey items

To reduce noise and increase interpretability, we combined pairs of related survey statements by averaging their responses for each user. The statements can be seen in Figure 5.3 and are labeled by the bold keyword in each statement. We combined Statements 1 and 6 (**intelligent** and **natural**) to create the “agency” category, which addresses both H1 and H2. Also for H1, we combined participants’ responses to Statements 2 and 3 (**detected** and **understood**) into the category of “touch perceptivity”. For H2, we combined the responses for Statements 4 and 5 (**active** and **feeling**) into the category of “ambient activity”. Finally, for H3, we combined the results of Statements 7 and 8 (**enjoyable** and **interact**) into “enjoyability”.

For all four categories, the statements merged had statistically significant correlations ( $p < 0.001$ ). The two statements which make up the agency category had a correlation coefficient of  $r = 0.78$ . The statements within the categories of ambient activity, enjoyability, and touch perceptivity had correlation coefficients of 0.59, 0.93, and 0.78, respectively.

### 5.2.2 Data processing

Next, we checked whether our data were normally distributed within each category for each video condition. We found that not all of the responses were normally distributed. For example, for the first video condition ●●, in which the robot gave no feedback at all, several participants rated the statements with a 0, a polarized answer at the bottom end of the provided rating scale. Therefore, we used the logit transformation to increase the normality of our data distributions:

$$\text{logit}(x) = \ln\left(\frac{x}{1-x}\right). \quad (5.1)$$

Here,  $x$  is a value between 0 and 1, and the function stretches out values at the two edges of this range. Since the transformation is undefined for 0 and 1, prior to merging the statements into their respective categories, any values rated as 0 were changed to 0.1, and ratings of 10 were changed to 9.9. The data were then divided by 10 before applying the logit transformation; the numerical range of our participants’ responses therefore shifted from 0 to 10 to a new range of  $-4.595$  ( $\text{logit}(0.01)$ ) to  $4.595$  ( $\text{logit}(0.99)$ ). Anderson-Darling tests confirmed that all four categories had sufficiently normal data distributions after transformation. Figure 5.4 uses box-and-whisker plots to showcase the combined, normalized responses of participants for the four statement categories separated by video condition.

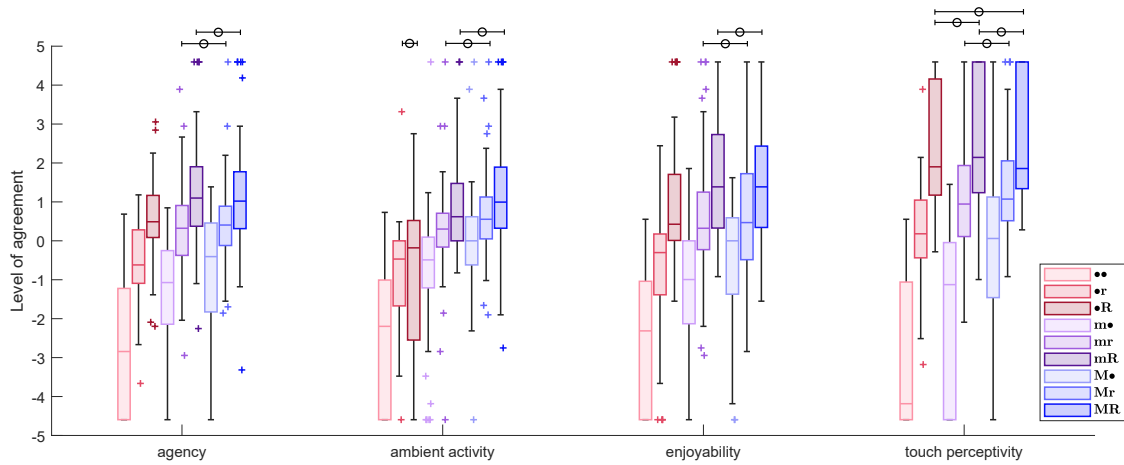


Figure 5.4: Box-and-whisker comparisons of combined, normalized user ratings of the video conditions. Lines above the boxplots marked with  $\ominus$  indicate pairwise comparisons that are NOT significantly different. For each distribution, the central line indicates the median, the box shows the IQR, the whiskers show the range up to 1.5 times the IQR, and + marks indicate outliers. The legend uses the abbreviations established in Table 5.1. Plot from [27], ©2023 IEEE.

### 5.2.3 RANOVA analysis

To prepare for repeated-measures analysis of variance (RANOVA), we then checked the four categories of data for sphericity. Mauchly’s Test of Sphericity indicated that the assumption of sphericity had not been violated for the agency category ( $\chi^2(9) = 10.892, p = 0.283$ ), the ambient activity category ( $\chi^2(9) = 12.875, p = 0.169$ ), or the enjoyability category ( $\chi^2(9) = 4.746, p = 0.856$ ). However, sphericity was violated for the touch perceptivity category; therefore, it was treated with a Greenhouse-Geisser correction ( $\hat{\epsilon} = 0.797$ ).

We then conducted four two-way RANOVAs. For all four categories, there was a significant main effect of mood and a significant main effect of reaction on the category’s rating, and there was also always a significant interaction between mood and reaction. We report the  $F$  statistics and  $p$  values for the main effects and interactions in Table 5.2.

For the agency and touch perceptivity categories, pairwise comparisons across the mood levels and across the reactions with Bonferroni correction ( $\alpha = 0.0083$ ) showed that there were significant differences between all three mood levels and all three reaction levels. For the ambient activity category, there was no significant difference between the **r** and **R** reaction levels. For the enjoyability category, there was no significant difference between the **m** and **M** mood levels.

Table 5.2: The results of two-way RANOVA analysis for each category.

	<b>mood</b>	<b>reaction</b>	<b>mood × reaction</b>
<b>agency</b>	$F(2, 100) = 59.06,$ $p < 0.001$	$F(2, 100) = 131.86,$ $p < 0.001$	$F(4, 200) = 6.71,$ $p < 0.001$
<b>ambient activity</b>	$F(2, 100) = 66.70,$ $p < 0.001$	$F(2, 100) = 53.77,$ $p < 0.001$	$F(4, 200) = 3.62,$ $p = 0.007$
<b>enjoyability</b>	$F(2, 100) = 52.80,$ $p < 0.001$	$F(2, 100) = 122.61,$ $p < 0.001$	$F(4, 200) = 6.35,$ $p < 0.001$
<b>touch perceptivity</b>	$F(1.79, 89.53) = 31.71,$ $p < 0.001$	$F(1.42, 71.04) = 130.64,$ $p < 0.001$	$F(3.19, 159.34) = 12.95,$ $p < 0.001$

### 5.2.4 Posthoc testing

After looking at the two independent variables separately with RANOVA testing, we conducted pairwise comparisons of all nine conditions listed in Table 5.1 using Tukey’s Test. All pairings had a significant difference unless otherwise noted. The results of this posthoc testing can be seen in Figure 5.4. Since the pairwise comparisons were predominantly statistically significant, we have chosen to highlight the pairs that were not significantly different.

For all four categories, there were two pairwise comparisons for which the difference was not significant: **mr** vs. **Mr** and **mR** vs. **MR**. This result indicates that so long as the visible reactions were the same, differing visible mood levels did not change how the user perceived the robot. Our category-specific findings are as follows:

*Agency:* No additional non-significant comparisons.

*Ambient activity:* The pair **•r** vs. **•R** also had no significant difference. If the robot did not perform a mood (i.e., no ambient activity) but reacted neutrally or emotionally to touch, the type of reaction did not affect the user’s impressions of the robot’s ambient activity.

*Enjoyability:* No additional non-significant pairs.

*Touch perceptivity:* There was also no significance between **•R** vs. **mR**, nor between **•R** vs. **MR**. Therefore, there was no significant difference in this category between any of the three conditions with the **R** reaction level. If the robot performed emotional reactions, then the mood level did not affect its touch perceptivity ratings.

## 5.3 Qualitative Results

Participants gave 500 qualitative short answers: 459 from the nine video conditions and 41 from the optional closing feedback text box. Their responses provided rich qualitative information that complemented our quantitative results, as summarized below.

### 5.3.1 Lack of reaction to touch is unsettling and unintelligent

If there was no immediate reaction, participants felt that the robot did not realize it had been touched, or they thought it could not discern whether the touch was positive or negative. Participants felt the robot was “*unaware of any differences in the types of touch*” (P41) and that “*it did not seem to understand the person’s action*” (P46). Furthermore, participants found the robot’s behavior unsettling if the robot moved ambiently but did not immediately react to touch. “*I did not like that the robot did not react at all to being slapped by the gloved hand,*” said P18. Participants called the lack of reaction “*unsettling*” (P1, P12) and “*unnatural*” (P3, P26, P35). “*It feels a lot more lifeless when it doesn’t respond*”, wrote P30.

### 5.3.2 Delayed reaction (i.e., mood) is better than no reaction

Participants really wanted the robot to react to touches and became upset if the robot did not. This phenomenon occurred in the conditions **●●** and **m●**. However, it occurred far less frequently in **M●**. Instead, the participants credited the emotional moods as very delayed responses to touch. Rather than call the robot unsettling, several participants said that the robot’s reactions (i.e., the emotional moods) were too slow. P10 said the robot was “*like a slow computer*”. P20 wrote, “*I liked the reactions to the hitting and the stroking. They seemed to convey actual emotion. I didn’t like that the reactions were delayed*”.

### 5.3.3 Neutral behaviors perceived as curiosity and awareness

People provided unique comments for conditions with a neutral mood and/or neutral reaction. Rather than the robot’s emotions, these comments were focused on the robot’s level of awareness about itself and its surroundings. Participants felt that the robot “*seem[ed] more aware about its environment*” (P11). They referred to the robot as “*curious*” (P9, 28, P37), “*naive*” (P7), and “*looking for*” or “*craving interaction*” (P35, P37). Some participants also took the neutral reaction as a questioning gesture, with the robot “*look[ing] for an explanation as to why it was touched*” (P3) or “*demanding for an answer*” (P28). Some participants even anthropomorphized what the robot might be thinking during the **mr** condition, writing “*Why did you do that? You woke me up*” (P14) and “*Why did you hit me?!? Let me think, was it something that I did?*” (P49).

### 5.3.4 Mood amplifies reaction and shows persisting memory

While participants strongly preferred when the robot showed a reaction, having a mood further improved their impressions of the robot. P10 echoed this sentiment, saying “*I liked that it reacted immediately, but I even liked it more when it started doing some*

*follow-up actions. It felt as if the robot thought about what just happened a bit more and became sadder or happy, just like humans.”*

Participants also appreciated that an emotional mood demonstrated that the touch interaction left a lasting impression on the robot. Even in the **M●** condition, with no immediate reaction, P1 wrote, “*After being hit, its overall mood changed to grumpy...I like that it seemed to remember something about the person, which made it more realistic*”. When there was no mood, participants felt that the robot’s behavior was less believable, with P10 calling it “*a forgetful robot*” and P42 writing, “*I don’t like that when being hit, after some seconds [the robot] seems pretty good with it*”.

### 5.3.5 Emotional reaction sounds were polarizing

The emotional reaction was the only reaction level with sound. Interestingly, the sounds evoked rather polarizing results. Some participants praised the sounds, calling them “*very endearing*” (P30) and writing comments such as, “*I liked the sound! Cool that sounds without words can still be so clearly understood*” (P38) and “*Wow! The sound adds so much!*” (P5). However, others wrote that they disliked the audio, with P31 calling it “*creepy*” and P33 writing, “*The voice was [a] bit brash and loud...I wasn’t sure if it liked to be touched or if it wanted me to leave it alone*”. Some participants also suggested that the sounds could be improved to better represent their emotional intents. Likewise, the emotional reaction movements could also be further tailored.

## 5.4 Discussion

### 5.4.1 Evaluating our hypotheses

As indicated in Figure 5.3, we evaluated H1 using both the agency and touch perceptivity categories and found that our predictions about immediate reactions matched well with the results. **Participants most prefer an immediate emotional reaction to touch.** Within each mood level, the robot was rated with the highest agency and touch perceptivity values when it displayed emotional reactions. Additionally, the emotional reactions received high touch perceptivity ratings regardless of the mood level displayed. The emotional reactions scored significantly higher than neutral reactions, which in turn scored significantly higher than no reaction. In our qualitative findings, participants perceived the robot as unsettling and unable to understand touch if it did not react, and they counted an emotional mood as a delayed emotional reaction.

Similarly, we evaluated H2 using the agency and ambient activity categories and found our predictions about ambient moods to be partially correct. If the robot showed no reaction to touch contacts, then the emotional mood scored higher than neutral mood, which scored higher than no mood. If the robot performed either neutral or emotional reactions to touch, then having some visible mood, either neutral or emotional, was rated signif-

icantly higher than conditions with no mood at all. However, there was no significant difference in ratings between the neutral and emotional mood. Qualitatively, participants viewed neutral behaviors as the robot being curious, and they noted that emotional moods enabled the robot to show a lasting memory about its touch interaction history. Therefore, we conclude that **demonstrating some mood, either neutral (perceived as curious or exploratory) or emotional, is better than no mood at all.**

We evaluated H3 using the enjoyability category to measure engagement and interest. Within each mood level, increasing reaction levels resulted in significantly higher enjoyability. With respect to increasing mood levels, if there was no reaction, then there was a significant difference in ratings between all mood levels. If there was a visible reaction (either neutral or emotional), then having a visible mood was rated significantly higher than having no mood, but there was no significant difference between the two visible mood options. **Participants most enjoyed the two conditions mR and MR.**

Overall, both the quantitative and qualitative results strongly support the importance of a robot responding to user touches on both shorter and longer time scales, using both an immediate emotional reaction and a visible mood.

## 5.4.2 Limitations

Though promising results were achieved, this study is not without limitations. First, some improvements could be made regarding the use of sounds in this study. The emotional reactions were accompanied by sounds, but the emotional moods were not. This difference may have affected participants' ratings. Future studies could feature some happy humming or sad mumbling for the emotional moods. However, the qualitative answers from participants show us that the emotional moods were clearly understood even without audio. Additionally, the sounds used for the emotional reaction vocalizations could be refined. If participants did not fully understand the emotional intent of the sounds, as P33 commented, then this ambiguity may have affected how they rated the survey statements. Future studies could incorporate a validation study to ensure the robot's sounds are perceived as intended.

Additionally, this study was conducted online rather than in person due to COVID-19 restrictions. Using videos of the robot prohibited direct physical interaction but provided us with a very controlled study. Each participant saw and rated the same emotional outputs in response to the same touch stimuli, and we did not have to worry about an in-person robot performing incorrectly during a trial. Participants were all presented with the eight survey statements in the same order, which could potentially have caused semantic bias. However, we chose this consistent ordering to make the survey as straightforward as possible for participants and to reduce the risk of respondent fatigue. Using an online approach also enabled us to recruit a wide range of diverse participants. While this diversity helped us to gain an understanding of user perspectives across many ages, countries, and robot expertise levels, not controlling these factors may have produced biases in our results. For example, previous research has found that users may react dif-

ferently to robots or have different interaction preferences according to their age [16]. Future work could investigate whether the findings we presented here still occur within narrower populations.

### 5.4.3 Implications for future research

In this study, the robot's reaction to touch did not depend on its current mood. However, this behavior could be made even more realistic. As P49 pointed out, "*After receiving the headstroke, the robot seemed happier, but the movements seemed a little exaggerated for someone who had just received a slap a few seconds before*". One could change the intensity of the robot's emotional reaction based on the current mood level. For example, a robot who had just been slapped could react less enthusiastically to stroking than a robot who had not recently received a negative touch.

Additionally, in this study, we used only three levels of ambient mood behaviors (negative, neutral, and positive) to present participants with distinct videos. In the emotional mood condition **M**, the robot quickly changed moods within one minute rather than over a slower, more realistic timescale, in order to show a mood change in a short video. To more closely imitate natural behavior, rather than representing mood using a set of discrete levels, we will utilize a more continuous approach by modeling the robot's internal affective state as a second-order dynamic system. This approach will enable us to fine-tune parameters controlling the robot's personality and to customize how its mood continuously changes over time.

Finally, while this study was a demonstration that utilized videos, we look forward to conducting future studies with in-person participants. Using HERA's tactile-perception system, we plan to conduct an in-person study in which HERA uses our aforementioned emotion model to display immediate reactions and longer-term moods in response to a variety of common social-touch gestures.

In summary, to investigate what impact robot reactions and moods play in user perceptions of robots receiving social touch, we conducted a video study with nine conditions. Indeed, the participants gave us a clear message: they prefer robots that respond to touch with immediate emotional reactions, and with visible moods between touches. In order to increase user interest and promote understanding of the robotic system, the robot should update its behavior based on the touch input it receives from its social interaction partner, the user. Having a robot react to touch in an emotional and customizable way, both with an immediate response and through a longer-term change in its mood, has the potential to promote a plethora of new experiences for human-robot interaction.



# Chapter 6

## Creating a robot that feels both touch and emotion

Using the insights gained from creating touch-perception guidelines, building custom tactile sensors, and investigating user preferences for robot emotions, we sought to combine all of our subsystems and create a fully autonomous version of HERA which can operate in real time. This chapter describes our efforts step-by-step. In Section 6.1, we describe our mathematical emotion model. In Section 6.2, we highlight the improvements made to our tactile sensors and microcontroller, and introduce the new touch dataset we collected. In Section 6.3, we introduce our improved gesture classification approach, evaluate and compare the results of two versions of the model, and describe the connection between the emotion model and the gesture classifier. Section 6.4 discusses the future research implications of our work.

### 6.1 Mathematical Emotion Model

Dynamic systems describe how natural phenomena evolve over time, such as the flow of water through pipes or the population of a community. Since such dynamics can easily be computed over time, they are a natural choice for a robot emotion model.

Widely used in psychology and HRI, Russell's Circumplex of Affect describes emotions across a continuous scale in two dimensions [109]. The first dimension, valence, refers to the positive or negative feeling of the emotion, and the second, arousal, refers to its energy level. We currently focus on modeling the valence dimension of a robot's emotional state over time; it can take any value between  $-100$  and  $100$ . For visualization, we draw the valence as the height of the robot's mouth. At maximum positive valence the robot has a perfect concave-up semicircle for a smile, and the inverse is true for maximum negative valence.

#### 6.1.1 Appropriate model order

Different natural phenomena can be best represented by dynamic systems of varying orders; we explore their suitability for HRI in Figure 6.1. Simple reactions, where the

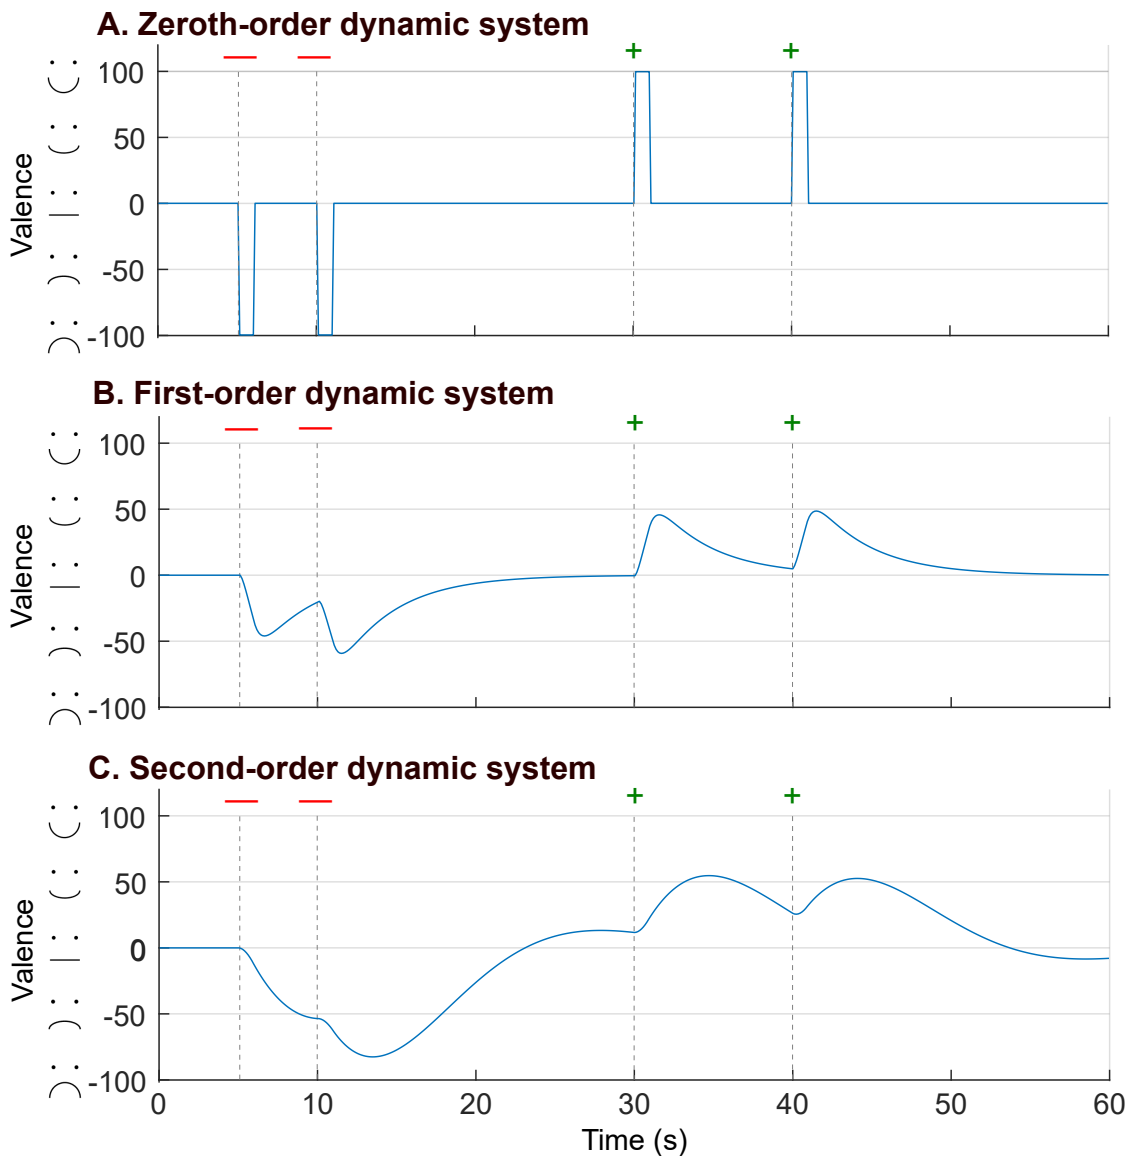


Figure 6.1: Modeling a robot’s internal affective state using dynamic systems of increasing orders. The – and + symbols indicate negative and positive stimuli, respectively.

robot’s affective state is a single response instantly called by the user’s action, can be represented as a zeroth-order system. The output (here, the robot’s valence) is proportional to the input (the external stimulus); a mechanical analogue is a linear spring with no mass or damper, which deflects as soon as a positive or negative force is applied. However, as noted in approach-avoidance theory, adding or removing a stimulus does not cause only a single, instantaneous, hard-coded response – rather, it also causes the organism’s internal state to be propelled toward a more-lasting mood. Therefore, a zeroth-order system




Figure 6.2: Our emotion model can be represented as a mass-spring-damper system. We illustrate its properties both at rest (left,  $y(t) = L$ ) and when past force pulses have moved the mass away from its neutral position (right,  $y(t) > L$ ).

is not the best way to represent natural emotions.

A first-order mechanical system has stiffness and damping, but it does not have mass. Therefore, it does not tend to oscillate. This order of system can represent richer emotions than a zeroth-order system because it maintains some memory of past experiences over time. As depicted in Figure 6.1, we need a second-order system to have the oscillation feature described in approach-avoidance theory.

### 6.1.2 Second-order dynamic system

The robot's valence is modeled as a linear second-order dynamic system with a point mass, a spring stiffness, a rest length for the spring, and damping (energy dissipation). The mass's position along the  $y$ -axis represents the robot's current affective state. External stimuli, such as visual user gestures, user dialogue, or touch contacts like petting or hitting, cause positive or negative force pulses that act on the mass and evoke an immediate emotional reaction from the robot. Furthermore, these pulses propel the mass along the  $y$ -axis and therefore also change the robot's mood over time. Importantly, these stimulus events influence the affective state toward a certain direction, but they do not teleport the robot instantly to a different affective state. This approach provides the robot with emotional memory, and its behavior is influenced by the entire recent history of interactions it has experienced.

A diagram of the mass-spring-damper system we use to represent the emotion model can be seen in Figure 6.2. Using Newton's second law, the sum of the forces currently acting on the mass,  $m$ , can be related to the mass's instantaneous acceleration,  $\ddot{y}$ , as:

$$F_s + F_d + F_p = m\ddot{y}, \quad (6.1)$$

where  $F_s$  represents the force applied by the spring,  $F_d$  is the force applied by the damper, and  $F_p$  is the fixed-duration force pulse currently being generated by external stimuli, if

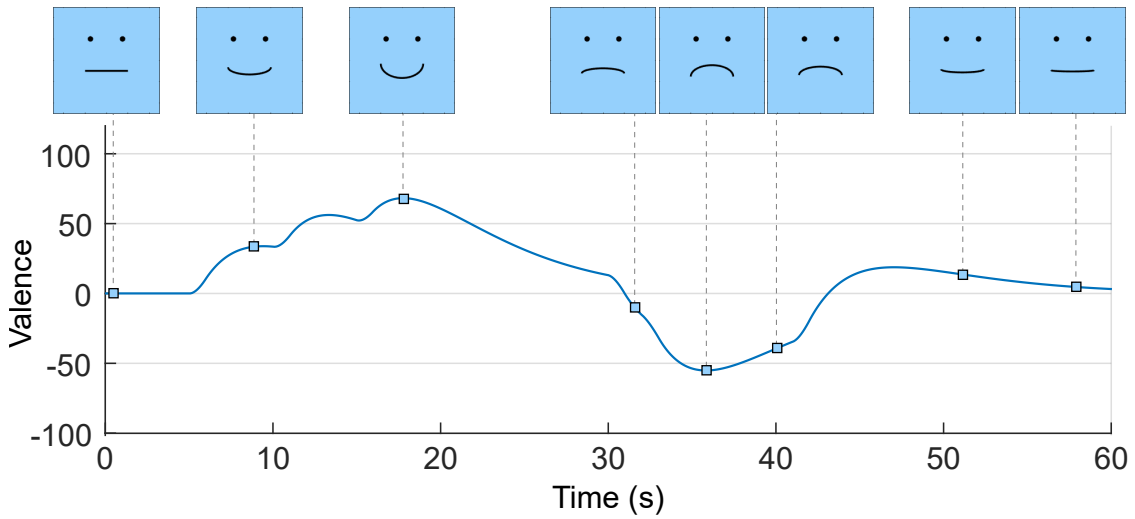


Figure 6.3: The height of the robot’s smile or frown directly corresponds to the internal valence. Blue squares indicate the valence that corresponds with the picture above it.

any exist. To show linear stiffness and damping, this equation can be expanded out as:

$$-k(y - L) - b\dot{y} + F_p = m\ddot{y}, \quad (6.2)$$

where  $k$  is the spring constant,  $y$  is the position of the mass,  $L$  is the mass position where the spring exerts no force,  $b$  is the damping coefficient, and  $\dot{y}$  is the mass’s velocity. Finally, this equation can be rewritten in the form of a second-order differential equation as:

$$\ddot{y} + \frac{b}{m}\dot{y} + \frac{k}{m}(y - L) = \frac{F_p}{m}. \quad (6.3)$$

In the absence of force pulses, and assuming  $b$ ,  $k$ , and  $m$  all have positive values, the mass will always return to the position  $L$ , pulled there by the spring, with oscillations calmed by the damping.

We provide a demonstration of our system in Figure 6.3, where one can see the internal emotion level of the robot over time. In this example, a small illustration of a robot face changes in real time simultaneously as the graph is plotted; we show the robot’s expression at eight timestamps. Though facial expression is convenient, a robot could use many other approaches to display its present valence level, such as ambient color, body posture, and sounds.

### 6.1.3 Customizing the robot’s personality

We can adjust selected parameters to customize how the robot’s emotional state changes in response to stimuli. First, we can calculate the natural frequency of this second-order

system as

$$\omega_n = \sqrt{\frac{k}{m}}. \quad (6.4)$$

We can then determine whether the system is over-, under-, or critically damped by calculating its damping ratio as

$$\zeta = \frac{b}{2m\omega_n}. \quad (6.5)$$

A robot with an overdamped emotion model ( $\zeta \gg 1$ ) may appear non-responsive, whereas one that is too underdamped ( $0 \leq \zeta \ll 1$ ) will experience large oscillations in its affective state and may appear erratic. As a small amount of oscillation is needed to mimic the bi-phasic response observed in approach-avoidance theory, we are most interested in damping ratios close to 1.

Depending on how the parameters of the emotion algorithm are tuned, the robot can display different personalities and responses. To display the versatility of this approach to modeling robot emotions, we present a fictional scenario in which a user interacts with a robot over the course of one minute. Figure 6.4 shows the robot's affective state in response to the described touch inputs. Each of the three subplots highlights how increasing and decreasing the value of a single parameter affects the robot's response. We provide a baseline parameter setting (the middle trial) that is identical across all three graphs. The user starts by stroking the robot three times (at 5 s, 10 s, and 15 s), which the robot perceives as positive touches. Then user then tries tickling the robot's feet (at 30 s and 32 s). The robot reacts negatively. Once the user realizes the robot didn't like this interaction, they attempt to console the robot by petting its head (at 41 s), to which the robot reacts positively.

We encourage the reader to carefully study Figure 6.4 to see how the robot's reactions in this story change as each parameter is adjusted. For example, the value of the spring rest length  $L$  determines the default valence to which the robot will always return. We associate this parameter with the **robot disposition**; setting a positive  $L$  creates a positive robot, whereas a robot with a negative  $L$  will have a negative valence when left alone. As we keep the mass constant, changing the natural frequency  $\omega_n$  adjusts the stiffness of the spring, which dictates how much the valence changes for the same input, as well as how quickly the system oscillates. We call this term **robot stoicism** – how strongly a robot resists reacting to external stimuli. A higher  $\omega_n$  leads to a more stoic robot, whereas a robot with a lower  $\omega_n$  reacts more and takes longer to return to its default  $L$ . Finally, the damping ratio  $\zeta$  calms oscillations in the affective state of the robot and by extension controls **robot calmness**. An underdamped system produces a robot whose mood shifts dramatically and who appears to overreact, whereas an overdamped system produces a robot who appears calm and measured.

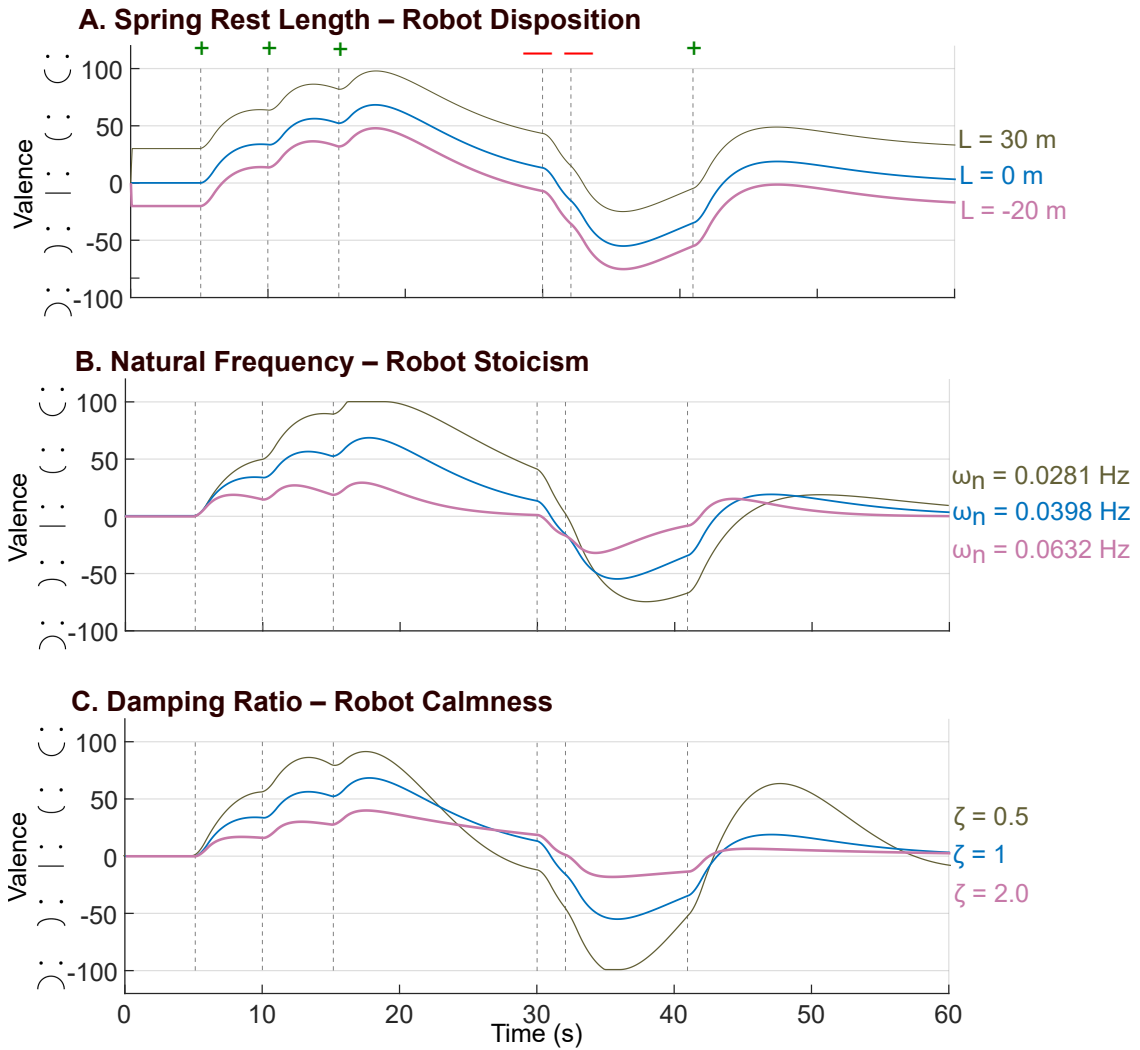


Figure 6.4: In each subplot, one parameter is changed to showcase how the robot’s response to stimuli can be customized. All graphs use  $m = 1\text{ kg}$ ,  $F_p = 30\text{ N}$ , and pulse duration  $T = 1.0\text{ s}$ . The middle blue trial is the same for all three subplots.

## 6.2 Upgrades to the Tactile-Perception System

Chapter 4 presented HERA’s tactile-perception system. In Section 4.5, we identified some aspects of the sensor’s physical design that could be changed to improve performance. Most notably, in our preliminary study, we needed to remove 6.5% of our recorded trials from analysis because neighboring sensor components came in contact with each other and experienced either momentary electrical shorting or coupling. We therefore took the knowledge we gained from our preliminary study to build an improved tactile-perception system. The following sections detail these improvements.

### 6.2.1 Sensor improvements

We built a new set of sensors following our study on the left arm of the NAO. For these new sensors, we introduced a finishing step of sealing the sensors’ perimeter edges with silicone. This thin silicone layer beneficially prevents neighboring sensors from electrically interfering with one another; it also reduces delamination and sensor motion. The soft feel of the silicone ensures that the sensors are still pleasant to touch. The sensors are also secured firmly onto the robot with double-sided tape to further reduce the risk of a sensor being repositioned after particularly vigorous touch by a user.

Additionally, we utilized new wiring for these improved sensors. The original sensors featured solid wires, which are inelastic and susceptible to breaking. The improved sensors use stranded wires, which are more robust and can withstand more movement and twisting. The wires were capped with two-pin female Molex connectors, which allow for easy and secure connection to the microcontroller. Additionally, we upgraded the metal clothing snaps that hold the wires in place. These buttons enable the sensor to be easily detached from the wires and the rest of the system if they need to be cleaned or repaired.

The original four sensor patterns were created by hand. We sought to develop a more systematic approach for the remainder of our sensors, which future researchers could implement for their own robots. To create the remaining sensor patterns, one researcher used aluminum foil to create a mold of each body part. The mold was traced, flattened, scanned, and then fine-tuned into a reproducible pattern using computer-aided design (CAD) software. From the initial base layer, and taking into account the thickness of each layer and the circumference of the body part, we then automatically generated the upper layers of the sensor, each of which were slightly larger than their predecessors. Using the resulting pattern CAD files, we were then able to use a laser cutter to produce the highly conductive fabric, low-conductivity foam, and heat-activated tape sensor layers. This streamlined approach enabled us to cut batches of layers in seconds and allowed for easy reproducibility. Cutting Neoprene foam in the laser cutter could release chlorine gas, and so instead these layers were cut by hand.

We iteratively designed the sensor patterns for each of the NAO’s body parts, and tested to see which laser cutter settings worked best. We now have sixteen sensor pat-



Figure 6.5: A view of all the sensors hidden under HERA’s koala suit. The inner NAO robot is covered with sixteen improved sensors, which provide HERA with whole-body tactile sensing. Moving joints were not covered to avoid pinching.

terns, which provide full-body tactile-sensing coverage to the NAO robot and can be seen in Figure 6.5. The sensors are divided into the following body parts, which include left and right versions for each: shoulder, upper arm, lower arm, hand, thigh, shin, and foot, along with the head and the chest. We opted not to have a tactile sensor on NAO’s back, as this is where the microcontroller is secured. We did not want to encourage users to touch near the microcontroller. The new sensors were also carefully designed not to lay on NAO’s joints, in order to avoid the sensor being pinched during robot movement. All sixteen sensor patterns are available via our open-access dataset [25].

Along with creating new sensors, we also improved the physical koala suit that surrounds them, with both hand- and machine-sewing. In Chapter 4, we discovered that the sensing range of the shoulder sensor had been reduced because the koala suit was too tight and was pre-compressing it. Therefore, we sought to better tailor HERA’s koala costume so that it hung evenly across the sensors, without causing any precompression. To accomplish this task, we repurposed fabric from a second, unused koala suit. Extra fabric was added through the back and bottom of the suit, which loosed the fabric in the shoulders to the appropriate tension. We also added a zipper in the back, so that the microcontroller was easily accessible. As the koala suit was originally purchased as a pair of children’s pajamas, the koala head was only a hood. We added additional fabric and snaps to the koala suit head, so that the NAO’s head could be fully enclosed.




Figure 6.6: The sender (right) is secured to HERA’s back under the koala suit, and the receiver (left) is connected to a computer.

### 6.2.2 Microcontroller improvements

We also made significant improvements to our microcontroller. Our previous microcontroller, an Arduino Uno, had only five analog input ports. As we aimed to cover the robot’s whole body with sensors, we needed more analog channels. We searched but did not find a commercial microcontroller that met our requirements: at least sixteen analog ports, wireless streaming, and a faster sampling rate than we currently had. Therefore, we upgraded from an Arduino Uno to a custom microcontroller board, built in-house by the MPI-IS Robotics Central Scientific Facility. The new M2 microcontroller [52] utilizes an ATmega32U4 processor and features sixteen analog input channels (one per sensor), a sampling rate of nearly 400 Hz (ten times faster to capture richer time-varying signals), and wireless data streaming via radio frequency (RF) transmission. This updated microcontroller is equipped with interface electronics, including a  $36\text{ k}\Omega \pm 10\%$  external resistor on each analog channel, which allows us to utilize the same voltage-divider approach as previously implemented with the Arduino. The updated microcontroller system consists of both a sender and a receiver module, as seen in Figure 6.6.

### 6.2.3 Collecting a new training dataset

In our user study, we captured a dataset of 1800 touches from 15 adult participants. However, the sensor data in this dataset was created using the old Arduino microcontroller, which had a sampling rate of 40 Hz, while our improved microcontroller had a sampling rate of nearly 400 Hz. It was possible that there was additional high-frequency information in participants’ touch signals that were not captured with the slower sampling rate. Since it was unknown what the touch signals looked like when read with a faster sampling rate, we could not extrapolate the signals of the old dataset to expand from a 40 Hz signal to a 400 Hz signal. Furthermore, there were differences in construction (i.e., the silicone lining) between the original sensors and the improved ones. Additionally, we now had sensors all over the body, rather than only on the arm. Touches conducted on the foot or chest might look different from a touch on the lower arm. For these reasons, we opted to record a new touch dataset using our improved whole-body set of sensors

and microcontroller.

We sought to automate and streamline the data collection process for this new dataset. For the original dataset, each participant was given touch instructions from a PowerPoint presentation, the ordering for which had been randomly generated and then individually assembled by an experimenter. Furthermore, the beginning and ending of touches were manually indicated by an experimenter during the data collection using a button press. We needed a way to automatically track which touch combinations had already been collected and to automatically label the beginning and ending of the gesture data.

To solve this challenge, we created a graphical user interface (GUI) in MATLAB, which can be seen in Figure 6.7. Three experimenters provided the touches for the new dataset, in order to both control the examples in the training data, and to provide a wide variety of examples. Before participating, the experimenters carefully studied videos of the touch performances from the original dataset so that they could perform a wide variety of touch styles and not just their own. They also met to discuss and define a general understanding of how they would perform the gentle and energetic force intensities. While different participants in the preliminary study all touched with different force intensities, we sought to train the robot with a single definition – the robot’s own “personal understanding” – of gentle and energetic touches.

The GUI provided a real-time view of the sensor signals as the experimenter touched, so that the experimenter could quickly verify whether they were satisfied with the training example they had produced. Instructions for the next type of gesture, touch intensity, and location on HERA’s body were given at the top of the screen. Two seconds after clicking the “Start trial” button (to give the experimenter time to move their hand from the mouse to the upcoming location), the red light would turn green, and the experimenter would conduct the touch for five seconds, after which the light returned to red, indicating the trial had finished. The sensor signals from all sixteen channels were recorded during this five-second period, and the active channel and type of touch were labeled on the resulting data file. If the experimenter saw that something had occurred to make the trial unusable, such as touching the wrong body part or a signal error, they could clear and repeat the trial.

The GUI was set to rotate through all possible trial combinations in a random order. We utilized the same five social touches as in the preliminary study – hitting, poking, stroking, squeezing, and tickling. We also used the same two force intensities, gentle and energetic. For locations, we used all sixteen locations across HERA’s body, along with a seventeenth location we called “nowhere”. When the experimenter received the instruction to perform a touch “nowhere”, they simply did not touch the robot for that five-second trial. This seventeenth location allowed us to periodically record what values the sensors showed when the system was completely undisturbed. Each experimenter went through the data collection process twice, for a total of  $3 \text{ experimenters} \times 2 \text{ full runs} \times 5 \text{ touch gestures} \times 2 \text{ force intensities} \times 17 \text{ touch locations} = 1020 \text{ unique trials}$ .

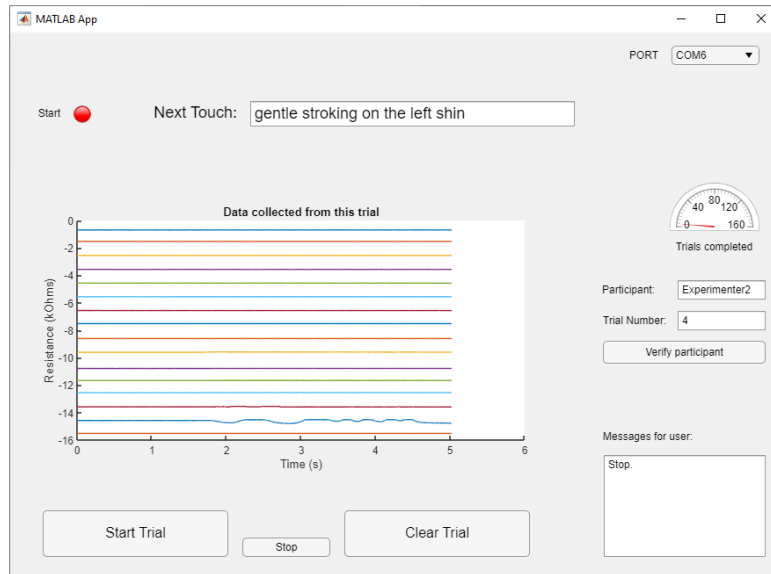


Figure 6.7: The GUI that the experimenters used to build the new training dataset. Experimenters could view upcoming touch instructions, watch plots of the sensor data in real time if necessary, and keep track of how many trials had been completed.

## 6.3 Upgrades to the Gesture Classification Algorithm

### 6.3.1 New gesture classifier approach

As we had collected a new dataset with an upgraded tactile-perception system, it was an opportune time to revisit our gesture classification approach. Our original classification algorithm was based on a random forest, and had been trained and tested using data from our four original arm sensors. However, it was difficult to improve the accuracy beyond what we reached in Chapter 4 without venturing into the territory of overfitting. Additional fine-tuning of the hyperparameters resulted in only marginal model improvement. With a new full-body sensor dataset, we would likely need to reoptimize model hyperparameters. To circumvent this process, we used a gradient-boosting classifier [12]. While random forests and gradient boosting are similar in that they are both tree-based ensemble algorithms, gradient boosting has the added benefit of onboard regularization, which decreases the necessity of hyperparameter tuning.

As with our previous gesture-identification model, the classification was conducted using Python and the Scikit-learn library. The data was segmented using a moving window size of  $w$  data points (where  $w = 385$  samples, about one second) and an overlap size of  $0.9w$  between successive windows. For each data segment, we evaluated a variety of features including the sum, max, min, average, median, standard deviation, variance, area under the curve, interquartile range, and number of peaks, as well as all ten of these metrics for the signal's first and second derivatives over time.

Various versions of the gradient-boosting model were trained using the new dataset, each with an 80-20% train-test split that was randomly sampled across all data. We trained the classifier on examples from all sixteen sensors because we wanted a general model that considers data from only a single sensor at a time. When the classifier looks at all sensors together, as done for the four sensors on the arm, it must choose at every time step which body part is being contacted, if any; touch can simultaneously occur at multiple body locations and should be perceived independently. We then used sixteen parallel instances of the trained classifier (one per sensor) to enable HERA to perceive touches of different kinds across its entire body in real time.

Another area of improvement was that our previous model was evaluated offline. We upgraded our classifier to provide real-time classification. The real-time approach utilizes sliding windows of data in a first-in, first-out queue, can run at 40 Hz, and processes all sixteen sensor channels in parallel, providing a separate output for every sensor at every time step. Next, while we could already see the offline test- and train- accuracy of any model we trained, we also needed to see how well a model would perform in real time. We therefore set out to create a test setup for analyzing the real-time accuracy of our new classification model.

It was important to create a test setup in which the user touches were as natural as possible, while still being very controlled. Touches were either four, six, or eight seconds long, in order to represent diverse durations of time that users might interact with the robot. Additionally, we wanted to avoid any sort of machine-learning bias that might have occurred if we had tested with five-second touches, which was the duration of each sample in the training data. There was a ten-second wait instructed in-between each touch. Two experimenters touched HERA at the same time. Due to the staggered touch durations, there were four possible scenarios: no one touching the robot, only experimenter 1 touching the robot, only experimenter 2 touching the robot, or both experimenters simultaneously touching the robot. This design allowed us to observe the classifier's accuracy when dealing with multiple touches occurring on the robot's body across different locations. The experimental setup for this real-time testing can be seen in Figure 6.8. The setup consisted of two trials. For each trial, each experimenter had a different set of ten touch instructions. These two trials were each performed twice – once with the balanced model running, and once with the unbalanced model running.

### **6.3.2 Mechanical coupling vs. electrical crosstalk**

Upon initial real-time testing, we noticed that more sensors seemed to be activated during touch than we were expecting. We investigated whether these additional touch detections were due to mechanical coupling or electrical crosstalk. In the event of electrical crosstalk, signal information could “bleed” over from one analog input to the next channel that is read. This would appear in the data as a cascading effect, with subsequent sensor channels displaying a sudden spike in values. However, we did not observe this phenomenon. Rather, sensors that were not next to one other electrically (i.e., not neigh-

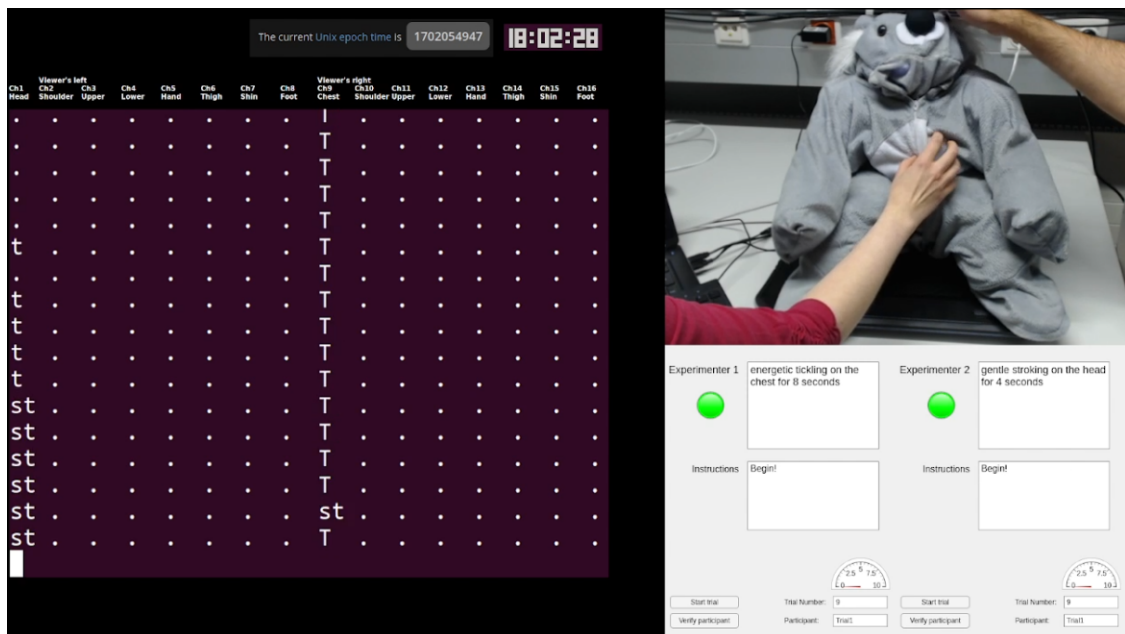


Figure 6.8: Real-time classifications can be seen for all sixteen sensors on the left, generated as the two experimenters touched HERA (top-right) based on instructions from the GUI (bottom-right). Lowercase and capitalized letters represent gentle and energetic touches, respectively. A period represents no touch. Each set of one or two letters corresponds to one of the five social touch gestures: “h” for hitting, “p” for poking, “st” for stroking, “sq” for squeezing, and “t” for tickling.

boring analog input channels) but were next to each other physically sometimes were both activated by a user's touch. This observation supported our theory of the sensors being mechanically coupled by the robot's body and the koala suit. A sample of our initial real-time testing results can be seen in Figure 6.9. The plotted sensor data was normalized by subtracting the mean value of each channel.

Further comparison of the trial videos to the trial data reaffirmed this finding. For example, we observed that the whole robot sometimes wobbled when energetically touched, which caused multiple sensors to be activated in similar patterns to the main sensor being touched. Additionally, during this real-time data collection trial, the robot's hand was slightly resting on its knee. When an experimenter energetically hit the left foot, we observed similar sensor activity for the left hand because it was being slightly bounced up and down. Conversely, when only one participant at a time was gently touching the robot, there was very little activity on the other sensors. The koala suit on top of the sensors also contributes to this mechanical coupling, as it transfers applied pressure to a larger region near the contact point. Along with the trials from our real-time data collection, inspection of our training dataset showed that this mechanical coupling phenomenon also appeared in our training data. This mechanical coupling was important for us to be aware of as we refined our updated classifier models.

### 6.3.3 Creating different models

While we iteratively explored several gradient-boosting models, here, we highlight two of our final models in particular. We refer to these as the “balanced” and “unbalanced” models. Originally, when training with the new dataset, we utilized only the labelled touches, i.e., only the data from the sensor that was actively being touched by the experimenter. The data from the other fifteen sensors that were not being touched were excluded from training. However, after observing the initial results of our real-time testing, we decided to utilize these additional “no touch” examples in training, as they show what the other areas of the robot's body feel when one body part is being touched. This approach enabled us to teach the classifier about the system's mechanical crosstalk and by extension provide a more realistic training set. For the balanced model, the loss was adjusted to account for the imbalance in training data that came from having so many more examples of “no touch” than of the touch examples. For the unbalanced model, each class was treated equally in the loss function.

### 6.3.4 Classification results: offline evaluation

We analyzed the success of our classifier across multiple levels. We first present the offline testing results of these two models. In this case, the gesture classification system predicted both the type of social touch and the force intensity. In Figure 6.10, we present the testing results for both the balanced and unbalanced models in row-normalized confusion matrices, also called true normalized. The values across a row add up to 1.00,

### 6.3 Upgrades to the Gesture Classification Algorithm

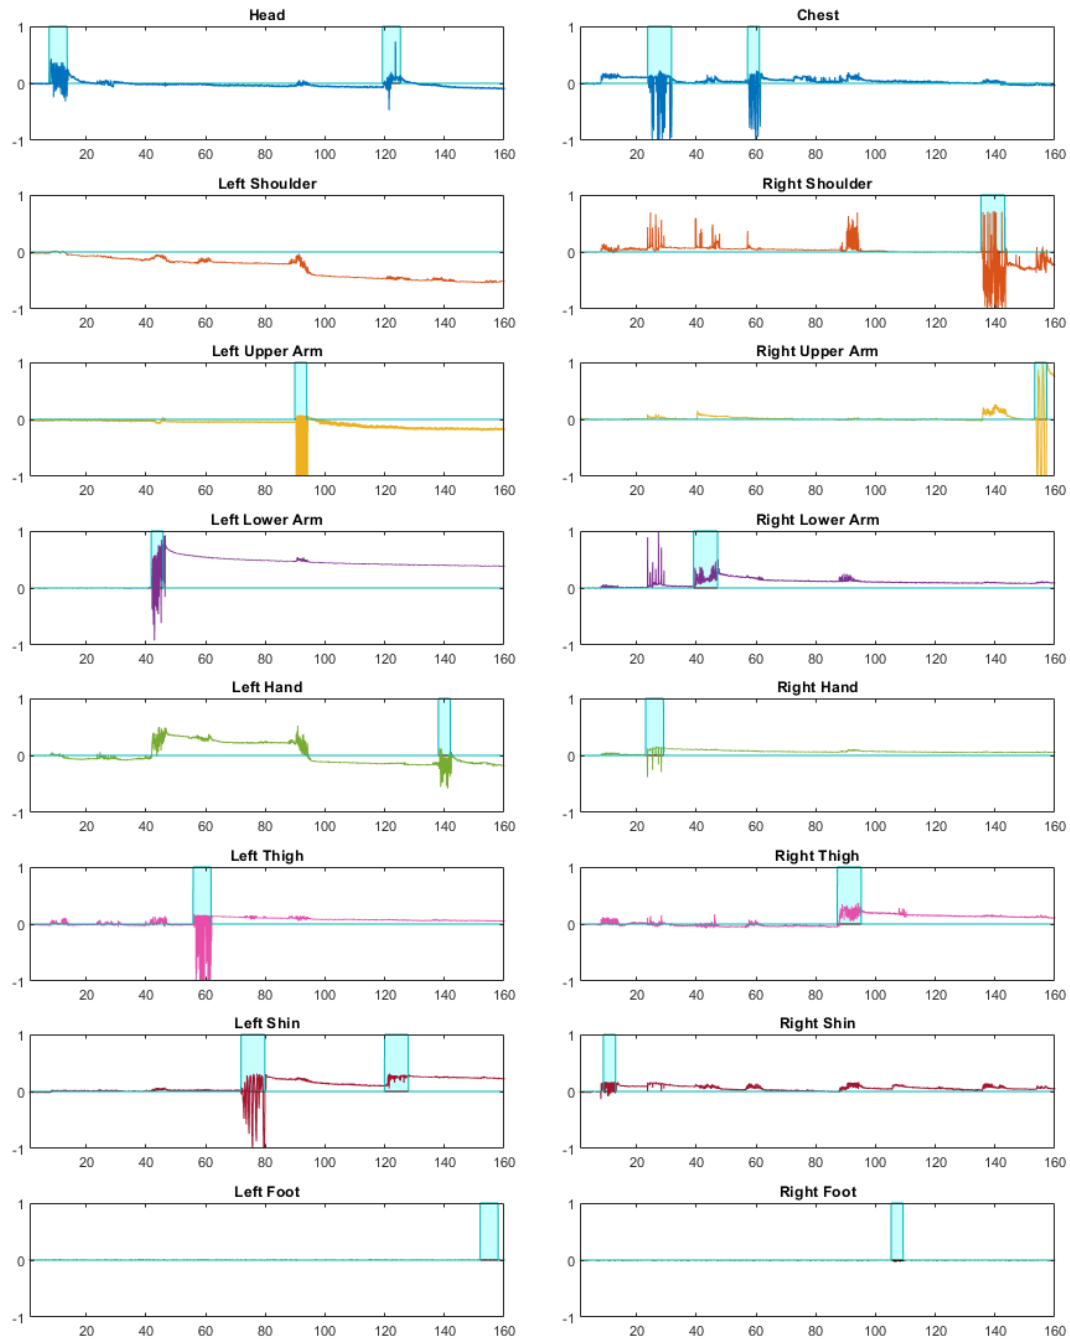


Figure 6.9: Normalized sensor data from the real-time test setup over a duration of 160 s. Light blue boxes indicate the time spans during which the experimenters were instructed to touch HERA. Instances of mechanical coupling across physically neighboring sensors can be observed. For example, one experimenter touched the robot's left lower arm at 40 s, which also caused changes in the left hand and the thigh sensor values.

or 100% of all the predictions the classifier made for one true label. Row normalization allows us to evaluate the system's **recall**, or the system's ability to predict true positives.

The balanced model demonstrated better recall than the unbalanced model. It had an overall average accuracy of 63%. With the chance of randomly guessing the correct label being  $1/11 = 9\%$ , the balanced model performs 7 times above the rate of chance. The unbalanced model had an average accuracy of 42.7%, or approximately 4.7 times better than chance.

Figure 6.11 presents the testing results for both models in a column-normalized format. These are also called prediction-normalized confusion matrices. For this style of confusion matrix, all the values in one column add up to 1.00. Column-normalization allows us to evaluate the system's **precision**, or the quality of its predictions. For example, for the balanced model, 82% of all windows predicted as gentle hitting were misidentified - the true label was actually nowhere (no touch). 12% of touches predicted as gentle hitting were truly gentle hitting, and 2% of touches predicted as gentle hitting were actually energetic hitting. In contrast, for the unbalanced model, 35% of touches predicted as gentle hitting were truly gentle hitting, 23% were energetic hitting, 13% were energetic poking (a misclassification that we also experienced with our old model), and only 7% were actually no touch, misclassified as gentle hitting.

While the case of gentle hitting is a rather extreme example, in this type of prediction-normalized evaluation, we see that overall the unbalanced model performs better than the balanced model. The balanced model had an overall average accuracy of 45.1%. With a random chance of correctly guessing the correct label being  $1/11 = 9\%$ , the balanced model performs roughly 5.0 times better than chance in terms of precision. The unbalanced model had an average accuracy of 51.9%, roughly 5.8 times better than chance. Importantly, in the unbalanced model, there was a dramatic decrease in true instances of no touch being mislabelled as a touch, i.e., fewer false positives.

### 6.3.5 Classification results: real-time evaluation

Next, we present our evaluation of the models' real-time performance. In order to address our findings of mechanical coupling, and based on the insights gathered from the offline testing, we evaluated the models' performance based on their ability to correctly classify whether or not touch had occurred.

A **true negative** occurred when the expected label and the predicted label were both no touch. For a **false positive**, the expected label was no touch, but the classifier predicted that touch had occurred. Conversely, a **false negative** meant that the expected label was touch, but the classifier predicted no touch. Finally, a **true positive** occurred when the expected label and the predicted label were both touch. Figure 6.12 shows the balanced and unbalanced models accuracy on the real-time datasets in terms of touch versus no touch, averaged across all sensor locations.

It is also important to note that the majority of samples from the real-time datasets are of instances where no touching occurred. As can be seen in Figure 6.9, the experimenters

### 6.3 Upgrades to the Gesture Classification Algorithm

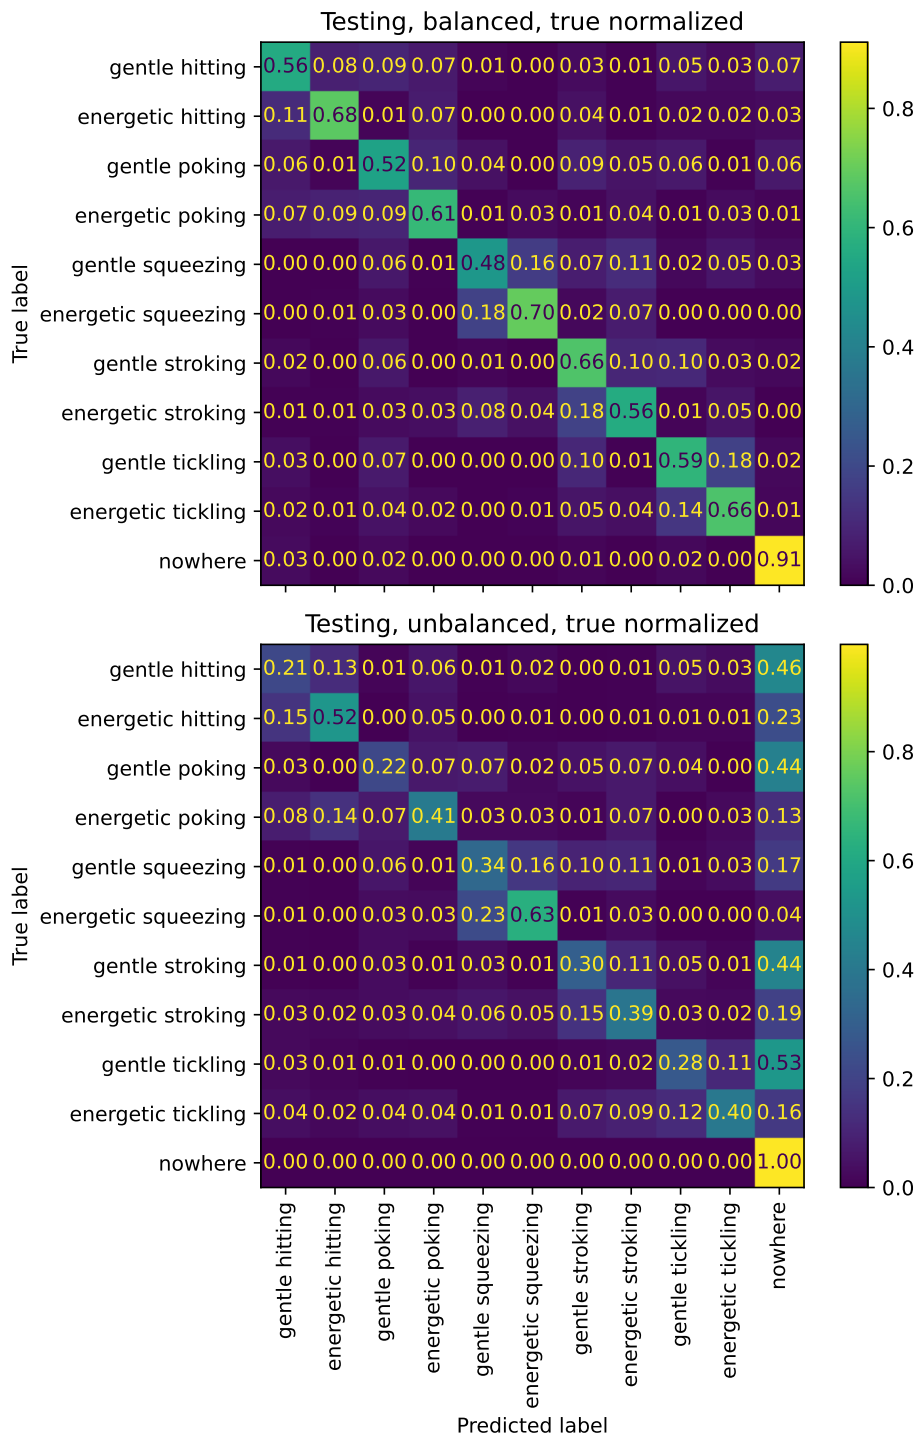


Figure 6.10: Row-normalized confusion matrices showing the testing performance of the balanced (top) and unbalanced (bottom) models. The unbalanced model has more false negatives, but also fewer false positive predictions than the balanced model.

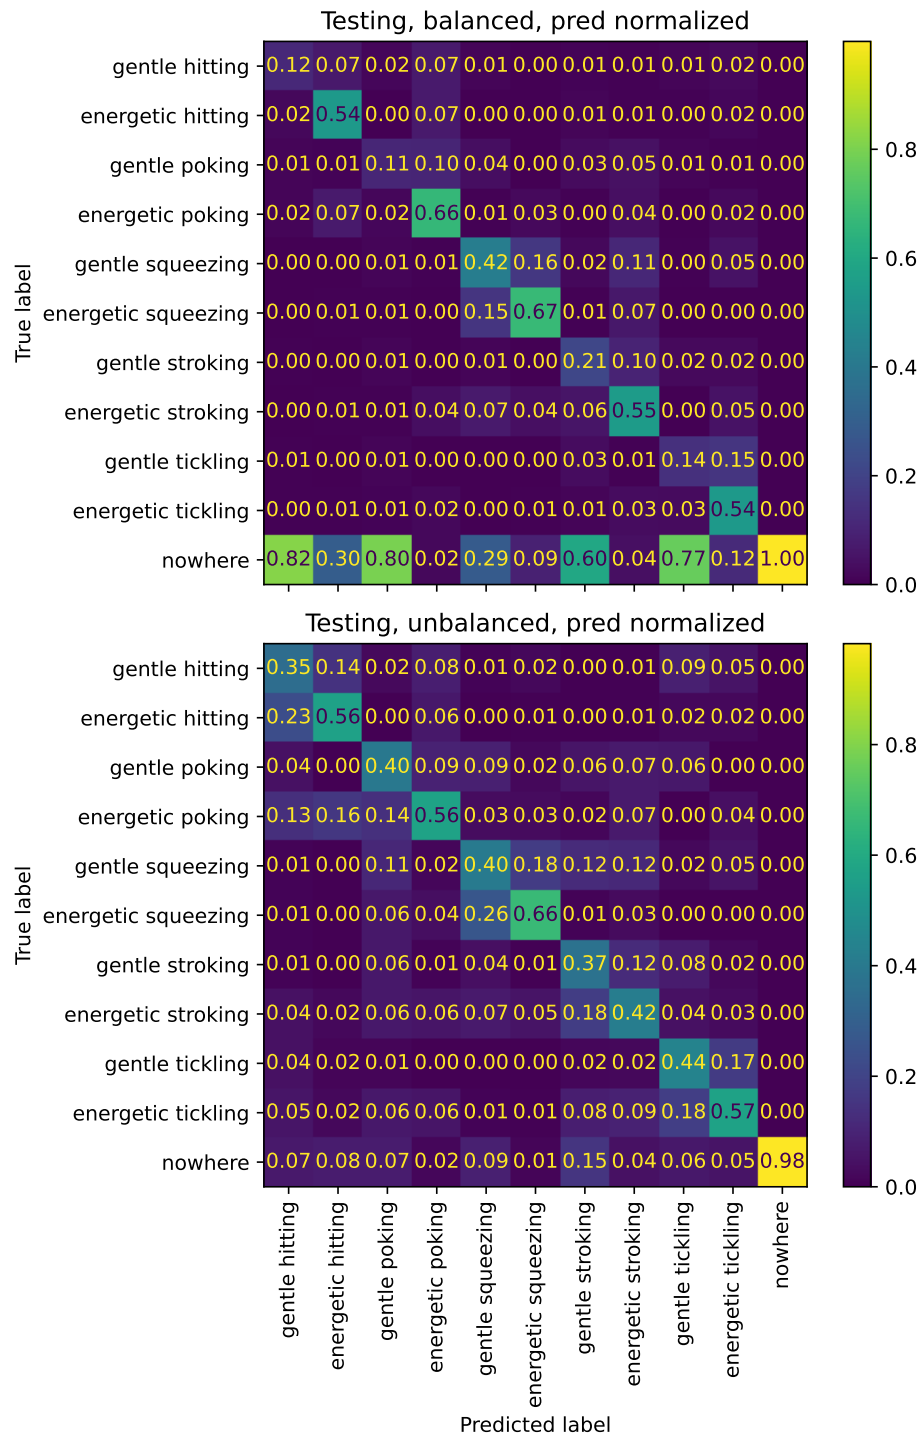


Figure 6.11: Column-normalized confusion matrices showing the testing performance of the balanced (top) and unbalanced (bottom) models.

### 6.3 Upgrades to the Gesture Classification Algorithm

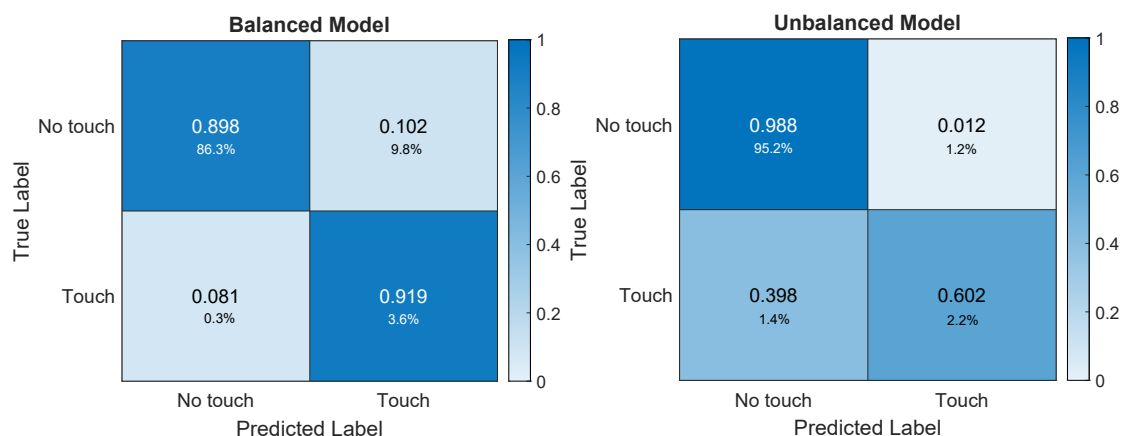


Figure 6.12: Confusion matrices showing the testing performance of the balanced (left) and unbalanced (right) models averaged across all sensor locations. Percentages indicate the proportion of the dataset that each quadrant represents.

were instructed to touch only two sensors, maximum, at one time. Furthermore, there were ten seconds of no touch between each new touch instruction, and the overall accuracy is measured across all sixteen sensors. For the two real-time trials we conducted for the balanced model, 96.1% of the dataset had a true label of no touch, whereas only 3.9% of the dataset had a true label of touch. For the two real-time trials we conducted the unbalanced model, 96.4% of the data had a true label of no touch, and 3.6% was touch.

Figure 6.13 and Figure 6.14 present confusion matrices for the real-time testing of the balanced and unbalanced models, separated by sensor location. These graphs enabled us to evaluate the performance of individual sensors. For example, we already knew that HERA’s foot sensors had lower sensitivity than the other sensors. This lack of sensitivity is confirmed with the balanced model classifier results, as while other sensors have near-perfect true positive scores, the left foot detected only 50% of touches.

In general, the unbalanced model had a higher rate of false negatives (i.e., missing touches that happened) than false positives (i.e., feeling touches when nothing happened). This meant that the robot was less likely to react to “ghost touches”. While in an ideal world, HERA would always correctly identify whether a touch has occurred or not, we felt it would be better for the robot to be slightly less reactive to touches than to react to non-existent touches. We believe that a robot that is slightly less reactive to touches could be seen by users as more socially acceptable – for example, perhaps the robot is shy, or stoic, or the user simply should touch slightly harder for the robot to react. On the opposite end, we believe that a robot that reacts to non-existent touches would be seen as defective, off-putting, or confusing.




Figure 6.13: Confusion matrices evaluating the performance of the balanced model by sensor location, based on sensor data from the real-time testing. This model outperforms the unbalanced model in terms of true positive predictions.

### 6.3 Upgrades to the Gesture Classification Algorithm

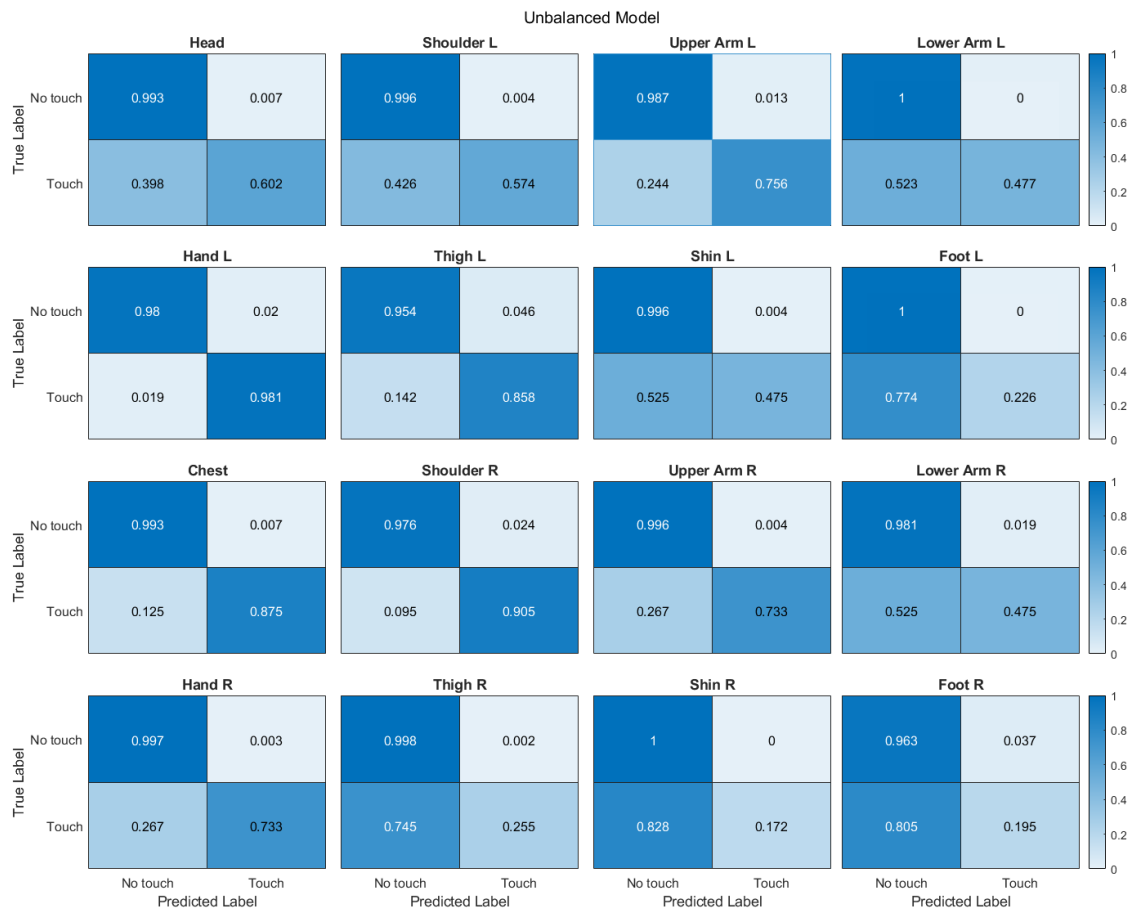


Figure 6.14: Confusion matrices evaluating the performance of the unbalanced model by sensor location, based on sensor data from the real-time testing. While touch accuracy is lower for this model than the balanced version, no touch conditions (which account for 96.4% of the dataset) are classified with an average accuracy of 98.8%.

### **6.3.6 Integrating the emotion model and gesture classifier**

The gesture classifier and emotion model are connected so that both subsystems can function in real time. The sensor data is transmitted wirelessly from the microcontroller sender to the microcontroller receiver. The sixteen parallel classifiers continually make predictions on successive windows of data coming in from the serial queue. The sixteen classification predictions are broadcast to a local Transmission Control Protocol (TCP) socket. The emotion model runs in parallel to the classifier and reads the prediction values coming from the TCP socket. When the emotion model starts, it reads in the weights of the ten combinations of social touch gestures and force intensities as an argument array, which can be customized. When no touch is detected across all sensors, no force pulse is applied to the mass representing the robot's emotional state. When any touch is detected, the force pulse applied to the emotion model at that moment is equal to the sum of these varying weights, pushing it in a positive or negative direction.

## **6.4 Discussion**

### **6.4.1 Limitations**

There are ways in which the work presented in this chapter can be further improved. For example, we evaluated the models' real-time performance by looking at touched vs. not touched conditions. A more detailed assessment of the models' success could be gained by evaluating narrower metrics, such as the system's performance based on force intensity, type of touch gesture, and the combination of these two features. We prioritized touch vs. no touch as our main assessment, as we thought it was most crucial to have a robot that did not react to "ghost touches". We also believe that a robot that misses some touches would be perceived better by users than a robot that reacts to non-existent touches. Future research could investigate these metrics and hypotheses.

Another way our real-time evaluation could be improved is with the real-time test setup we provided. For each of the two models, the experimenters conducted the same touches with the same timing, to the best of their ability. However, this method resulted in a comparison that was not exactly one-to-one. For example, it is possible that the experimenters may have performed the touches differently, even if slightly so, between the two models. In the future, the models should be compared with exactly the same data. The classifier code could be updated so that the same sensor data which was predicted on in real time can also be used again later with another model. That later model could also analyze the data as if it was being received in real time.

Additionally, with our current setup of the emotion model, a predicted touch will generate the same weight force pulse regardless of the sensor location touched. These weights could be customized depending on the sensor location. For example, tickling the robot on the foot could be registered as a stronger weight than tickling it on the head. This augmentation could allow for greater customization of the robot's responses. For

our model, we wanted to create an initial demonstration to show that these subsystems could communicate in real time. We will expand on the emotion weights based on location in our future steps.

Finally, we used a gradient boosting classifier and analyzed a fixed set of commonly used hand-crafted features. However, the classifier's performance might be improved if it was given additional features. For example, one could include features showing how strongly a sensor is being stimulated compared to other sensors. As the classifier is run independently for each of the sixteen sensors, each local classifier could then know how large its current signal fluctuation is relative to the largest overall. This spatial comparison might enable the classifier to better mask out touches that are activated due to mechanical coupling. Furthermore, it is reminiscent of tactile masking in human skin. Further research in this direction could enable our tactile-perception system to take on an even greater bio-inspired approach than its current design.

### **6.4.2 Implications for future research**

We are excited to present HERA and its real-time touch-perception and emotion capabilities. While this thesis has focused on HERA's technological preparation, creation, and testing, we also have plans for next steps.

Most immediately, we will conduct a user study in which participants interact with HERA in real time. While the positive or negative feeling of individual touches may be a personal preference, we will assign values based on the participant feedback we received in Section 4.3.5. Future research could also utilize the feedback of experts. For example, autism therapists should directly assign the robot's touch reactions (as well as its emotional model parameters) when working with children with autism, customizing the settings for each child's unique needs. The aim of this planned user study will be to demonstrate the functionality of the integrated system and investigate whether users can clearly perceive HERA's changing emotional state, based on outputs such as verbalizations and movements. We will also evaluate users' opinions of interacting with HERA, as we have done previously in Chapters 4 and 5.



# Chapter 7

## Conclusion

### 7.1 Summary

The dissertation concludes with a summary of what was presented and a discussion on potential future research directions involving robots that perceive and react to social touch.

#### **7.1.1 Establishing guidelines for a touch-perceiving robot for children with autism**

In Chapter 3, we established the first-ever touch-perception guidelines for robots interacting with children with autism. We accomplished this by thoroughly searching the existing literature to create a preliminary set of guidelines, which we then asked eleven autism specialists to evaluate and comment on through detailed interviews. We utilized thematic analysis to find three overarching themes within their feedback. From these themes, we refined our initial guidelines into seven qualitative requirements, which we also translated into quantitative specifications for engineers. Along with creating these touch-perception guidelines, we also gained meaningful insights on what types of touches children with autism were most likely to use, what locations on the robots they were most likely to touch, and in what sort of roles a social robot companion would be most beneficial for the autism support community.

#### **7.1.2 Endowing a robot companion with practical social-touch perception**

In Chapter 4, we introduced custom fabric-and-foam-based tactile sensors to meet the autism specialists' requests, and we provided open-access DIY instructions and materials so that others can create sensors for their own robot systems. We also introduced HERA, our prototype robot companion. We ran a preliminary study in which fifteen adult participants touched HERA, using five touch types, four sensor locations, and two forces intensities. We evaluated our touch dataset using a gesture classification algorithm based on a random forest. We found that when classifying both the type of touch and

force intensity combined, touches were identified with an average accuracy of 74.1%. We demonstrated that using our simple sensor design, with only four time-series sensor channels from the four sensors, one can produce rich tactile data signals with promising accuracy for social contexts.

### **7.1.3 Investigating user preferences for robot emotions in response to touch**

In Chapter 5, we investigated users' preferences on whether and how a robot should emotionally react to touch, over short and long intervals. We conducted a video study in which participants watched, rated, and provided comments on nine randomly ordered videos. In these videos, HERA showed varying levels of immediate reactions and long-term moods in response to a positive stroking touch and a negative hitting touch. Our 51 participants were diverse, with a wide variety of ages, home countries, experience with robots, and gender, in order to establish a general finding across audiences. Overall, we found that users most highly rated the scenarios in which HERA demonstrated immediate emotion reactions and either neutral or emotional moods. Additionally, we provided key qualitative feedback about users' preferences for reaction timing, how robot mood helps demonstrate persisting memory, and how neutral behaviors were perceived as a curious or self-aware robot.

### **7.1.4 Creating a robot that feels both touch and emotion**

In Chapter 6, we utilized our findings to make HERA fully operational. We proposed a mathematical emotion model which enables social robots to respond to external stimuli (such as touch) with immediate and long-term emotions. We modeled the robot's affective state as a second-order dynamic system, analogous to a mass connected to ground by a parallel spring and damper, which enabled us to represent emotions in a naturalistic way and to capture the essential features of approach-avoidance theory. We also introduced extensive improvements to HERA's tactile-perception system, including upgrades to the sensors, microcontroller, and gesture classification approach. Finally, we evaluated the success of two updated classifier models, explained the connections between the classifier and emotion subsystems, and discussed future research implications.

## **7.2 Future Work**

For our next steps, we will conduct a user study in which adult participants interact with HERA in a playful and lightly structured setting. We will update HERA's audio cues and physical movements based on the feedback we received from the participants in Chapter 5, along with feedback from pilot participants. We will investigate whether HERA's emotional state can be understood by adults based on these cues, before moving

on to child participants. Future studies with both neurotypical children and children with autism will be done after careful consultation and approval from the Max Planck Ethics Council. Furthermore, longitudinal studies are essential to proving the efficacy of socially assistive robots and for the clinical community to adopt their use. If HERA is used as an educational platform for the autism support community, it will be important to work closely with a therapy team over multiple sessions, in order to evaluate whether working with HERA indeed produces long-term benefits.

In addition to further developing social touch education with autism therapists, giving robots a sense of social touch has far-reaching benefits which can expand to many other domains both within and outside of medicine. For example, robots that understand and react to touch could be beneficial for individuals experiencing loneliness or stress due to being separated from their loved ones, such as college students, older adults in nursing homes, or individuals in extended hospital stays. In the future, everyday hospital and household robots should also be able to feel social touch gestures across their entire bodies, to facilitate communication and increase safety.

The research presented in this dissertation has focused on passive social touch, i.e., the robot receives the touch. However, active touch, wherein the robot touches the user, is also a burgeoning topic of interest in the HRI community. Research in this area has already begun, with researchers investigating how it feels to play clapping games [54], to high-five [108], and hug robots [18]. Researchers might also want to investigate how passive and active touch influence the user's sense of trust toward the robot system, as trust is currently a high-priority topic in the HRI research community.

In this dissertation, we identified that a myriad of opportunities would be available if robots can understand and react to social touch (Chapter 3). We found that participants greatly enjoyed touching a soft and squishy robot, even before it was emotionally reactive (Chapter 4). We also discovered that users find a robot most intelligent and enjoyable when it reacts to touch with both immediate emotional reactions and ambient moods (Chapter 5). We combined these findings to create a Haptic Empathetic Robot Animal (Chapter 6). We hope to inspire other researchers to get involved in creating similar touch-perceiving intelligent systems, perhaps using the DIY technology and open-access tools we have provided. Such robots could have the potential to touch the lives of humans in ways yet to be discovered.



# Symbols

$b$	Damping coefficient
$^{\circ}\text{C}$	Degrees Celsius
$\hat{\epsilon}$	Greenhouse-Geisser correction
$\text{€}$	Euro
$F_d$	Force applied by the damper
$F_s$	Force applied by the spring
$F_p$	Fixed duration force pulse generated by external stimuli
$k$	Spring stiffness
$k\Omega$	Kiloohm(s)
$L$	Mass position at rest
$m$	Mass
$n$	Number of samples/participants
$\Omega$	Ohm(s)
$p$	Probability
$R$	Resistance
$R_0$	Initial resistance
$w$	Window size
$\omega_n$	Natural frequency
$y$	Position
$\dot{y}$	Velocity
$\ddot{y}$	Acceleration
$\zeta$	Damping ratio



# Abbreviations

3D	Three-dimensional
AAC	Alternative and augmentative communication
ABA	Applied behavior analyst
AI	Artificial intelligence
ASD	Autism spectrum disorder
BCBA	Board-certified behavior analyst
CAD	Computer-aided design
COVID-19	Coronavirus disease
cm	Centimeter(s)
DIY	Do-it-yourself
DTP	Deep-touch pressure
FSR	Force-sensing resistor
GUI	Graphical user interface
HERA	Haptic Empathetic Robot Animal
HRI	Human-robot interaction
Hz	Sampling rate in Hertz
IMPRS-IS	International Max Planck Research School for Intelligent Systems
IQR	Interquartile range
kg	kilogram(s)
kPa	Kilopascal(s)
mm	Millimeter(s)
N	Newton(s)
OT	Occupational therapist
P	Participant
PT	Physical therapist
RANOVA	Repeated-measures analysis of variance
RF	Radio frequency
s	second(s)
SAR	Socially assistive robotics
SD	Standard deviation
SLP	Speech-language pathologist
SNR	Signal-to-noise ratio
SVM	Support vector machine
TCP	Transmission Control Protocol

## *Abbreviations*

---

N	Newton(s)
ms	millisecond(s)

# List of Tables

3.1	Initial list of six touch-sensing requirements that we derived from the literature. Specific sources that inspired each requirement are cited in the “Motivation” column. For each requirement, we also show the follow-up question that the experimenter asked to spark additional feedback after completion of the ranking task. . . . .	19
3.2	Summary of participant experience and backgrounds. . . . .	21
3.3	Our final qualitative guidelines for a touch-sensing robot companion for children with autism, as derived from the recommendations of 11 autism specialists. Requirements are listed in descending order of importance, with 1 indicating most important. . . . .	38
3.4	The minimum quantitative specifications we propose for a touch-perceiving companion robot for children with autism. These quantitative specifications were translated from our final qualitative requirements (Table 3.3). . . . .	39
4.1	The seven questions in the robot acceptance survey asked in both the opening and closing evaluations. . . . .	53
4.2	The parameter values that were used for the gesture-classification algorithms. Unless otherwise specified, the same value was used for classifying contact location, intensity, gesture, and gesture and intensity together. . . . .	56
5.1	A comparison of the nine robot behavior conditions shown in the study, each including a mood level and a reaction level. . . . .	73
5.2	The results of two-way RANOVA analysis for each category. . . . .	79



# List of Figures

3.1	The prototype robot companion for children with autism that was presented to the therapists in our interview study. It is a NAO robot (Soft-Bank Robotics); the NAO robot shown on the left is wearing a thin fabric shirt, and the one on the right is enclosed in a soft koala suit. Photo redistributed from [22] under a Creative Commons license – Attribution 4.0 International (CC BY 4.0). . . . .	23
3.2	The requirement-ranking interface enabled participants to prioritize tactile-sensing requirements for the robot as they preferred. Participants could reorder the six provided requirements by clicking on and dragging the respective slides in the navigation pane on the left. The main view enabled the participant to clearly see the individual requirements and the motivations that inspired them. Illustration redistributed from [22], CC BY. . . . .	24
3.3	A photograph of the prototype fabric-based tactile sensor for a robot companion that was shown to participants. Participants were told that several such sensors could be placed across the robot’s body, and that they could be used to detect the intensity and/or type of touch being enacted on the sensor. Participants were asked to give feedback and share any concerns they might have with such a design. Photo from [22], CC BY. . . . .	25
3.4	A diagram illustrating the flow of our process and the links between our literature review, interview study, and resulting themes. The thematic analysis portion of our study is highlighted in green. The identified themes were used to refine our initial requirement list into finalized qualitative requirements, which we then translated into quantitative specifications. Figure adapted from [22], CC BY. . . . .	26
3.5	Locations that participants specifically prioritized to equip with touch sensing, displayed on a diagram of a generalized robot companion. The number of therapists who suggested each region is shown in parentheses. Most of our eleven participants requested whole-body touch sensing, but they also provided these specific locations in case whole-body sensing is not possible. Illustration from [22], CC BY. . . . .	29

3.6	A visualization of the touch gestures that specialists recommended (or did not recommend) a touch-perceiving robot companion should be able to detect from children with autism. An asterisk above a gesture indicates that it was selected as a “top five” gesture by more than three participants. Figure from [22], CC BY. . . . .	34
3.7	The participants’ individual responses for the requirement-ranking task. A ranking of 1 indicates the most important requirement, and 6 indicates the least important. A summary of each requirement’s rankings is also displayed at the top of each panel using a green opacity overlay – the darkness of a number’s green background indicates the number of participants who selected that rank for that requirement. White means that no participants chose that ranking. The asterisk above each overall summary indicates that requirement’s median ranking. Participant 10 completed two different versions of the ranking – one for using the robot in a school setting (P10 <sub>S</sub> ), and the second for using it in a home (P10 <sub>H</sub> ). Figure from [22], CC BY. . . . .	37
4.1	Left: Custom tactile sensors were built from fabric and foam to wrap around the four rigid segments of NAO’s arm. Right: As seen in this photo-merge, the sensors are secured on top of the robot’s plastic exterior and are hidden underneath a soft koala suit to create the robot companion HERA. Photos redistributed from [24] under a Creative Commons License, CC BY. . . . .	47
4.2	A visual guide highlighting the eleven key steps required for building our fabric-based tactile sensors. Illustration from [24], CC BY. . . . .	50
4.3	Real-time sensor readings show how the four sensor resistance values change (left) as the participant interacts with the robot (right). The four × symbols on the plot mark the experimenter’s timestamps. In this example, the participant performed an energetic tickling gesture on the robot’s hand between 90 and 96 seconds (s) and performed energetic stroking on the same body part between 101 and 107 s. Figure from [24], CC BY. . . . .	54
4.4	Normalized confusion matrices showing the classification results for location, force intensity, and gesture. Graphs from [24], CC BY. . . . .	57
4.5	Normalized confusion matrices showing the classification results for both force intensity and gesture together. Figure from [24], CC BY. . . . .	58
4.6	A box-and-whisker plot comparing participants’ level of agreement with seven survey statements before and after interacting with the robot. For each distribution, the central line indicates the median, the box shows the interquartile range (IQR), the whiskers show the range up to 1.5 times the IQR, and + marks indicate outliers. Plot from [24], CC BY. . . . .	60

4.7	A box-and-whisker plot showcasing participants' ratings of how they expect the robot would react in the future to the various touch gestures. A rating of 10 indicates a very positive response, and a rating of 0 indicates a very negative response. Touches with a gentle force intensity are shown in blue, and energetic gesture variations are colored red. Plot from [24], CC BY. . . . .	61
4.8	Box-and-whisker plots indicating the range of sensor resistance readings during all participants' cumulative gesture performances. Gentle trials are shown in a pale color directly next to the more darkly colored energetic trials of the same type. Figure redistributed from [24], CC BY. . .	62
4.9	Normalized comparisons of individual sensor's resistances across the tested force range with and without the koala suit, averaged across all trials of each physical experiment. The initial and final sensitivities are represented using semi-transparent lines oriented at the appropriate slopes. Each plot also lists the corresponding numerical values for sensitivity and hysteresis. Graphs from [24], CC BY. . . . .	64
5.1	A screenshot from one of the nine videos in our study, which showcased three levels of immediate reaction to social touch and three long-term mood responses by HERA. Photo redistributed from [27], ©2023 IEEE.	72
5.2	Timeline illustrating the order of stimuli and responses that viewers saw in each of the nine video conditions. While the robot's reaction and mood levels varied by condition, the timing and sequence of events remained the same across all nine videos. Illustration from [27], ©2023 IEEE. . .	75
5.3	The statements that participants rated for each of the videos, numbered by presentation order. The boxes show how the eight statements were grouped into categories, and the arrows show the categories used to evaluate each hypothesis. Figure from [27], ©2023 IEEE. . . . .	76
5.4	Box-and-whisker comparisons of combined, normalized user ratings of the video conditions. Lines above the boxplots marked with $\ominus$ indicate pairwise comparisons that are NOT significantly different. For each distribution, the central line indicates the median, the box shows the IQR, the whiskers show the range up to 1.5 times the IQR, and $+$ marks indicate outliers. The legend uses the abbreviations established in Table 5.1. Plot from [27], ©2023 IEEE. . . . .	78
6.1	Modeling a robot's internal affective state using dynamic systems of increasing orders. The $-$ and $+$ symbols indicate negative and positive stimuli, respectively. . . . .	86
6.2	Our emotion model can be represented as a mass-spring-damper system. We illustrate its properties both at rest (left, $y(t) = L$ ) and when past force pulses have moved the mass away from its neutral position (right, $y(t) > L$ ). . . . .	87

6.3	The height of the robot’s smile or frown directly corresponds to the internal valence. Blue squares indicate the valence that corresponds with the picture above it. . . . .	88
6.4	In each subplot, one parameter is changed to showcase how the robot’s response to stimuli can be customized. All graphs use $m = 1$ kg, $F_p = 30$ N, and pulse duration $T = 1.0$ s. The middle blue trial is the same for all three subplots. . . . .	90
6.5	A view of all the sensors hidden under HERA’s koala suit. The inner NAO robot is covered with sixteen improved sensors, which provide HERA with whole-body tactile sensing. Moving joints were not covered to avoid pinching. . . . .	92
6.6	The sender (right) is secured to HERA’s back under the koala suit, and the receiver (left) is connected to a computer. . . . .	93
6.7	The GUI that the experimenters used to build the new training dataset. Experimenters could view upcoming touch instructions, watch plots of the sensor data in real time if necessary, and keep track of how many trials had been completed. . . . .	95
6.8	Real-time classifications can be seen for all sixteen sensors on the left, generated as the two experimenters touched HERA (top-right) based on instructions from the GUI (bottom-right). Lowercase and capitalized letters represent gentle and energetic touches, respectively. A period represents no touch. Each set of one or two letters corresponds to one of the five social touch gestures: “h” for hitting, “p” for poking, “st” for stroking, “sq” for squeezing, and “t” for tickling. . . . .	97
6.9	Normalized sensor data from the real-time test setup over a duration of 160 s. Light blue boxes indicate the time spans during which the experimenters were instructed to touch HERA. Instances of mechanical coupling across physically neighboring sensors can be observed. For example, one experimenter touched the robot’s left lower arm at 40 s, which also caused changes in the left hand and the thigh sensor values. . . . .	99
6.10	Row-normalized confusion matrices showing the testing performance of the balanced (top) and unbalanced (bottom) models. The unbalanced model has more false negatives, but also fewer false positive predictions than the balanced model. . . . .	101
6.11	Column-normalized confusion matrices showing the testing performance of the balanced (top) and unbalanced (bottom) models. . . . .	102
6.12	Confusion matrices showing the testing performance of the balanced (left) and unbalanced (right) models averaged across all sensor locations. Percentages indicate the proportion of the dataset that each quadrant represents. . . . .	103

6.13	Confusion matrices evaluating the performance of the balanced model by sensor location, based on sensor data from the real-time testing. This model outperforms the unbalanced model in terms of true positive predictions. . . . .	104
6.14	Confusion matrices evaluating the performance of the unbalanced model by sensor location, based on sensor data from the real-time testing. While touch accuracy is lower for this model than the balanced version, no touch conditions (which account for 96.4% of the dataset) are classified with an average accuracy of 98.8%. . . . .	105



# Bibliography

- [1] 3M – Electronics Materials Solutions Division (2015). 3M thermal bonding film 583 - technical data. <https://multimedia.3m.com/mws/media/169553O/3m-thermal-bonding-film-583.pdf>.
- [2] Altun, K. and MacLean, K. E. (2015). Recognizing affect in human touch of a robot. *Pattern Recognition Letters*, **66**, 31–40.
- [3] Andreasson, R., Alenljung, B., Billing, E., and Lowe, R. (2018). Affective touch in human–robot interaction: Conveying emotion to the NAO robot. *International Journal of Social Robotics*, **10**(4), 473–491.
- [4] Arduino (2018). Digital pins — Arduino. <https://www.arduino.cc/en/Tutorial/Foundations/DigitalPins>.
- [5] Argall, B. D. and Billard, A. G. (2010). A survey of tactile human–robot interactions. *Robotics and Autonomous Systems*, **58**(10), 1159–1176.
- [6] Armbruster, D. A. and Pry, T. (2008). Limit of blank, limit of detection and limit of quantitation. *The Clinical Biochemist Reviews*, **29**(Suppl 1), S49.
- [7] Ashburner, J., Ziviani, J., and Rodger, S. (2008). Sensory processing and classroom emotional, behavioral, and educational outcomes in children with autism spectrum disorder. *American Journal of Occupational Therapy*, **62**(5), 564–573.
- [8] Aylett, R. S., Louchart, S., Dias, J., Paiva, A., and Vala, M. (2005). Fearnot!—an experiment in emergent narrative. In *International Workshop on Intelligent Virtual Agents*, pages 305–316, Kos, Greece. Springer.
- [9] Bass, M. M., Duchowny, C. A., and Llabre, M. M. (2009). The effect of therapeutic horseback riding on social functioning in children with autism. *Journal of Autism and Developmental Disorders*, **39**(9), 1261–1267.
- [10] Beck, A., Cañamero, L., and Bard, K. A. (2010). Towards an affect space for robots to display emotional body language. In *Proceedings of the IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, pages 464–469, Viareggio, Italy. IEEE.

- [11] Begum, M., Serna, R. W., and Yanco, H. A. (2016). Are robots ready to deliver autism interventions? A comprehensive review. *International Journal of Social Robotics*, **8**(2), 157–181.
- [12] Bentéjac, C., Csörgő, A., and Martínez-Muñoz, G. (2021). A comparative analysis of gradient boosting algorithms. *Artificial Intelligence Review*, **54**, 1937–1967.
- [13] Bestbier, L. and Williams, T. I. (2017). The immediate effects of deep pressure on young people with autism and severe intellectual difficulties: demonstrating individual differences. *Occupational Therapy International*, **2017**.
- [14] Bethel, C. L., Henkel, Z., Darrow, S., and Baugus, K. (2018). Therabot – an adaptive therapeutic support robot. In *Proceedings of the IEEE World Symposium on Digital Intelligence for Systems and Machines (DISA)*, pages 23–30.
- [15] Bilyea, A., Seth, N., Nesathurai, S., and Abdullah, H. (2017). Robotic assistants in personal care: A scoping review. *Medical Engineering & Physics*, **49**, 1–6.
- [16] Biswas, M., Romeo, M., Cangelosi, A., and Jones, R. B. (2020). Are older people any different from younger people in the way they want to interact with robots? scenario based survey. *Journal of Multimodal User Interfaces*, **14**(1), 61–72.
- [17] Block, A. E., Christen, S., Gassert, R., Hilliges, O., and Kuchenbecker, K. J. (2021). The six hug commandments: Design and evaluation of a human-sized hugging robot with visual and haptic perception. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 380–388, New York, NY, USA. ACM.
- [18] Block, A. E., Seifi, H., Hilliges, O., Gassert, R., and Kuchenbecker, K. J. (2022). In the arms of a robot: Designing autonomous hugging robots with intra-hug gestures. *ACM Transactions on Human-Robot Interaction (THRI)*, **12**(2), 1–49.
- [19] Braun, V. and Clarke, V. (2012). Thematic analysis. In *APA Handbook of Research Methods in Psychology, Vol 2: Research Designs: Quantitative, Qualitative, Neuropsychological, and Biological.*, pages 57–71. American Psychological Association.
- [20] Breazeal, C. (2009). Role of expressive behaviour for robots that learn from people. *Philosophical Transactions of the Royal Society B: Biological Sciences*, **364**(1535), 3527–3538.
- [21] Burns, R., Jeon, M., and Park, C. H. (2018). Robotic motion learning framework to promote social engagement. *Applied Sciences*, **8**(2), 241.
- [22] Burns, R. B., Seifi, H., Lee, H., and Kuchenbecker, K. J. (2021a). Getting in touch with children with autism: Specialist guidelines for a touch-perceiving robot. *Paladyn, Journal of Behavioral Robotics*, **12**(1), 115–135.

- [23] Burns, R. B., Seifi, H., Lee, H., and Kuchenbecker, K. J. (2021b). A haptic empathetic robot animal for children with autism. In *Companion of the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 583—585. IEEE.
- [24] Burns, R. B., Lee, H., Seifi, H., Faulkner, R., and Kuchenbecker, K. J. (2022a). Endowing a NAO robot with practical social-touch perception. *Frontiers in Robotics and AI*, **9**(840335), 1–17.
- [25] Burns, R. B., Lee, H., Seifi, H., Faulkner, R., and Kuchenbecker, K. J. (2022b). Sensor patterns dataset for endowing a NAO robot with practical social-touch perception. Dataset.
- [26] Burns, R. B., Lee, H., Seifi, H., Faulkner, R., and Kuchenbecker, K. J. (2022c). User study dataset for endowing a NAO robot with practical social-touch perception. Dataset.
- [27] Burns, R. B., Ojo, F., and Kuchenbecker, K. J. (2023). Wear your heart on your sleeve: Users prefer robots with emotional reactions to touch and ambient moods. In *Proceedings of the IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, pages 1914–1921. IEEE.
- [28] Büscher, G. H., Kõiva, R., Schürmann, C., Haschke, R., and Ritter, H. J. (2015). Flexible and stretchable fabric-based tactile sensor. *Robotics and Autonomous Systems*, **63**, 244–252.
- [29] Cabibihan, J.-J., Javed, H., Ang, M., and Aljunied, S. M. (2013). Why robots? a survey on the roles and benefits of social robots in the therapy of children with autism. *International Journal of Social Robotics*, **5**(4), 593–618.
- [30] Cang, X. L., Bucci, P., Strang, A., Allen, J., MacLean, K., and Liu, H. S. (2015). Different strokes and different folks: Economical dynamic surface sensing and affect-related touch recognition. In *Proceedings of the ACM International Conference on Multimodal Interaction*, pages 147–154, New York, NY, USA. Association for Computing Machinery.
- [31] Cascio, C. J., Moana-Filho, E. J., Guest, S., Nebel, M. B., Weisner, J., Baranek, G. T., and Essick, G. K. (2012). Perceptual and neural response to affective tactile texture stimulation in adults with autism spectrum disorders. *Autism Research*, **5**(4), 231–244.
- [32] Cascio, C. J., Lorenzi, J., and Baranek, G. T. (2016). Self-reported pleasantness ratings and examiner-coded defensiveness in response to touch in children with asd: Effects of stimulus material and bodily location. *Journal of Autism and Developmental Disorders*, **46**(5), 1528–1537.

- [33] Cascio, C. J., Moore, D., and McGlone, F. (2019). Social touch and human development. *Developmental Cognitive Neuroscience*, **35**, 5–11.
- [34] Cavallo, F., Semeraro, F., Fiorini, L., Magyar, G., Sinčák, P., and Dario, P. (2018). Emotion modelling for social robotics applications: a review. *Journal of Bionic Engineering*, **15**(2), 185–203.
- [35] Centers for Disease Control and Prevention (2020). Data & statistics on autism spectrum disorder. <https://www.cdc.gov/ncbddd/autism/data.html>.
- [36] Chang, J., MacLean, K., and Yohanan, S. (2010). Gesture recognition in the haptic creature. In A. M. L. Kappers, J. B. F. van Erp, W. M. B. Tiest, and F. C. T. van der Helm, editors, *Haptics: Generating and Perceiving Tangible Sensations, Proceedings of the EuroHaptics, Part I*, Lecture Notes in Computer Science, pages 385–391. Springer.
- [37] Choi, S. and Kuchenbecker, K. J. (2013). Vibrotactile display: Perception, technology, and applications. *Proceedings of the IEEE*, **101**(9), 2093–2104.
- [38] Churamani, N., Barros, P., Gunes, H., and Wermter, S. (2022). Affect-driven learning of robot behaviour for collaborative human-robot interactions. *Frontiers in Robotics and AI*, **9**, 20.
- [39] Cisco WebEx (2019). Four key security features of Cisco web conferencing. <https://blog.webex.com/video-conferencing/four-key-security-features-of-cisco-webex/>.
- [40] Cooney, M. D., Becker-Asano, C., Kanda, T., Alissandrakis, A., and Ishiguro, H. (2010). Full-body gesture recognition using inertial sensors for playful interaction with small humanoid robot. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2276–2282.
- [41] Dael, N., Mortillaro, M., and Scherer, K. R. (2012). Emotion expression in body action and posture. *Emotion*, **12**(5), 1085.
- [42] Dahiya, R. S., Metta, G., Valle, M., and Sandini, G. (2010). Tactile sensing: From humans to humanoids. *IEEE Transactions on Robotics*, **26**(1), 1–20.
- [43] Day, N., Penaloza, J., Santos, V. J., and Killpack, M. D. (2018). Scalable fabric tactile sensor arrays for soft bodies. *Journal of Micromechanics and Microengineering*, **28**(6), 064004.
- [44] Domjan, M. P. (2014). *The principles of learning and behavior*. Cengage Learning.

- 
- [45] Duquette, A., Michaud, F., and Mercier, H. (2008). Exploring the use of a mobile robot as an imitation agent with children with low-functioning autism. *Autonomous Robots*, **24**(2), 147–157.
- [46] Egges, A., Kshirsagar, S., and Magnenat-Thalmann, N. (2004). Generic personality and emotion simulation for conversational agents. *Computer Animation and Virtual Worlds*, **15**(1), 1–13.
- [47] Elliot, A. J., Eder, A. B., and Harmon-Jones, E. (2013). Approach–avoidance motivation and emotion: Convergence and divergence. *Emotion Review*, **5**(3), 308–311.
- [48] Etzi, R., Spence, C., and Gallace, A. (2014). Textures that we like to touch: An experimental study of aesthetic preferences for tactile stimuli. *Consciousness and Cognition*, **29**, 178–188.
- [49] Feil-Seifer, D. and Mataric, M. J. (2005). Defining socially assistive robotics. In *Proceedings of the IEEE International Conference on Rehabilitation Robotics (ICORR)*, pages 465–468.
- [50] Ferrari, E., Robins, B., and Dautenhahn, K. (2009). Therapeutic and educational objectives in robot assisted play for children with autism. In *Proceedings of the IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, pages 108–114. IEEE.
- [51] Ferretti, M., Morgavi, G., and Veruggio, G. (2018). How should a robot caregiver for elderly people be? *International Journal of Control Systems and Robotics*, **3**.
- [52] Fiene, J. (2011). Meam.design: The maevarm (m1 & m2) project. <http://medesign.seas.upenn.edu/index.php/Guides/MaEvArM>.
- [53] Fischer, P., Daniel, R., and Siva, K. (1990). Specification and design of input devices for teleoperation. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 540–545.
- [54] Fitter, N. T. and Kuchenbecker, K. J. (2020). How does it feel to clap hands with a robot? *International Journal of Social Robotics*, **12**(1), 113–127.
- [55] Fitter, N. T., Mohan, M., Kuchenbecker, K. J., and Johnson, M. J. (2020). Exercising with Baxter: Preliminary support for assistive social-physical human-robot interaction. *Journal of NeuroEngineering and Rehabilitation*, **17**(19).
- [56] Foxlin, E. *et al.* (2002). Motion tracking requirements and technologies. *Handbook of Virtual Environment Technology*, **8**, 163–210.

- [57] Gaus, Y. F. A., Olugbade, T., Jan, A., Qin, R., Liu, J., Zhang, F., Meng, H., and Bianchi-Berthouze, N. (2015). Social touch gesture recognition using random forest and boosting on distinct feature sets. In *Proceedings of the ACM International Conference on Multimodal Interaction*, pages 399–406.
- [58] Gordon, G., Spaulding, S., Westlund, J. K., Lee, J. J., Plummer, L., Martinez, M., Das, M., and Breazeal, C. (2016). Affective personalization of a social robot tutor for children’s second language skills. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 30, pages 1–7, Phoenix, USA. AAAI Press.
- [59] Greczek, J., Kaszubski, E., Atrash, A., and Matarić, M. (2014). Graded cueing feedback in robot-mediated imitation practice for children with autism spectrum disorders. In *Proceedings of the IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, pages 561–566. IEEE.
- [60] Harlow, H. F. and Zimmermann, R. R. (1959). Affectional responses in the infant monkey. *Science*, **130**(3373), 421–432.
- [61] Hertenstein, M. J., Holmes, R., McCullough, M., and Keltner, D. (2009). The communication of emotion via touch. *Emotion*, **9**(4), 566–573.
- [62] Isbister, K., Cottrell, P., Cecchet, A., Dagan, E., Theofanopoulou, N., Bertran, F. A., Horowitz, A. J., Mead, N., Schwartz, J. B., and Slovak, P. (2022). Design (not) lost in translation: A case study of an intimate-space socially assistive “robot” for emotion regulation. *ACM Transactions on Computer-Human Interaction (TOCHI)*, **29**(4), 1–36.
- [63] Javed, H. and Park, C. H. (2019). Interactions with an empathetic agent: Regulating emotions and improving engagement in autism. *IEEE Robotics & Automation Magazine*, **26**(2), 40–48.
- [64] Javed, H., Burns, R., Jeon, M., Howard, A. M., and Park, C. H. (2019). A robotic framework to facilitate sensory experiences for children with autism spectrum disorder: A preliminary study. *ACM Transactions on Human-Robot Interaction (THRI)*, **9**(1), 1–26.
- [65] Jeong, S., Santos, K. D., Graca, S., O’Connell, B., Anderson, L., Stenquist, N., Fitzpatrick, K., Goodenough, H., Logan, D., Weinstock, P., *et al.* (2015a). Designing a socially assistive robot for pediatric care. In *Proceedings of the International Conference on Interaction Design and Children*, pages 387–390.
- [66] Jeong, S., Logan, D. E., Goodwin, M. S., Graca, S., O’Connell, B., Goodenough, H., Anderson, L., Stenquist, N., Fitzpatrick, K., Zisook, M., *et al.* (2015b). A social robot to mitigate stress, anxiety, and pain in hospital pediatric care. In *Companion of*

- 
- the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 103–104, New York, USA. IEEE.
- [67] Jung, M. M., Poppe, R., Poel, M., and Heylen, D. K. (2014). Touching the void – introducing cost: Corpus of social touch. In *Proceedings of the International Conference on Multimodal Interaction*, pages 120–127.
- [68] Jung, M. M., Poel, M., Poppe, R., and Heylen, D. K. (2017). Automatic recognition of touch gestures in the corpus of social touch. *Journal on Multimodal User Interfaces*, **11**(1), 81–96.
- [69] Keenan, M., Dillenburger, K., Röttgers, H. R., Dounavi, K., Jónsdóttir, S. L., Moderato, P., Schenk, J., Virués-Ortega, J., Roll-Pettersson, L., and Martin, N. (2015). Autism and ABA: The gulf between North America and Europe. *Review Journal of Autism and Developmental Disorders*, **2**(2), 167–183.
- [70] Keshmiri, S., Shiomi, M., Sumioka, H., Minato, T., and Ishiguro, H. (2020). Gentle versus strong touch classification: Preliminary results, challenges, and potentials. *Sensors*, **20**(11), 3033.
- [71] Kim, E. S., Berkovits, L. D., Bernier, E. P., Leyzberg, D., Shic, F., Paul, R., and Scassellati, B. (2013). Social robots as embedded reinforcers of social behavior in children with autism. *Journal of Autism and Developmental Disorders*, **43**(5), 1038–1049.
- [72] Kim, Y.-M., Koo, S.-Y., Lim, J. G., and Kwon, D.-S. (2010). A robust online touch pattern recognition for dynamic human-robot interaction. *IEEE Transactions on Consumer Electronics*, **56**(3), 1979–1987.
- [73] Kirby, R., Forlizzi, J., and Simmons, R. (2010). Affective social robots. *Robotics and Autonomous Systems*, **58**(3), 322–332.
- [74] Kozima, H., Michalowski, M. P., and Nakagawa, C. (2009). Keepon. *International Journal of Social Robotics*, **1**(1), 3–18.
- [75] Krichmar, J. L. and Chou, T.-S. (2018). A tactile robot for developmental disorder therapy. In *Proceedings of Technology, Mind, and Society*, pages 1–6. ACM.
- [76] Kuchenbecker, K. J. (2018). Haptics and haptic interfaces. In *Encyclopedia of Robotics*. Springer.
- [77] Lee, H., Park, K., Kim, J., and Kuchenbecker, K. J. (2019). Internal array electrodes improve the spatial resolution of soft tactile sensors based on electrical resistance tomography. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 5411–5417, Montreal, Canada. Equal contribution by Hyosang Lee and Kyungseo Park.

- [78] Legal Information Institute (2020). 20 U.S. code § 1401 definitions. <https://www.law.cornell.edu/uscode/text/20/1401>.
- [79] Lehmann, H., Rojik, A., Friebe, K., and Hoffmann, M. (2023). Hey, robot! An investigation of getting robot’s attention through touch. In *Proceedings of the International Conference on Social Robotics (ICSR)*, pages 388–401. Springer.
- [80] Leite, I., Martinho, C., Pereira, A., and Paiva, A. (2009). As time goes by: long-term evaluation of social presence in robotic companions. In *Proceedings of the IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, pages 669–674, Toyama, Japan. IEEE.
- [81] Lim, A. and Okuno, H. G. (2014). The MEI robot: Towards using motherese to develop multimodal emotional intelligence. *IEEE Transactions on Autonomous Mental Development*, **6**(2), 126–138.
- [82] Logan, D. E., Breazeal, C., Goodwin, M. S., Jeong, S., O’Connell, B., Smith-Freedman, D., Heathers, J., and Weinstock, P. (2019). Social robots for hospitalized children. *Pediatrics*, **144**(1).
- [83] Losinski, M., Sanders, S. A., and Wiseman, N. M. (2016). Examining the use of deep touch pressure to improve the educational performance of students with disabilities: A meta-analysis. *Research and Practice for Persons with Severe Disabilities*, **41**(1), 3–18.
- [84] Lowe, R., Andreasson, R., Alenljung, B., Lund, A., and Billing, E. (2018). Designing for a wearable affective interface for the NAO robot: a study of emotion conveyance by touch. *Multimodal Technologies and Interaction*, **2**(1), 2.
- [85] Luo, S., Bimbo, J., Dahiya, R., and Liu, H. (2017). Robotic tactile perception of object properties: A review. *Mechatronics*, **48**, 54–67.
- [86] Marco, E. J., Hinkley, L. B., Hill, S. S., and Nagarajan, S. S. (2011). Sensory processing in autism: A review of neurophysiologic findings. *Pediatric Research*, **69**(8), 48–54.
- [87] Mastrogiuseppe, M., Capirci, O., Cuva, S., and Venuti, P. (2015). Gestural communication in children with autism spectrum disorders during mother–child interaction. *Autism*, **19**(4), 469–481.
- [88] Melia, M., Schmidt, M., Geissler, B., König, J., Krahn, U., Ottersbach, H. J., Letzel, S., and Muttray, A. (2015). Measuring mechanical pain: The refinement and standardization of pressure pain threshold measurements. *Behavior Research Methods*, **47**(1), 216–227.

- [89] Mey, S. C. (2017). Animal assisted therapy for children with autism. *International Journal of Child Development and Mental Health*, **5**(1), 29–42.
- [90] Milo, R., Jorgensen, P., Moran, U., Weber, G., and Springer, M. (2010). BioNumbers – the database of key numbers in molecular and cell biology. *Nucleic Acids Research*, **38**(suppl\_1), D750–D753. BNID 110800.
- [91] Miyata, N., Yamaguchi, K., and Maeda, Y. (2007). Measuring and modeling active maximum fingertip forces of a human index finger. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2156–2161.
- [92] Moshkina, L. (2012). Improving request compliance through robot affect. In *AAAI Conference on Artificial Intelligence*, pages 2031–2037, Toronto, Canada. AAAI Press.
- [93] Moshkina, L. and Arkin, R. C. (2005). Human perspective on affective robotic behavior: a longitudinal study. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1444–1451, Edmonton, Canada. IEEE.
- [94] Moshkina, L., Park, S., Arkin, R. C., Lee, J. K., and Jung, H. (2011). TAME: Time-varying affective response for humanoid robots. *International Journal of Social Robotics*, **3**(3), 207–221.
- [95] Neal, J., Bigby, L., and Nicholson, R. (2004). Occupational therapy, physical therapy, and orientation and mobility services in public schools. *Intervention in School and Clinic*, **39**(4), 218–222.
- [96] Neilson, P. D. (1972). Speed of response or bandwidth of voluntary system controlling elbow position in intact man. *Medical and Biological Engineering*, **10**(4), 450–459.
- [97] O’Haire, M. E., McKenzie, S. J., McCune, S., and Slaughter, V. (2014). Effects of classroom animal-assisted activities on social functioning in children with autism spectrum disorder. *The Journal of Alternative and Complementary Medicine*, **20**(3), 162–168.
- [98] Ojha, S., Vitale, J., and Williams, M.-A. (2021). Computational emotion models: a thematic review. *International Journal of Social Robotics*, **13**(6), 1253–1279.
- [99] O’Haire, M. E. (2013). Animal-assisted intervention for autism spectrum disorder: A systematic literature review. *Journal of Autism and Developmental Disorders*, **43**(7), 1606–1622.

- [100] Pakkar, R., Clabaugh, C., Lee, R., Deng, E., and Matarić, M. J. (2019). Designing a socially assistive robot for long-term in-home use for children with autism spectrum disorders. In *Proceedings of the IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, pages 1–7, New Delhi, India. IEEE.
- [101] Peca, A., Simut, R., Pintea, S., Costescu, C., and Vanderborght, B. (2014). How do typically developing children and children with autism perceive different social robots? *Computers in Human Behavior*, **41**, 268–277.
- [102] Peranich, L., Reynolds, K. B., O’Brien, S., Bosch, J., and Cranfill, T. (2010). The roles of occupational therapy, physical therapy, and speech/language pathology in primary care. *The Journal for Nurse Practitioners*, **6**(1), 36–43.
- [103] PPS UK Limited (2021). SingleTact | easy to use, accurate, miniature force sensors. <https://www.singletact.com/force-sensors/>.
- [104] Pyo, S., Lee, J., Kim, W., Jo, E., and Kim, J. (2019). Multi-layered, hierarchical fabric-based tactile sensors with high sensitivity and linearity in ultrawide pressure range. *Advanced Functional Materials*, page 1902484.
- [105] Robins, B. and Dautenhahn, K. (2014). Tactile interactions with a humanoid robot: Novel play scenario implementations with children with autism. *International Journal of Social Robotics*, **6**(3), 397–415.
- [106] Robins, B., Otero, N., Ferrari, E., and Dautenhahn, K. (2007). Eliciting requirements for a robotic toy for children with autism – results from user panels. In *Proceedings of the IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, pages 101–106. IEEE.
- [107] Robins, B., Amirabdollahian, F., Ji, Z., and Dautenhahn, K. (2010). Tactile interaction with a humanoid robot for children with autism: A case study analysis involving user requirements and results of an initial implementation. In *Proceedings of the IEEE International Symposium in Robot and Human Interactive Communication (RO-MAN)*, pages 704–711. IEEE.
- [108] Romano, J. M. and Kuchenbecker, K. J. (2011). Please touch the robot. *IEEE/RJS Conference on Intelligent Robots and Systems (IROS)*. Hands-on demonstration presented at IEEE/RJS Conference on Intelligent Robots and Systems (IROS).
- [109] Russell, J. A. (1980). A circumplex model of affect. *Journal of Personality and Social Psychology*, **39**(6), 1161.
- [110] Saito, T., Shibata, T., Wada, K., and Tanie, K. (2003). Relationship between interaction with the mental commit robot and change of stress reaction of the elderly. In *Proceedings of the IEEE International Symposium on Computational Intelligence in Robotics and Automation*, volume 1, pages 119–124, Kobe, Japan. IEEE.

- 
- [111] Scassellati, B., Boccanfuso, L., Huang, C.-M., Mademtzi, M., Qin, M., Salomons, N., Ventola, P., and Shic, F. (2018). Improving social skills in children with ASD using a long-term, in-home social robot. *Science Robotics*, **3**(21), 1–9.
- [112] Schmitz, A., Maiolino, P., Maggiali, M., Natale, L., Cannata, G., and Metta, G. (2011). Methods and technologies for the implementation of large-scale robot tactile sensors. *IEEE Transactions on Robotics*, **27**(3), 389–400.
- [113] Schneider, M. and Adamy, J. (2014). Towards modelling affect and emotions in autonomous agents with recurrent fuzzy systems. In *IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, pages 31–38, San Diego, USA. IEEE.
- [114] Sefidgar, Y. S., MacLean, K. E., Yohanan, S., Van der Loos, H. M., Croft, E. A., and Garland, E. J. (2015). Design and evaluation of a touch-centered calming interaction with a social robot. *IEEE Transactions on Affective Computing*, **7**(2), 108–121.
- [115] Shamsuddin, S., Yussof, H., Ismail, L., Hanapiah, F. A., Mohamed, S., Piah, H. A., and Zahari, N. I. (2012). Initial response of autistic children in human-robot interaction therapy with humanoid robot nao. In *Proceedings of the IEEE International Colloquium on Signal Processing and its Applications*, pages 188–193.
- [116] Shibata, T. and Wada, K. (2011). Robot therapy: A new approach for mental healthcare of the elderly – a mini-review. *Gerontology*, **57**(4), 378–386.
- [117] Shiomi, M., Minato, T., and Ishiguro, H. (2017). Subtle reaction and response time effects in human-robot touch interaction. In *Proceedings of the International Conference on Social Robotics (ICSR)*, pages 242–251. Springer.
- [118] Silvera Tawil, D., Rye, D., and Velonaki, M. (2012). Interpretation of the modality of touch on an artificial arm covered with an EIT-based sensitive skin. *The International Journal of Robotics Research*, **31**(13), 1627–1641.
- [119] Song, Y., Huang, W., Mu, C., Chen, X., Zhang, Q., Ran, A., Peng, Z., Sun, R., and Xie, W. (2019). Carbon nanotube-modified fabric for wearable smart electronic-skin with exclusive normal-tangential force sensing ability. *Advanced Materials Technologies*, **4**(5), 1800680.
- [120] Stanton, C. M., Kahn, P. H., Severson, R. L., Ruckert, J. H., and Gill, B. T. (2008). Robotic animals might aid in the social development of children with autism. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 271–278.
- [121] Stiehl, W. D. and Breazeal, C. (2005). Affective touch for robotic companions. In *Proceedings of the International Conference on Affective Computing and Intelligent Interaction*, pages 747–754. Springer.

- [122] Stock-Homburg, R. (2022a). Survey of emotions in human–robot interactions: Perspectives from robotic psychology on 20 years of research. *International Journal of Social Robotics*, **14**, 389–411.
- [123] Stock-Homburg, R. (2022b). Survey of emotions in human–robot interactions: Perspectives from robotic psychology on 20 years of research. *International Journal of Social Robotics*, **14**, 389–411.
- [124] Suzuki, R. and Lee, J. (2016). Robot-play therapy for improving prosocial behaviours in children with autism spectrum disorders. In *Proceedings of the IEEE International Symposium on Micro-NanoMechatronics and Human Science (MHS)*, pages 1–5.
- [125] Tan, P. and Tong, L. (2008). Integrated and discontinuous piezoelectric sensor/actuator for delamination detection. In *Delamination Behaviour of Composites*, pages 141–168. Elsevier.
- [126] Tanaka, F., Cicourel, A., and Movellan, J. R. (2007). Socialization between toddlers and robots at an early childhood education center. *Proceedings of the National Academy of Sciences*, **104**(46), 17954–17958.
- [127] Taneja, P., Olausson, H., Trulsson, M., Svensson, P., and Baad-Hansen, L. (2019). Defining pleasant touch stimuli: A systematic review and meta-analysis. *Psychological Research*, pages 1–16.
- [128] Tapus, A., Peca, A., Aly, A., Pop, C., Jisa, L., Pintea, S., Rusu, A. S., and David, D. O. (2012). Children with autism social engagement in interaction with NAO, an imitative robot: A series of single case experiments. *Interaction Studies*, **13**(3), 315–347.
- [129] Teyssier, M., Parilusyan, B., Roudaut, A., and Steimle, J. (2021). Human-like artificial skin sensor for physical human-robot interaction. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 3626–3633. IEEE.
- [130] Tielman, M., Neerincx, M., Meyer, J.-J., and Looije, R. (2014). Adaptive emotional expression in robot-child interaction. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 407–414, Bielefeld, Germany. IEEE.
- [131] Van Erp, J. B. and Toet, A. (2015). Social touch in human–computer interaction. *Frontiers in Digital Humanities*, **2**, 2.
- [132] Wada, K. and Shibata, T. (2007). Social effects of robot therapy in a care house—change of social network of the residents for two months. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 1250–1255.

- [133] Watling, R., Deitz, J., Kanny, E. M., and McLaughlin, J. F. (1999). Current practice of occupational therapy for children with autism. *American Journal of Occupational Therapy*, **53**(5), 498–505.
- [134] Westlund, J. M. K., Park, H. W., Williams, R., and Breazeal, C. (2018). Measuring young children’s long-term relationships with social robots. In *Proceedings of the ACM Conference on Interaction Design and Children (IDC)*, pages 207–218.
- [135] Wood, L. J., Zaraki, A., Robins, B., and Dautenhahn, K. (2019). Developing Kaspar: A humanoid robot for children with autism. *International Journal of Social Robotics*, pages 1–18.
- [136] Yohanan, S. and MacLean, K. E. (2012). The role of affective touch in human-robot interaction: Human intent and expectations in touching the haptic creature. *International Journal of Social Robotics*, **4**(2), 163–180.
- [137] Yoshikai, T., Fukushima, H., Hayashi, M., and Inaba, M. (2009). Development of soft stretchable knit sensor for humanoids’ whole-body tactile sensibility. In *Proceedings of the IEEE/RAS International Conference on Humanoid Robots*, pages 624–631.