Socially Intelligent Robots
that Understand and Respond to Human Touch

Merel Jung

SOCIALLY INTELLIGENT ROBOTS THAT UNDERSTAND AND RESPOND TO HUMAN TOUCH

Merel Jung

GRADUATION COMMITTEE

Chairman and Secretary:
   Prof. dr. P.M.G. Apers
Supervisor:
   Prof. dr. D.K.J. Heylen
Co-supervisor:
   Dr. M. Poel
Committee Members:
   Prof. dr. K.E. MacLean
   Dr. M. Ammi
   Prof. dr. ir. B.J.A. Kröse
   Prof. dr. ir. M.C. van der Voort
   Prof. dr. J.B.F. van Erp

SOCIALLY INTELLIGENT ROBOTS THAT
UNDERSTAND AND RESPOND TO HUMAN TOUCH

DISSERTATION

to obtain
the degree of doctor at the University of Twente,
on the authority of the rector magnificus,
prof. dr. T.T.M. Palstra,
on account of the decision of the graduation committee,
to be publicly defended
on Wednesday June 28, 2017 at 16:45

by

Merel Madeleine Jung
born on May 28, 1987
in Rotterdam, the Netherlands

This dissertation has been approved by:

   Prof. dr. D.K.J. Heylen, University of Twente, NL (supervisor)
   Dr. M. Poel, University of Twente, NL (co-supervisor)

SUMMARY

Touch is an important nonverbal form of interpersonal interaction which is used to communicate emotions and other social messages. As interactions with social robots are likely to become more common in the near future these robots should also be able to engage in tactile interaction with humans. Therefore, the aim of the research presented in this dissertation is to work towards socially intelligent robots that can understand and respond to human touch. To become a socially intelligent actor a robot must be able to *sense*, *classify* and *interpret* human touch and *respond* to this in an appropriate manner. To this end we present work that addresses different parts of this interaction cycle.

After the introduction in Part I of the dissertation, we have taken a data-driven approach in Part II. We have focused on the *sense* and *classify* steps of the interaction cycle to automatically recognize social touch gestures such as pat, stroke and tickle from pressure sensor data.

In Chapter 2 we present CoST: Corpus of Social Touch, a dataset containing 7805 captures of 14 different social touch gestures. All touch gestures were performed in three variants: gentle, normal and rough on a pressure sensitive mannequin arm. Recognition of these 14 gesture classes using various classifiers yielded accuracies of up to 60%; moreover, gentle gestures proved to be harder to classify than normal and rough gestures. We further investigated how different classifiers, interpersonal differences, gesture confusions and gesture variants affected the recognition accuracy.

In Chapter 3 we describe the outcome of a machine learning challenge on touch gesture recognition. This challenge was extended to the research community working on multimodal interaction with the goal of sparking interest in the touch modality and to promote exploration of the use of data processing techniques from other more mature modalities for touch recognition. Two datasets were made available containing labeled pressure sensor data of social touch gestures: the CoST dataset presented in Chapter 2 and the Human-Animal Affective Robot Touch (HAART) gesture set. The most important outcomes of the challenges were: (1) transferring techniques from other modalities, such as image processing, speech, and human action recognition provided valuable feature sets; (2) gesture classification confusions were similar despite the various data processing methods that were used.

In Part III of the dissertation we present three studies on the use of social touch in interaction with robot pets. We have mainly focused on the *interpret* and *respond* steps of the interaction cycle to identify which touch gestures a robot pet should understand, how touch can be interpreted within a social context and in which ways a robot can respond to human touch.

In Chapter 4 we present a study of which the aim was to gain more insight into the factors that are relevant to interpret the meaning of touch within a social context. We elicited touch behaviors by letting participants interact with a robot pet companion in different affective scenarios. In a contextualized lab setting, participants acted as if they were coming home in different emotional states (i.e., stressed, depressed, relaxed and excited) without being given specific instructions on the kinds of behaviors that they should display. Based on video footage of the interactions and interviews we explored the use of touch behaviors, the expressed social messages and the expected robot pet responses. Results show that emotional state influenced the social messages that were communicated to the robot pet as well as the expected responses. Furthermore, it was found that multimodal cues were used to communicate with the robot pet, that is, participants often talked to the robot pet while touching it and making eye contact. Additionally, the findings of this study indicate that the categorization of touch behaviors into discrete touch gesture categories based on dictionary definitions is not a suitable approach to capture the complex nature of touch behaviors in less controlled settings.

In Chapter 5 we describe a study in which we evaluated the expressive potential of breathing behaviors for 1-DOF zoomorphic robots. We investigated the extent to which researcher-designed emotional breathing behaviors could communicate four different affective states. Additionally, we were interested in the influence of robot form on the interpretation of these breathing behaviors. For this reason two distinct robot forms were compared: a rigid wood-based form resembling a rib cage called 'RibBit' and a flexible, plastic-based form resembling a ball of fur called 'Flexi-Bit'. In the study, participants rated for each robot how well the different breathing behaviors reflected each of four affective states: stressed, depressed, relaxed and excited. The results show that both robot forms were able to express high and low arousal states through breathing behavior, whereas valence could not be expressed reliably. Low arousal states could be communicated by low frequency breathing behavior and higher frequency breathing conveyed high arousal. In contrast, context might play a more important role in the interpretation of different levels of valence. Unexpectedly, robot form did not influence the perception of the behavior

that was expressed. These findings can help to inform future design of affective behavior for robot pet companions.

In Chapter 6 we present a study in which we explored in what ways people with dementia could benefit from interaction with a robot pet companion with more advanced touch recognition capabilities and which touch gestures would be important in their interaction with such a robot. In addition, we explored which other target groups might benefit from robot pets with more advanced interaction capabilities. We administered a questionnaire and conducted interviews with two groups of health care providers who all worked in a geriatric psychiatry department. One group had experience with robotic seal Paro while the other group had no experience with the use of robot pets. The results show that health care providers perceived Paro as an effective intervention to improve the well-being of people with dementia. Furthermore, the care providers indicated that people with dementia (would) use mostly positive forms of touch and speech to interact with Paro. Paro's auditory responses were criticized because they can over-stimulate the patients. Additionally, the care providers argued that social interactions with Paro are currently limited and therefore the robot does not meet the needs of a broader audience such as healthy elderly people that still live in their own homes. The development of robot pets with more advanced social capabilities such as touch and speech recognition might result in more intelligent interactions which could help to better adapt to the needs of people with dementia and could make interactions more interesting for a broader audience. Moreover, the robot's response modalities and its appearance should match the needs of to the target group.

To conclude, the contributions of this dissertation are the following. We have made a touch gesture dataset available to the research community and have presented benchmark results. Furthermore, we have sparked interest into the new field of social touch recognition by organizing a machine learning challenge and have pinpointed directions for further research. Also, we have exposed potential difficulties for the recognition of social touch in more naturalistic settings. Moreover, the findings presented in this dissertation can help to inform the design of a behavioral model for robot pet companions that can understand and respond to human touch. Additionally, we have focused on the requirements for tactile interaction with robot pets for health care applications.

SAMENVATTING

Aanraking is een belangrijke vorm van non-verbale intermenselijke interactie die gebruikt wordt om emoties en andere sociale boodschappen te communiceren. Omdat interactie met sociale robots in de nabije toekomst hoogstwaarschijnlijk meer gebruikelijk zal worden moeten deze robots kunnen omgaan met aanraking tijdens interacties met mensen. Daarom is het doel van het onderzoek, dat in dit proefschrift wordt gepresenteerd, om toe te werken naar sociaal intelligente robots die in staat zijn om menselijke aanraking te begrijpen en erop te kunnen reageren. Om op een sociaal intelligente manier te kunnen acteren moet een robot in staat zijn om menselijk aanraking te kunnen *waarnemen*, *classificeren* en *interpreteren* en hier op een gepaste manier op kunnen *reageren*. De verschillende onderdelen van deze interactie cyclus zullen aan bod komen in dit proefschrift.

Na de introductie in Deel I van het proefschrift, hanteren we een datagedreven benadering in Deel II. We hebben daarbij de focus gelegd op de stappen *waarnemen* en *classificeren* uit de interactie cyclus voor het automatisch herkennen van verschillende soorten aanrakingen zoals aaien, kietelen en het geven van een klopje op basis van sensor data.

In Hoofdstuk 2 presenteren we een data verzameling van sociale aanraking genaamd CoST: 'Corpus of Social Touch'. Deze data verzameling bevat 7805 voorbeelden van 14 verschillende soorten sociale aanrakingen. Alle aanrakingen zijn uitgevoerd op een voor aanraking gevoelige paspop arm in drie intensiteiten: zacht, gemiddeld en ruig. Deze 14 verschillende soorten aanrakingen werden onderscheiden van elkaar door middel van verschillende classificatie methodes wat resulteerde in accuratesses van maximaal 60%. Daarbij bleken zachte aanrakingen moeilijker te onderscheiden dan gemiddelde en ruigere aanrakingen. Verder hebben we de invloed van verschillende soorten classificatie methodes, interpersoonlijke verschillen, verwarringen tussen aanrakingen en de verschillende intensiteiten op de mate waarin de aanrakingen herkend konden worden onderzocht.

In Hoofdstuk 3 beschrijven we de uitkomst van een *machine learning challenge* die we hebben georganiseerd. Hiervoor hebben we onderzoekers uit het veld van multimodale interactie uitgedaagd om verschillende aanrakingen te herkennen door middel van *machine learning* technieken. Het doel van deze uitdaging was om meer aandacht te genereren voor onder-

zoek op het gebied van aanraking en om te exploreren of data verwer-
kingstechnieken die nu gebruikt worden voor het herkennen van andere
modaliteiten ook toe te passen zijn voor het herkennen van aanrakingen.
Twee data verzamelingen met gelabelde druk sensor data van verschil-
lende sociale aanrakingen zijn beschikbaar gesteld aan de deelnemers:
de CoST data verzameling die gepresenteerd is in Hoofdstuk 2 en de
'Human-Animal Affective Robot Touch' (HAART) data verzameling. De
belangrijkste uitkomsten waren dat: (1) gebruikelijke technieken voor de
herkenning van beeld, spraak en menselijke activiteiten ook kunnen wor-
den ingezet voor het herkennen van aanraking; (2) verwarringen tussen
aanrakingen vergelijkbaar waren ondanks de verschillende data bewer-
kingstechnieken die waren gebruikt.

In Deel III van het proefschrift presenteren we drie studies op het ge-
bied van aanraking in interactie met robot dieren. De focus ligt hierbij
voornamelijk op de stappen *interpreteren* en *reageren* uit de interactie cy-
clus om te onderzoeken welke aanrakingen een robot dier zou moeten
kunnen begrijpen, hoe aanraking geïnterpreteerd kan worden in een so-
ciale context en op welke manieren een robot kan reageren op menselijke
aanraking.

In Hoofdstuk 4 presenteren we een studie waarvan het doel is om meer
inzicht te krijgen in de factoren die relevant zijn voor het interpreteren
van de betekenis van aanraking in een sociale context. We hebben aan-
rakingsgedrag uitgelokt door participanten te laten interacteren met een
robot dier in verschillende emotioneel geladen scenario's. In een in het lab
nagebouwde woonkamer omgeving hebben participanten gedaan alsof ze
thuis kwamen in verschillende emotionele stemmingen (dat wil zeggen:
gestrest, neerslachtig, ontspannen en enthousiast) zonder specifieke in-
structies over wat voor gedrag ze moeten vertonen. We hebben het gebruik
van aanrakingen, de uitgedrukte sociale boodschappen en de verwachte
reactie van de robot onderzocht op basis van video opnames en inter-
views. De resultaten laten zien dat de emotionele stemming van invloed
was op zowel de sociale boodschap die werd gecommuniceerd naar het
robot dier als op de verwachte reactie. Daarnaast bleek dat participan-
ten gebruik maakte van multimodale signalen om te communiceren met
het robot dier, dat wil zeggen, deelnemers praatten vaak tegen het robot
dier terwijl ze deze aanraakten en oogcontact maakten. Bovendien duiden
de bevindingen van deze studie erop dat het categoriseren van aanra-
kingen in discrete categorieën op basis van woordenboek definities niet
een geschikte benadering lijkt voor het beschrijven van de complexe aard
van aanrakingen in een minder gecontroleerde omgeving.

In Hoofdstuk 5 beschrijven we een studie waarin we de expressieve mogelijkheden van ademhalingspatronen evalueren voor dierachtige robots met 1 vrijheidsgraad. We onderzoeken in hoeverre deze robots vier verschillende emotionele stemmingen kunnen communiceren door middel van door onderzoekers ontwikkelde ademhalingspatronen. Daarnaast waren we geïnteresseerd in de invloed die de vorm van een robot dier heeft op de interpretatie van de ademhalingspatronen. Om deze reden hebben we twee verschillende robot dieren vergeleken: een rigide robot dier van hout dat op een ribbenkast lijkt genaamd 'RibBit' en een flexibel robot dier van plastic dat op een balletje met vacht lijkt genaamd 'Flexi-Bit'. In de studie beoordeelden participanten voor elke robot in hoeverre de verschillende ademhalingspatronen elke emotionele stemming representeerde: gestrest, neerslachtig, ontspannen en enthousiast. De resultaten lieten zien dat beide robots in staat waren om een laag en hoog activatie niveau over te brengen door middel van ademhaling terwijl valentie niet betrouwbaar kon worden gecommuniceerd. Een staat van lage activatie kan worden gecommuniceerd door middel van laag frequente ademhaling en hoog frequente ademhaling kan een staat van hoge activatie overbrengen. Daarentegen speelt context waarschijnlijk een belangrijkere rol in het interpreteren van verschillende niveaus van valentie. In tegenstelling tot onze verwachting bleek dat de vorm van de robot geen invloed had op de perceptie van de ademhalingspatronen. Deze bevindingen kunnen bijdragen aan het ontwerp van affectieve gedragingen voor toekomstige robot dieren.

In Hoofdstuk 6 presenteren we een studie waarin we onderzoeken op welke manier mensen met dementie kunnen profiteren van interactie met een robot dier met meer geavanceerde mogelijkheden op het gebied van aanraking en welke aanrakingen belangrijk zijn in hun interactie met een robot dier. Daarnaast onderzoeken we welke andere doelgroepen nog meer profijt kunnen hebben van robot dieren met meer geavanceerde interactie mogelijkheden. Voor dit onderzoek hebben we vragenlijsten en interviews afgenomen bij twee groepen verzorgers die allen werkzaam waren op een psychogeriatrische afdeling. Een groep had ervaring met robot zeehond Paro terwijl de andere groep geen ervaring had met het werken met robot dieren. De resultaten laten zien dat de verzorgers Paro als een effectieve interventie zien om het welzijn van mensen met dementie te bevorderen. Daarnaast geven de verzorgers aan dat mensen met dementie voornamelijk positieve aanrakingen en spraak (zouden) gebruiken in hun interactie met Paro. Er werd kritiek geuit op de auditieve reacties van Paro omdat deze tot overstimulatie kunnen leiden bij de patiënten. Bovendien beargumenteerden de verzorgers dat de sociale interacties met Paro

nu beperkt zijn en dat de robot in zijn huidige staat daarom ongeschikt is voor een breder publiek zoals gezonde ouderen die nog zelfstandig wonen. De ontwikkeling van robot dieren met meer geavanceerde sociale mogelijkheden zoals aanraking en spraak herkenning kan resulteren in intelligentere interacties die beter kunnen aansluiten bij de behoefte van mensen met dementie en die van een breder publiek. Daarnaast is het belangrijk om het uiterlijk en de reactie mogelijkheden van een robot dier af te stemmen op de doelgroep.

Ter afsluiting, de bijdragen van dit proefschrift zijn de volgende. We hebben een data verzameling met verschillende aanrakingen beschikbaar gemaakt voor onderzoek en hebben *benchmark* resultaten gepresenteerd. Daarnaast hebben we aandacht gegenereerd voor het nieuwe veld van aanraking herkenning door middel van het organiseren van een *machine learning challenge* en hebben we richtingen aangegeven voor verder onderzoek. Ook hebben we potentiële problemen bij het herkennen van sociale aanraking in meer natuurgetrouwe omgevingen aan het licht gesteld. Tevens kunnen de bevindingen die in dit proefschrift zijn gepresenteerd helpen bij het ontwerp van gedragsmodellen voor robot dieren die menselijk aanraking kunnen begrijpen en gepast kunnen reageren. Bovendien hebben we ons ook gefocust op de benodigdheden voor tactiele interactie met robot dieren voor toepassingen in de gezondheidszorg.

ACKNOWLEDGMENTS

It was in the second year of my bachelor's studies that I heard about PhD programs. At the time, it sounded like something that I would want to do after getting my master's degree. As I really enjoyed carrying out research for both my bachelor's and master's theses I became even more sure that I wanted to pursue a PhD. While I was finishing up my master's thesis I contacted Dirk to talk about the possibilities of doing a PhD at the Human Media Interaction department. After some months I got offered a PhD position on affective touch, which resulted in this dissertation about 4.5 years later. Thank you Dirk for giving me the opportunity to work on this interesting topic.

Although, doing PhD research can feel like a lonely journey at times, a lot of people have supported me over the years. First of all, I would like to thank my supervisors. Betsy, thank you for helping me to define my PhD research. I really enjoyed our meetings and I'm thankful for your support. Thank you Mannes for immediately agreeing to take over the supervision from Betsy. I really appreciate everything you have taught me about machine learning and your feedback on my experiments and papers has helped me to become a better researcher. Also, thank you for your understanding and the moral support when I was stressed out about upcoming deadlines. Thank you Ronald for also stepping in as one of my supervisors. You have helped me a lot with setting up my first data collection and writing my first papers. Dirk, thank you for your feedback when I was setting up a new experiment and your comments on my papers. Also, thank you for giving me the freedom to give direction to my research.

Gijs and Christian, I'm glad that we were able to shared some of the weird moments that seemed to be inevitable when working on the topic of social touch. Also, Saskia and Lisa thank you for the nice collaboration.

I have also received a lot of support from other HMI staff members. Thank you Dennis and Mariët for giving me advice regarding the supervision of students. Also, thank you Dennis for lending your expertise on video annotation. Boris, thank you for the talks we have had about machine learning and my dataset, you enthusiasm for the field is really inspiring. Rieks, thank you for the time you took to teach me about mathematical concepts and machine learning techniques. Thank you Charlotte, Alice and Wies for all your help throughout the years. Lynn, thank you for reading my work and improving my English writing.

I would also like to thank all my (part-time) office mates for making me feel at home at HMI. Jeroen, I'm glad to have shared an office with you for 4 years. I really appreciate your down to earth personality. I will remember our discussions about one of our most important first world problems: having not enough time to play all the awesome games that are available while still finishing our PhD research within an acceptable time frame. Although, it is not a definitive solution to the aforementioned problem, thank you for taking the initiative to start the tradition of the HMI gaming nights. Cristina, thank you for making our office a more cheerful place with your little robots and kawaii stuff. Please take good care of the Furbies for me. Khiet, one of our part-time office mates, often showing up with a cup of coffee on one hand and a cup of tea in the other, thank you for our morning talks. I really appreciate all your support throughout the years and the nice dinners we have had. Jan, fellow internet citizen, always stopping by our office around 11am, thank you for spreading happiness with your contagious laughter and kazoo music and for teaming up with me during fussball matches.

I also want to thank all my colleagues at HMI for making the department such a nice place to work and for all the fun times we have had hanging out after work. Thank you Khiet, Randy and Jeroen for the semi-healthy?? lunches we have enjoyed together. Special thanks to Randy for driving us to the various famous American restaurants that Enschede has to offer. Also, thank you Jelte for all the fun we have had playing co-op games. I have really enjoyed my time at HMI: having lunch together, playing fussball, playing video and tabletop games, going to conferences/ summer schools together and building things at thunder Thursdays. Thank you all: Jeroen, Cristina, Khiet, Randy, Jered, Daniel, Daphne, Gijs, Christian, Michiel, Roelof, Jelte, Robby, Merijn, Lorenzo, Bob, Jan, Jamy, Jaebok, Vicky, Aduén, Alejandro C., Alejandro M. and everyone else I forgot.

I would also like to thank Karon for giving me the opportunity to visit her lab. It has been a pleasure to work with you. I really enjoyed working together on the CuddleBits project, thank you Laura, Paul, Jussi, Oliver and the others from the SPIN lab. Also, Laura I have enjoyed organizing the touch challenge together and thank you for all the nice dinners we have had during my stay in Vancouver.

Ik wil ook graag mijn moeder bedanken voor al haar steun. Ik vind het fijn dat je altijd in mij bent blijven geloven. Verder wil ik ook Bert en Emelien bedanken voor al hun steun in de afgelopen jaren. Also, thank you Ivor for all the fun movie nights and everything that you have taught me over the years.

And last, but definitely not least, thank you Tiemo for you unconditional support throughout the years. Thank you for being there for me during stressful times and for always believing in me. I have had a hard time being away from you when traveling for work but we have also enjoyed nice holidays together when you came to visit me afterwards. I have fond memories of our first road trip together. I really appreciate your endless curiosity, thank you for everything that you have taught me.

Merel Jung
Enschede, June 2017

CONTENTS

Part I

AUTOMATIC UNDERSTANDING OF HUMAN
TOUCH: INTRODUCTION AND MOTIVATION

1

INTRODUCTION

1.1 TOUCH IN SOCIAL INTERACTION

People express themselves through social signals in the form of verbal and nonverbal behaviors. Touch is one of the important nonverbal forms of social interaction as are visual cues such as facial expressions, gaze, body posture and air gestures [116]. However, compared to vision and audition (as in vocal cues), interpersonal touch does not generally receive much research attention yet [40, 50]. Similarly, the touch modality is often overlooked in human-computer interaction such as remote communication and in interactions with embodied or virtual agents [112]. As interactions with social robots are likely to become more common in the near future these robots are expected to engage in tactile interaction with humans [112]. Therefore the aim of the research presented in this dissertation is to work towards socially intelligent robots that can understand and respond to human touch.

Touch behavior is seen in many different forms of social interaction: a handshake as a greeting, a high-five to celebrate a joint accomplishment, a tap on the shoulder to gain someone's attention, a comforting hug from a friend, or holding hands with a romantic partner. In contrast to functional touch, which can be used to explore our environment and manipulate objects such as tools, Haans and IJsselsteijn described social touch as all instances of interpersonal touch, whether this is accidental (e.g. bumping into someone on the street) or conscious (e.g. hugging someone who is upset) [46]. In interpersonal interaction touch is important for establishing and maintaining social interaction [40]. Touch can be used to generate and communicate both positive and negative emotions [48, 50] as well as to express intimacy [3], power and status [40]. Furthermore, there is research that indicates that a brief touch can result in a more positive evaluation of the toucher [36] and can increase the willingness to comply with a request such as filling out a questionnaire [44]. Additionally, the positive effects of touch on well-being are extensively described in the literature [34, 80]. For

example, a five-day touch intervention was found to significantly reduce anxiety in intensive care patients compared to standard treatment (i.e., a rest hour) [47].

On the physiological level the human skin serves an important function as a sense organ for discriminating different tactile sensations such as whether a surface is smooth or rough [58]. Apart from the discriminative function of touch, the human sense of touch also plays an important role in affective experiences. Caress-like stroking touches have been found to selectively activate specific receptors called C-Tactile (CT) afferents in the hairy skin, which respond particularly strongly to stroking at a velocity of about 3 cm/s [1, 78, 82]. Strokes at this velocity also result in the highest subjective pleasantness ratings [77, 78]. Moreover, on the cortical level, nerves related to discriminative touch mostly activate the somatosensory cortex, whereas the CT-afferent nerves mainly activate areas that are involved in affective processing (i.e., the posterior insular cortex and the orbitofrontal cortex) [79, 82]. Interestingly, third person observations of stroking touches in a social setting have been shown to result in similar pleasantness ratings and similar brain activation in the posterior insula as experienced touch [81, 82, 120]. However, these pleasantness ratings of stroking touches have been found to be sensitive to top-down social cues such as the gender of the toucher [42]. These findings indicate that there are specialized pathways for both experienced and observed social touch interactions [82].

The lack of research on social touch can be in part explained by its private nature which makes it more difficult to gather data during natural interactions [50]. In order to study touch behavior, researchers have to rely on different strategies. Common methods are self-reports (e.g. questionnaires or dairy studies), observations and controlled experiments [110]. Additionally, touch is a complex modality: the sense of touch is the combined effort of input from different receptors which register touch (e.g. pressure, vibration and skin stretch), pain, temperature and limb proprioception [50, 68]. Moreover, there are many types of touch (e.g. stroke, hit and tickle) and the social context (e.g. concurrent verbal and nonverbal behavior, the type of interpersonal relationship and the situation in which the touch takes place) influences how these different types of touch should be interpreted [50, 51, 59, 107]. The complexity of interpersonal touch along with technical difficulties make it challenging to transfer the touch modality to remote interaction and human-robot interaction [40, 46, 112].

1.2 SOCIAL TOUCH IN HUMAN-COMPUTER INTERACTION

When moving from interpersonal touch to social touch in human-computer interaction one of the challenges is for a computer to understand and respond to human touch [112]. Additionally, social actors such as robots and virtual agents should be able to simulate social touches [17, 54, 112]. The development of social agents that can engage in social touch interaction is part of the larger research area aimed at automatic understanding of social behavior that is *Social Signal Processing (SSP)* [115] and development of artificial social intelligence that is the field of *affective computing* [88]. In these fields social behavior is currently mostly studied in the form of vocal behavior using speech/ audio analysis and nonverbal behaviors including facial expressions, body postures and air gestures the detection of which can be automated with the help of computer vision [109, 115]. Sensors such as microphones and cameras have been found to be able to capture social signals that can be interpreted through machine learning techniques and statistical analysis [115]. As touch is also important in social interaction we will focus specifically on the touch modality to enable robots to automatically understand and respond to human touch.
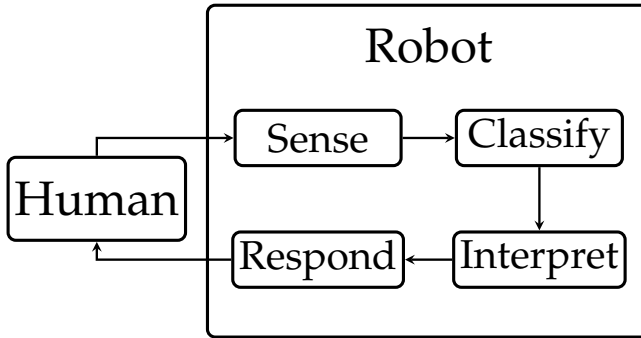
Figure 1: Steps in the interaction cycle for a socially intelligent robot that can understand and respond to human touch.

Extending social touch interaction to include interaction with social agents can result in more natural interaction, providing opportunities for various applications. For example, the addition of tactile interaction can benefit robot therapy in which robots are used to comfort people in stressful environments, for instance, children in hospitals [56] and elderly people in nursing homes [117]. Furthermore, the addition of haptic technology to a training scenario involving a virtual patient could help medical students to learn how to use social touch appropriately in a health-care

setting [74, 75]. However, just equipping a robot or interface with touch sensors to mimic the human somatosensory system is not enough. To become a socially intelligent actor a robot should be able to *sense*, *classify* and *interpret* human touch and *respond* to this in a socially appropriate manner (see Figure 1). The model in Figure 1 is based on the traditional *Sense-Think-Act cycle* for intelligent agent behavior from the field of artificial intelligence ([94], p. 51). In this dissertation we have broken down the 'think' step of the traditional model into two steps, namely 'classify' and 'interpret'. Similar models have been used in the touch literature before, all with a slightly different focus [103, 124]. The model proposed by Yohanan and MacLean [124] focuses on the recognition and expression steps of the interaction cycle on both the robot and the human side whereas the model used by Silvera-Tawil et al. [103] focuses mostly on the recognition and interpretation of social touch by a robot. The work presented in this dissertation will contribute to all the steps in the interaction cycle as illustrated in Figure 1.

1.3   MAIN CONTRIBUTIONS

The main contributions of the research reflected in this dissertation are the following:

A publicly available dataset of social touch gestures (Chapter 2)
To the best of our knowledge there were no publicly available datasets on social touch which are necessary for research and benchmarking. First, we give a systematic overview of the characteristics of available studies on the sensing and recognition of social touch up to August 2015. Second, we present a corpus of social touch gestures which is called Corpus of Social Touch (CoST). Third, we compare the performance of different classifiers to provide a baseline for touch gesture recognition within CoST and evaluate the factors that influence the recognition accuracy.

Moving forward the new field of social touch recognition (Chapter 3)
As the recognition of touch behavior has received far less research attention than recognition of behaviors in the visual and auditory modalities, we aimed to spark interest into this relatively new field by organizing a machine learning challenge. Researchers with expertise in other sensory modalities were able to try out their processing techniques on two touch datasets which included CoST. In this dissertation we present the outcome of this undertaking and pinpoint further research directions.

A first step towards the automatic understanding of social touch for naturalistic human-robot interaction (Chapter 4)
Current studies in the domain of social touch for human-robot interaction focused mainly on highly controlled settings in which users were requested to perform different touch behaviors, one at a time, according to predefined labels. However, as context is important for the interpretation of touch behavior we explore the use of touch during interactions with a robot pet in a scenario in which participants acted as it they were coming home in different emotional states. No specific instructions were given to the participants on the kinds of behaviors that they should display. In this dissertation we reflect on the challenges of segmentation and labeling of touch behaviors in a less controlled setting.

Informing the design of a behavioral model for robot pet companions that can understand and respond to human touch (Chapters 4 and 5)
In a contextualized lab setting, participants acted as if they were coming home in different emotional states (i.e., stressed, depressed, relaxed and excited) without being given specific instructions on the kinds of behaviors that they should display. We explore the use of touch and other social behaviors, the expressed social messages and the expected robot pet responses.

In addition, we explore a haptic response in the form of a simulated breathing mechanism for one degree of freedom (1-DOF) robot pets which are collectively called the 'CuddleBits' [21]. Contrary to previous studies we focus specifically on breathing behavior and explore the expressive space of various breathing patterns. In this dissertation we evaluate whether 1-DOF robot movements can communicate different valence and arousal states. Furthermore, we investigate the influence of robot materiality on the interpretation of the affective robot behaviors.

Requirements for tactile interaction with robot pets for health care applications (Chapter 6)
Robot pet companions such as robotic seal Paro are increasingly used in care for the elderly due to the positive effects that interaction with these robots can have on the well-being of patients with dementia. As touch is one of the most important interaction modalities for patients with dementia this can be a natural way to interact with these robots. However, currently commercially available companion robots do not focus specifically on touch interaction, which seems like a missed opportunity. In this dissertation we explore in what ways people with dementia could benefit from interaction with a robot pet companion with more advanced touch

recognition capabilities and which touch gestures would be important in their interaction with such a robot. In addition, we explore which other target groups might benefit from robot pets with more advanced interaction capabilities.

1.4   OUTLINE OF THIS DISSERTATION

This dissertation consist of 4 parts. We have introduced the field of social touch and motivated the need to enable social agents to understand and respond to human touch in Part I. In Part II we will take a data-driven approach. We will focus on the *sense* and *classify* steps of Figure 1 to automatically recognize social touch gestures such as pat, stroke and tickle from pressure sensor data. In Chapter 2 we will present the Corpus of Social Touch (CoST) and discuss the performance results of several classifiers for the recognition of the touch gestures in CoST. Then, we will describe the outcome of a machine learning challenge on touch gesture recognition which was hosted in conjunction with the 2015 ACM International Conference on Multimodal Interaction (ICMI) in Chapter 3. In Part III we will study social touch within the context of human-robot interaction. We will mainly focus on the *interpret* and *respond* steps of Figure 1 to identify which touch gestures a robot pet should understand, how touch can be understood within a social context and ways in which a robot can respond to human touch. We will present work towards the interpretation of social touch in a more naturalistic setting in Chapter 4. Next, in Chapter 5 we will describe a study on the design of affective behavior for robot pets. Then, we will present a study in which we explore the benefits of robot pet companion with more advanced touch interaction capabilities for health care applications in Chapter 6. Finally, we will reflect on the work presented in this dissertation in Part IV. In Chapter 7 conclusions will be drawn based on the findings presented in this dissertation and we will provide directions for further research.

Part II

SENSING AND RECOGNIZING SOCIAL TOUCH GESTURES

The focus of this section will be on the use of sensors to register human touch and the use of machine learning techniques to automatically recognize different touch gestures from the sensor data. Firstly, we will present the Corpus of Social Touch (CoST) and the touch gesture recognition results for this dataset. Secondly, we will present the protocol and the findings from a machine learning challenge to recognize social touch gestures.

2

AUTOMATIC RECOGNITION OF TOUCH GESTURES IN
THE CORPUS OF SOCIAL TOUCH

The following chapter[1] covers research which was carried out by Merel
Jung under the supervision of Mannes Poel, Ronald Poppe and Dirk
Heylen. The content of this chapter is identical to that of the published
paper with some minor textual adaptations to embed the content into this
dissertation. The future work described in the paper has been moved to
Chapter 7 of the dissertation.

To understand human touch a robot needs sensors to register these
touches. Next, machine learning algorithms can be trained to automati-
cally distinguish between different types of touch. The focus of this chap-
ter and Chapter 3 will be on the recognition of touch gestures with social
meaning that are performed by hand on a pressure-sensitive surface; we
call these 'social touch gestures'. In this chapter we will present the Cor-
pus of Social Touch (CoST) and the performance results of several classi-
fiers for the recognition of the touch gestures in this dataset.

2.1 INTRODUCTION

Touch gestures can be used in social interaction to communicate and ex-
press different emotions [50, 48]. For example, love can be communicated
by hugging and stroking while anger can be expressed by pushing and
shaking [48]. Socially intelligent robots should be able to automatically
detect and recognize touch gestures in order to respond appropriately.

Equipping a robot with touch sensors is the first step towards touch
interaction based on human touch input. Once the sensor registers the
touch, we need to recognize the type of touch and interpret its meaning.

---

1 Based on Jung, M. M., Poel, M., Poppe, R., and Heylen, D. K. J., Automatic recognition of
touch gestures in the corpus of social touch, Journal on Multimodal User Interfaces, vol. 11,
no. 1, pp. 81–96, 2016.

Moreover, a robust touch recognition system should be perceived as working in real time and should be participant independent to avoid training sessions for new users. Some promising attempts have been made to recognize different sets of touch gestures (e.g. stroke, poke, and hit) recorded on various interfaces. However, as recognition rates vary depending on the degree of similarity between the touch gestures it is difficult to judge the relative strengths of one approach over the other.

To work towards reliable touch gesture recognition we recorded a corpus of social touch hand gestures to characterize various touch gestures. We will focus on the recognition of a list of relevant social touch gestures. The interpretation of the social meaning of these touch gestures is beyond the scope of this chapter. To the best of our knowledge there are no publicly available datasets on social touch for research and benchmarking. The contribution of this chapter is three-fold: first, we will give a systematic overview of the characteristics of available studies on the recognition of social touch; second, we will present the Corpus of Social Touch (CoST); third, we will compare the performance of different classifiers to provide a baseline for touch gesture recognition within CoST and evaluate the factors that influence the recognition accuracy.

The remainder of the chapter is organized as follows: in the next section we will discuss related work on the recognition of social touch, in Section 2.3 we will describe the CoST dataset. Next, touch gesture recognition results will be presented and discussed in Section 2.4 and Section 2.5, respectively. The chapter will conclude in Section 2.6.

2.2   RELATED WORK ON SOCIAL TOUCH RECOGNITION

There have been a number of studies on social touch recognition. We will briefly discuss the different characteristics of these studies. A summary of previous studies is presented in Table 1. Please note that we have only considered the studies that reported details on classification and studies published up to August 2015.

2.2.1   *Touch surface and sensors*

In these studies, touch was performed on various surfaces such as robots (e.g. [71]), sensor sheets (e.g. [85]) or human body parts such as arms [102]. Physical appearances of interfaces for touch interaction included robotic animals (e.g. [124]), full body humanoid robots (e.g. [71]), partial embodiments such as a mannequin arm (e.g. [103]) and a balloon interface [83]. Several techniques were used for the sensing of touch, each having its own

Table 1: Results of literature on social touch recognition

Paper	Touch surface	Sensor(s)	Touch recognition of...	n	Classifier	Design	Accuracy
Altun and MacLean [2]	Haptic Creature	force sensing resistors, accelerometer	26 gestures	31	random forest	between-subjects	33%
Altun and MacLean [2]	Haptic Creature	force sensing resistors, accelerometer	9 emotions	31	random forest	between-subjects	36%
Altun and MacLean [2]	Haptic Creature	force sensing resistors, accelerometer	9 emotions	31	random forest	within-subjects	48%
Altun and MacLean [2]	Haptic Creature	force sensing resistors, accelerometer	9 emotions using gesture recog.	31	random forest	between-subjects	36%
Bailenson et al. [6]	force-feedback joystick	2d accelerometer	7 emotions	16	classification by human	1 subject rates 1 other	33%
Bailenson et al. [6]	force-feedback joystick	2d accelerometer	7 emotions	16	SVM[a] RBF[b] kernel	between-subjects	36%
Bailenson et al. [6]	other subject's hand	/	7 emotions	16	classification by human	1 subject rates 1 other	51%
Chang et al. [25]	Haptic Creature	force sensing resistors	4 gestures	1	custom recognition software	real-time	up to 77%
Cooney et al. [26]	Sponge (humanoid) robot	accelerometer, gyro sensor	13 full-body gestures	21	SVM[a] RBF[b] kernel	between-subjects	77%
Cooney et al. [27]	humanoid robot 'mock-up'	photo-interrupters	20 full-body gestures	17	k-NN[c]	between-subjects	63%
Cooney et al. [27]	humanoid robot 'mock-up'	photo-interrupters	20 full-body gestures	17	SVM[a] RBF[b] kernel	between-subjects	72%
Cooney et al. [27]	humanoid robot 'mock-up'	Microsoft Kinect	20 full-body gestures	17	k-NN[c]	between-subjects	67%
Cooney et al. [27]	humanoid robot 'mock-up'	Microsoft Kinect	20 full-body gestures	17	SVM[a] RBF[b] kernel	between-subjects	78%
Cooney et al. [27]	humanoid robot 'mock-up'	photo-interrupters, Microsoft Kinect	20 full-body gestures	17	k-NN[c]	between-subjects	82%
Cooney et al. [27]	humanoid robot 'mock-up'	photo-interrupters, Microsoft Kinect	20 full-body gestures	17	SVM[a] RBF[b] kernel	between-subjects	91%
Flagg et al. [37]	furry lap pet	conductive fur sensor, piezoresistive fabric pressure sensors	9 gestures	16	neural network	between-subjects	75%
Flagg et al. [37]	furry lap pet	conductive fur sensor, piezoresistive fabric pressure sensors	9 gestures	16	logistic regression	between-subjects	72%
Flagg et al. [37]	furry lap pet	conductive fur sensor, piezoresistive fabric pressure sensors	9 gestures	16	Bayes network	between-subjects	68%
Flagg et al. [37]	furry lap pet	conductive fur sensor, piezoresistive fabric pressure sensors	9 gestures	16	random forest	between-subjects	86%
Flagg et al. [37]	furry lap pet	conductive fur sensor, piezoresistive fabric pressure sensors	9 gestures	16	random forest	within-subjects	94%
Flagg et al. [38]	fur sensor	conductive fur sensor	3 gestures	7	linear regression	between-subjects	82%
Ji et al. [57]	KASPAR (hand section)	capacitive pressure sensors	4 gestures	1	SVM[a] intersection kernel	within-subject	up to 96%
Ji et al. [57]	KASPAR (hand section)	capacitive pressure sensors	4 gestures	1	SVM[a] RBF[b] kernel	within-subject	up to 93%

Results of literature on social touch recognition (cont.)

Paper	Touch surface	Sensor(s)	Touch recognition of...	n	Classifier	Design	Accuracy
Jung [60]	mannequin arm	piezoresistive fabric pressure sensors	14 gestures	31	Bayesian classifier	subject-independent	53%
Jung [60]	mannequin arm	piezoresistive fabric pressure sensors	14 gestures	31	SVM[a] linear kernel	subject-independent	46%
Jung et al. [65]	mannequin arm	piezoresistive fabric pressure sensors	14 rough gestures	31	Bayesian classifier	subject-independent	54%
Jung et al. [65]	mannequin arm	piezoresistive fabric pressure sensors	14 rough gestures	31	SVM[a] linear kernel	subject-independent	53%
Kim et al. [71]	KaMERo	charge-transfer touch sensors, accelerometer	4 gestures	12	temporal decision tree	real-time	83%
Knight et al. [73]	sensate bear	electric field sensor, capacitive sensors	4 gestures	11	Bayesian networks + k-NN[c]	real-time	20-100%
Nakajima et al. [83]	Emoballoon	barometric pressure sensor, microphone	6 gestures + 'no touch'	9	SVM[a] RBF[b] kernel	between-subjects	75%
Nakajima et al. [83]	Emoballoon	barometric pressure sensor, microphone	6 gestures + 'no touch'	9	SVM[a] RBF[b] kernel	within-subjects	84%
Naya et al. [85]	sensor sheet	pressure-sensitive conductive ink	5 gestures	11	k-NN[c] + Fisher's linear discriminant	between-subjects	87%
Silvera-Tawil et al. [101]	sensor sheet	pressure sensing based on EIT[d]	6 gestures	1	logitboost algorithm	within-subject	91%
Silvera-Tawil et al. [101]	sensor sheet	pressure sensing based on EIT[d]	6 gestures	35	logitboost algorithm	between-subjects	74%
Silvera-Tawil et al. [101]	experimenter's back	/	6 gestures	35	classification by human	between-subjects	86%
Silvera-Tawil et al. [102]	mannequin arm	pressure sensing based on EIT[d], force sensor	8 gestures + 'no touch'	2	logitboost algorithm	within-subjects	88%
Silvera-Tawil et al. [102]	experimenter's arm	/	8 gestures	2	classification by human	within-subjects	75%
Silvera-Tawil et al. [102]	mannequin arm	pressure sensing based on EIT[d], force sensor	8 gestures + 'no touch'	40	logitboost algorithm	subject-independent	71%
Silvera-Tawil et al. [102]	other subject's arm	/	8 gestures	40	classification by human	1 subject rates 1 other	90%
Silvera-Tawil et al. [103]	mannequin arm	pressure sensing based on EIT[d], force sensor	6 emotions + 'no touch'	2	logitboost algorithm	within-subjects	88%
Silvera-Tawil et al. [103]	mannequin arm	pressure sensing based on EIT[d], force sensor	6 social messages + 'no touch'	2	logitboost algorithm	within-subjects	84%
Silvera-Tawil et al. [103]	mannequin arm	pressure sensing based on EIT[d], force sensor	6 emotions + 'no touch'	2	logitboost algorithm	between-subjects	32%
Silvera-Tawil et al. [103]	mannequin arm	pressure sensing based on EIT[d], force sensor	6 social messages + 'no touch'	2	logitboost algorithm	between-subjects	51%
Silvera-Tawil et al. [103]	mannequin arm	pressure sensing based on EIT[d], force sensor	6 emotions + 'no touch'	42	logitboost algorithm	subject-independent	47%
Silvera-Tawil et al. [103]	other subject's arm	/	6 emotions	42	classification by human	1 subject rates 1 other	52%

Results of literature on social touch recognition (cont.)

Paper	Touch surface	Sensor(s)	Touch recognition of...	n	Classifier	Design	Accuracy
Silvera-Tawil et al. [103]	mannequin arm	pressure sensing based on EIT[d], force sensor	6 social messages + 'no touch'	42	logitboost algorithm	subject-independent	50%
Silvera-Tawil et al. [103]	other subject's arm	/	6 social messages	42	classification by human	1 subject rates 1 other	62%
Stiehl et al. [106]	The Huggable (arm section)	electric field sensor, force sensors, thermistors	8 gestures (disregarding 'slap')	1	neural network	within-subject	79%
van Wingerden et al. [113]	mannequin arm	piezoresistive fabric pressure sensors	14 rough gestures	31	neural network	between-subjects	64%

[a]SVM = Support Vector Machine, [b]RBF = Radial Basis Function, [c]k-NN = k-Nearest Neighbor, [d]EIT = Electrical Impedance Tomography

advantages and drawbacks for example, low cost vs. large hysteresis in force sensing resistors [29]. These sensing techniques were implemented in the form of artificial robot skins (e.g. [102]) or by following a modular approach using sensor tiles (e.g. [57]) or individual sensors to cover the robot's body (e.g. [25]). Designing an artificial skin entails extra requirements such as flexibility and stretchability to cover curved surfaces and moving joints [101, 104] but has the advantage of providing equal sensor density for detection across the entire surface which can be hard to achieve using individual sensors [25]. The approach of using computer vision to register touch is noteworthy [27].

2.2.2   *Touch recognition*

Previous research on the recognition of touch has included hand gestures (e.g. stroke [65]), full body gestures (e.g. hug [26]), emotions (e.g. happiness [103]), and social messages (e.g. affection [103]). Data was gathered from a single subject to test a proof of concept (e.g. [25]) or from multiple subjects to allow for the training of a subject independent model (e.g. [103]). Classification results show that it is harder to recognize emotions or social messages than the touch itself. This can be explained by the nontrivial nature of mapping touch to an emotional state or an intention for example, a single touch gesture can be used to communicate various emotions [48, 124]. Also, as expected, results of a within-subjects design were better than classification between-subjects (e.g. [2]) meaning that there was a larger inter-person variance than intra-person variance. Human classification of touch out-performed automatic classification (e.g. [101]). However, when touch was mediated by technology, human performance decreased. Bailenson et al. [6] found that emotions were better recognized by participants when performing a real handshake with another person compared to when the handshake with the other person was mediated through a force-feedback joystick. Classification was mostly off-line however, some promising attempts have been made with real-time classification, which is a prerequisite for real-time touch interaction (e.g. [71]). Real-time systems come with extra requirements such as gesture segmentation and ensuring adequate processing speed. Combining computer vision with touch sensing yielded better touch recognition results than relying on a single modality [27].

Direct comparison of touch recognition between studies based on reported accuracies is difficult because of differences in the number and nature of touch classes, sensors, and classification protocols. Furthermore, some reported accuracies were the result of a best-case scenario intending

to be a proof of concept (e.g. [25]). Some studies focused on the location of the touch rather than the touch gesture, such as distinguishing between 'head-pat' and 'foot-rub' [73]. While information on body location can enhance touch recognition, Silvera-Tawil et al. showed that comparable accuracies can be achieved by limiting the touch location to a single arm [102].

2.3 COST: CORPUS OF SOCIAL TOUCH

To address the need for social touch datasets, we recorded a corpus of social touch gestures (CoST) which was introduced in [65]. This dataset is publicly available [63].
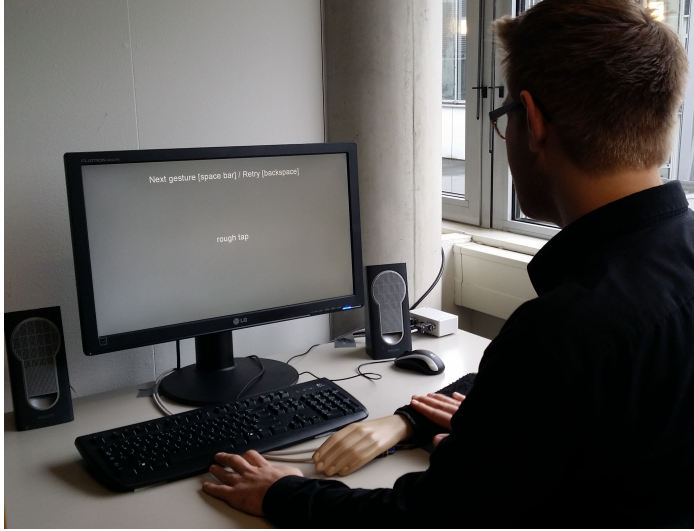
Figure 2: Participant performing the instructed touch gesture on the pressure sensor (the black fabric) wrapped around the mannequin arm

2.3.1 *Touch gestures*

CoST consists of the pressure sensor data of 14 different touch gestures performed on a sensor grid wrapped around a mannequin arm (see Figure 2). The touch gestures (see Table 2) included in the data collection were chosen from a touch dictionary composed by [124] based on the literature on touch interaction between humans and between humans and animals. The list of gestures was adapted to suit interaction with a mannequin arm.

Touch gestures involving physical movement of the arm itself, such as lift, push and swing, were omitted because the movement of the mannequin arm could not be sensed by the pressure sensors. All touch gestures were performed in three variants: gentle, normal and rough to increase the variety of ways a gesture could be performed by each individual.

Table 2: Touch dictionary, adapted from Yohanan and MacLean [124]

Gesture label	Gesture definition
Grab	Grasp or seize the arm suddenly and roughly.
Hit	Deliver a forcible blow to the arm with either a closed fist or the side or back of your hand.
Massage	Rub or knead the arm with your hands.
Pat	Gently and quickly touch the arm with the flat of your hand.
Pinch	Tightly and sharply grip the arm between your fingers and thumb.
Poke	Jab or prod the arm with your finger.
Press	Exert a steady force on the arm with your flattened fingers or hand.
Rub	Move your hand repeatedly back and forth on the arm with firm pressure.
Scratch	Rub the arm with your fingernails.
Slap	Quickly and sharply strike the arm with your open hand.
Squeeze	Firmly press the arm between your fingers or both hands.
Stroke	Move your hand with gentle pressure over arm, often repeatedly.
Tap	Strike the arm with a quick light blow or blows using one or more fingers.
Tickle	Touch the arm with light finger movements.

2.3.2   *Pressure sensor grid*

For the sensing of the gestures, an 8×8 pressure sensor grid (PW088-8x8/HIGHDYN from *plug-and-wear*[2], see Figure 3) was connected to a Teensy 3.0 USB Development Board (by *PJRC*[3]). The sensor was made of textile consisting of five layers. The two outer layers were protective lay-

---

2 www.plugandwear.com
3 www.pjrc.com
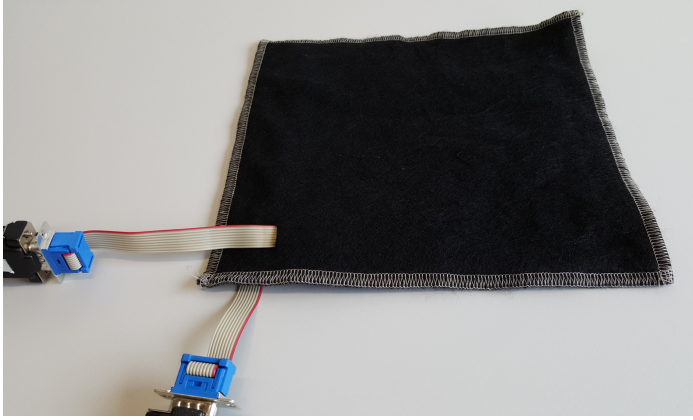
Figure 3: 8×8 pressure sensor grid

ers made of felt. Each outer layer was attached to a layer containing eight strips of conductive fabric separated by non-conductive strips. Between the two conductive layers was the middle layer which comprised a sheet of piezoresistive material. The conductive layers were positioned orthogonally so that they formed an 8 by 8 matrix. The sensor area was 160×160 mm with a thickness of 4 mm and a spatial resolution of 20 mm.

One of the conductive layers was attached to the power supply while the other was attached to the A/D converter of the Teensy board. After A/D conversion, the sensor values of the 64 channels ranged from 0 to 1,023 (i.e., 10 bits). Figure 4 displays the relationship between the sensor values and the pressure in kg/cm² for both the whole range (0-1,023) and the range used in the data collection (0-990). Pressure used during human touch interaction typically ranges from 30 g/cm² to 1,000 g/cm² [104], which corresponds to sensor values between 25 and 800. From the plots it can be seen that the sensor's resolution is accurate within this range but decreases at higher pressure levels. Sensor data was sampled at 135 Hz.

Our sensor meets the requirements set by Silvera-Tawil et al. [104] for optimal touch sensing in social human-robot interaction as the spatial resolution falls within the recommend range of 10-40 mm and the sample rate exceeds the required minimum (20 Hz). However, the human somatosensory system is more complex than this sensor as receptors in the skin register not only pressure but also pain and temperature and receptors in the muscles, joints and tendons register body motion [40, 104]. The sensor grid produces artifacts in the signal such as crosstalk, wear out and hysteresis (i.e., the influence of the previous and current input, which is discussed in Section 2.3.4). For demonstration purposes, we illustrated the
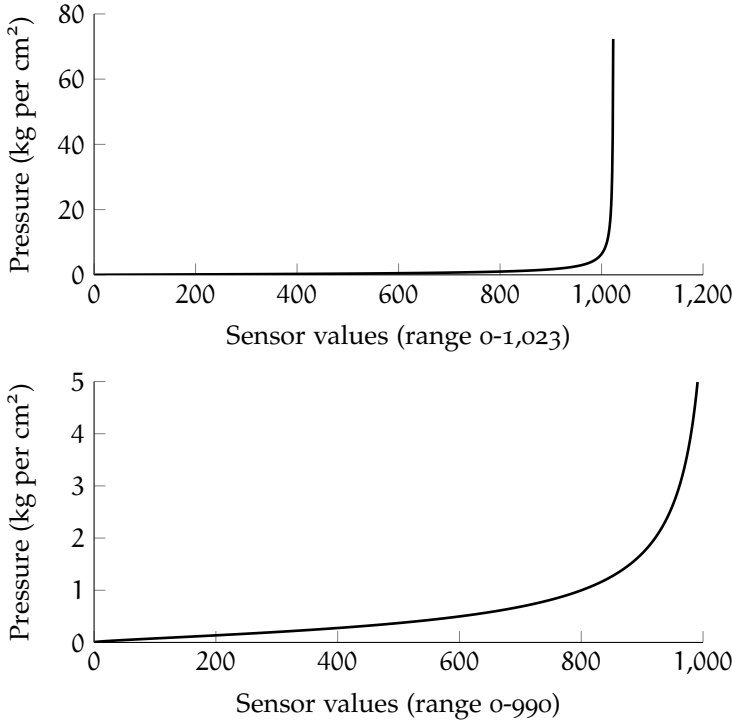
Figure 4: Plot of the relationship between the sensor output after A/D conversion and pressure in kg/cm² for both the whole range (top) and the range used (bottom)

sensor's crosstalk by pushing down with the end of a pencil perpendicular to the sensor grid to create a concentrated load (see Figure 5). The sensor was wrapped around the mannequin arm to create a setup similar to the one used for the data collection. We did not compensate for the artifacts in the data.

2.3.3   *Data acquisition*

2.3.3.1   *Setup*

The sensor was attached to the forearm of a full size rigid mannequin arm consisting of the left hand and the arm up to the shoulder (see Figure 2). The arm was chosen as the contact surface because this is one of the body locations that is often used to communicate emotions [48]. Also, the arm is one of the least invasive body areas on which to be touched [51] and pre-
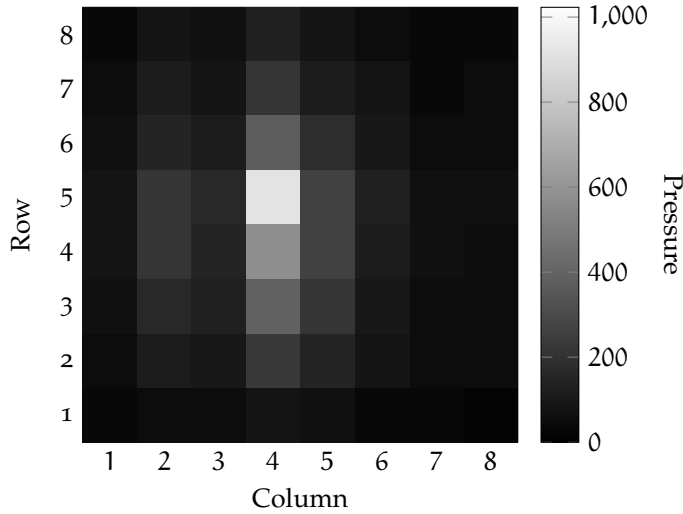
Figure 5: Crosstalk visualization showing the sensor data of a single frame, a pencil was pressed down on the sensor grid (light pressure point) effecting the pressure level of adjacent channels

sumably a neutral body location to touch others. The mannequin arm was fastened to the right side of the table to prevent it from slipping. Instructions for which gesture to perform had been scripted using *PsychoPy*[4] and were displayed to the participants on a computer monitor. Video recordings were made during the data collection as verification of the sensor data and the instructions given.

2.3.3.2 *Procedure*

Upon entering the experiment room, the participant was welcomed and was asked to read and sign an informed consent form. After filling in demographic information, the participant was provided with a written explanation of the data collection procedure. Participants were instructed to use their right hand to perform the touch gestures and use their left hand on the keyboard. Then an instruction video was shown of a person performing all 14 gestures on the mannequin arm based on the definitions from Table 2. Participants were instructed to repeat every gesture from the video to practice. No video examples were shown during the actual data collection. Next, example instructions were given to perform a stroke gesture in all three variants (i.e., gentle, normal and rough). After each

---

4 A module written for Python, see www.psychopy.org

gesture the participant could press the *spacebar* to continue to the next gesture or *backspace* to retry the current gesture. Once everything was clear to the participant the data collection started.

During the data collection each participant was prompted with 14 different touch gestures 6 times in 3 variants resulting in 252 gesture captures. In the instructions of the gesture to perform, the participants were shown only the gesture variant combined with the name of the gesture (e.g. 'gentle grab'), not the definition from Table 2. The order of instructions was pseudo-randomized into three blocks. Each instruction was given two times per block but the same instruction was not given twice in consecutive order. A single fixed list of instructions was constructed using these criteria. This list and the reversed order of the list were used as instructions in a counterbalanced design. After each block, there was a break and the participant was asked to report any difficulty in performing the instructions. Finally, participants were asked to describe the gestures and manners in their own words. The entire procedure took approximately 40 minutes for each participant.

2.3.3.3 *Participants*

A total of 32 people volunteered to participate in the data collection. Data of one participant was omitted due to technical difficulties. The remaining participants, 24 male and 7 female, all studied or worked at the University of Twente in the Netherlands. Most (26) had the Dutch nationality (1 British/Dutch), others were Ecuadorean, Egyptian, German (2x) and Italian. The age of the participants ranged from 21 to 62 years ($M = 34, SD = 12$) and 29 were right-handed.

2.3.4 *Data preprocessing*

The raw data was segmented into gesture captures based on the keystrokes of the participants marking the end of a gesture. Segmentation between keystrokes still contained many additional frames from before and after the gesture was performed. Removing these additional frames is especially important to reduce noise in the calculation of features that contain a time component, such as features that average over frames in time. See Figure 6 for an example of a gesture capture of 'normal tap' as segmented between keystrokes. Further segmentation is indicated by dashed lines. This plot also illustrates that the sensor values remain nonzero (the absolute minimum) when the sensor is not touched and that hysteresis occurs. In this case the sensor values are higher after the touch gesture is performed compared to before.
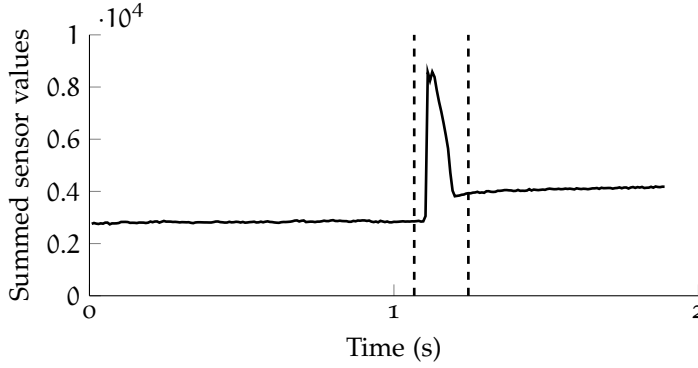
Figure 6: Gesture capture of 'normal tap' as segmented between keystrokes, further segmentation based on pressure difference is indicated by the dashed lines

Further segmentation of the gesture captures was based on the change in the gesture's intensity (i.e., the summed pressure over all 64 channels) over time using a sliding window approach. The first window starts at the beginning of the gesture capture and includes the number of frames corresponding to the window size parameter. The next window remains the same size but is shifted a number of frames corresponding to the step size parameter. The pressure intensity of each window is compared to that of the previous window. This procedure continues till the end of the gesture capture. Parameters (i.e., threshold of minimal pressure difference, step size, window size and offset) were optimized by visual inspection to ensure that all gestures were captured within the segmented part. The optimized parameters were fixed for all recordings.

After visual inspection it turned out that six gesture captures could not be automatically segmented because differences in pressure were too small (i.e., below the threshold parameter). The video recordings revealed that the gestures were either skipped or were performed too fast to be distinguishable from the sensor's noise. One other gesture capture was of notably longer duration (over a minute) than all other instances because the instructions were unclear at first. These seven gesture captures were instances of the variants 'gentle massage', 'gentle pat', 'gentle stroke', 'normal squeeze', 'normal tickle', 'rough rub', and 'rough stroke'. The instances of these gesture variants were removed from the dataset. The remaining dataset consists of 7,805 touch gesture captures in total: 2,601 gentle, 2,602 normal and 2,602 rough gesture captures.

Table 3: Mean and standard deviation (in parentheses) of the duration, mean and maximum pressure and contact area per touch variant and for all data

Variant	Gentle	Normal	Rough	All
Mean pressure (g/cm^2)	115 (61)	136 (82)	189 (157)	147 (112)
Max pressure (g/cm^2)	894 (511)	1,260 (629)	1,983 (813)	1,379 (802)
Contact area (% of sensor)	.21 (.16)	.22 (.18)	.26 (.21)	.23 (.19)
Duration (ms)	1,385 (1,303)	1,377 (1,257)	1,500 (1,351)	1,421 (1,305)

2.3.5 *Descriptive statistics*

To get an idea of the differences between touch gestures and the variants, descriptive statistics were calculated on three important characteristics of touch: intensity (g/cm²), contact area (% of sensor area) and gesture duration (ms). Pressure intensity was calculated as the mean pressure of all channels averaged over time and the maximum channel value of the gesture over all channels. Contact area was calculated for the frame with the highest summed pressure over all channels (corresponds to feature 21). Means and standard deviations of the touch data after segmentation are displayed for each variant and in total in Table 3 and per gesture in Table 4. It is notable that the mean and maximum pressure used per variant follow the expected pattern: gentle variants < normal variant < rough variants, indicating that participants used pressure to distinguish between the different variants. Figure 7 illustrates that there was a lot of overlap in duration between the different gestures (e.g. between hit and slap) and a lot of variance within each gesture, especially within massage and tickle. The tables and figure illustrate that the challenge of touch gesture recognition is complex and that it is not possible to distinguish between these different touch gestures using only these descriptive statistics. Table 5 shows the touch characteristics for males and females separately. Based on these characteristics there seems to be no significant differences between male and female touch gestures.

2.3.6 *Self reports*

In the self reports the most common difficulties (mentioned by at least 5 out of 31 participants) of distinguishing between gestures (disregarding the variants) were reported on pat vs. tap (12), grab vs. squeeze (10), rub vs. stroke (7), hit vs. slap (5) and pinch vs. squeeze (5). Furthermore, some combinations of gestures with variants were perceived as less logical. The most commonly mentioned gesture variants were: rough tickle (4), gentle

Table 4: Mean and standard deviation (in parentheses) of the duration (ms), mean and maximum pressure (g/cm^2) the contact area (% of sensor area) per touch gesture

Gesture	Grab	Hit	Massage	Pat	Pinch	Poke	Press
Mean pressure	349 (191)	101 (32)	172 (77)	100 (33)	126 (45)	95 (27)	188 (99)
Max pressure	1,774 (919)	1,643 (854)	1,621 (800)	1,057 (568)	1,701 (892)	1,258 (793)	1,660 (802)
Contact area	.59 (.17)	.15 (.05)	.36 (.20)	.15 (.07)	.12 (.08)	.08 (.08)	.23 (.16)
Duration	1,373 (715)	337 (403)	3,538 (1,898)	709 (753)	1,132 (597)	650 (502)	1,181 (608)

Gesture	Rub	Scratch	Slap	Squeeze	Stroke	Tap	Tickle
Mean pressure	131 (45)	106 (28)	95 (30)	286 (180)	116 (35)	92 (30)	96 (26)
Max pressure	1,282 (671)	1,064 (524)	1,165 (557)	1,980 (946)	1,135 (623)	1,055 (610)	911 (497)
Contact area	.21 (.10)	.17 (.08)	.15 (.06)	.47 (.24)	.20 (.08)	.12 (.07)	.18 (.09)
Duration	2,170 (1142)	2,205 (1,268)	321 (462)	1,502 (813)	1,722 (829)	564 (486)	2491 (1,446)
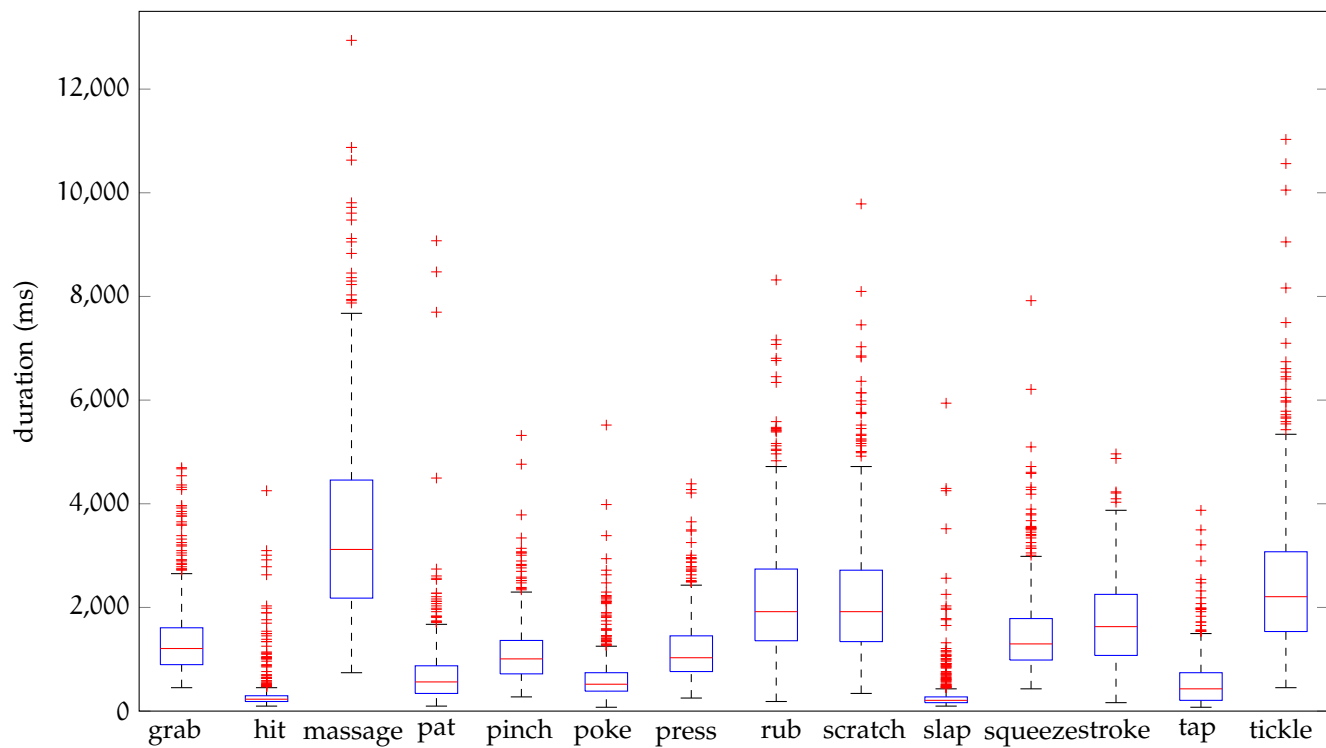
Figure 7: Boxplot of the duration (ms) for all 7,805 captures per touch gesture

hit (3) and gentle slap (3). Also, three participants raised concerns about breaking the setup when performing gestures too roughly.

At the end of the experiment participants were asked to provide their own descriptions. The most common keywords used to describe the gentle gesture variants were: soft (mentioned by 8 participants), slow (6), less force (6), less pressure (5), and light (3) while the rough variants were defined as: more force (12), hard (7), more pressure (4), and fast (3), energetic (3). 'Normal' was described as: the default/ regular (7), without thinking (4) and neutral (3).

Table 5: Mean and standard deviation (in parentheses) of the duration, mean and maximum pressure and contact area per touch variant and for all data for male and female subjects

Variant	Gentle	Normal	Rough	All
Male				
Mean pressure (g/cm^2)	117 (63)	137 (85)	193 (163)	149 (117)
Max pressure (g/cm^2)	885 (518)	1,245 (629)	1,981 (828)	1,370 (811)
Contact area (% of sensor)	.21 (.16)	.22 (.18)	.27 (.22)	.23 (.19)
Duration (ms)	1,358 (1,296)	1,349 (1,249)	1,491 (1,357)	1,399 (1,303)
Female				
Mean pressure (g/cm^2)	112 (50)	130 (72)	175 (133)	139 (96)
Max pressure (g/cm^2)	925 (485)	1,310 (624)	1,990 (763)	1,409 (772)
Contact area (% of sensor)	.20 (.15)	.21 (.17)	.24 (.20)	.21 (.17)
Duration (ms)	1,477 (1,325)	1,476 (1,281)	1,528 (1,330)	1,494 (1,312)

2.4 RECOGNITION OF SOCIAL TOUCH GESTURES

In this section we will present the performance results of several classifiers for the recognition of touch gestures in CoST. To establish the benchmark performance for CoST we compared the performance of four different commonly used classifiers. Two simple classifiers were chosen: a statistical model (Bayesian classifier) and a decision tree which allows for more insight into the classification process (e.g. which features are most important). Furthermore, we chose two more complex classifiers: a Support Vector Machine (SVM) which uses a single decision boundary and a neural network which allows for more complex decision boundaries.

2.4.1   *Feature extraction*

The dataset from the pressure sensor consists of a pressure value (i.e., the intensity) per channel (i.e., the location) at 135 fps. From the recorded sensor dataset, features were extracted for every gesture capture. The majority of features were based on the literature. The first features (1-28) were taken from previous work on this dataset [60, 65] which were based on social touch recognition literature, differences are indicated between parentheses. Features used for video classification can be applied to this dataset because the dataset of CoST is a grid of pressure values that are updated at a fixed rate which is similar to a low-resolution gray scale video. Features 29-43 were slight adaptations of the features used in [113] which were based on video classification literature. Feature numbers are indicated in parentheses.

- *Mean pressure* is the mean over channels and time (1).

- *Maximum pressure* is the maximum value over channels and time (2).

- *Pressure variability* is the mean over time of the sum over channels of the absolute value of difference between two consecutive frames (3).

- *Mean pressure per row* is the mean over columns and time resulting in one feature per row which are in the direction of the mannequin arm's length (from top to bottom, 4-11).

- *Mean pressure per column* is the mean over rows and time resulting in one feature per column which are in the direction of the mannequin arm's width (from left to right, 12-19).

- *Contact area* per frame is the fraction of channels with a value above 50% of the maximum value. Mean contact area is the mean over time of contact area (20) and the maximum pressure contact area is the contact area of the frame with the highest mean pressure over channels (21). The size of the contact area indicated whether the whole hand was used for a touch gesture, as would be expected for grab, or for example only one finger, as would be expected for a poke.

- *Temporal peak count* indicated how many times there was a significant increase in pressure level. That is, whether a touch gesture consisted of continuous touch contact as would be expected for grab or alternating pressure levels which would be expected for a tickle. One feature counts the number of frames for which the average pressure

of a frame was larger than that of its neighboring frames (22). (This feature replaced the previous version of feature 22 from [60, 65]). The other feature was calculated as the number of positive crossings of the threshold. The threshold was the mean over time of the pressure summed over all channels (23).

- *Traveled distance* (previously called 'displacement' in [60, 65]) indicated the amount of movement of the hand across the contact area. For example, for a squeeze less movement across the sensor grid would be expected than for a stroke. Center of mass (i.e., the average channel weighted by pressure) was used to calculate the movement on the contact surface in both the row and column directions. Two features were calculated in the row direction: the mean traveled distance of the center of mass over time (24) and the summed absolute difference of the center of mass over time (25). The same features were calculated for the column direction (26-27).

- *Duration* of the gesture measured in frames (28).

- *Pressure distribution* (previously called 'histogram-based features' in [113]) is the normalized histogram over all channels and time of the pressure values. The histogram contains eight bins equally spaced between 0 and 1,023 (29-36).

- *Spatial peaks* (previously called 'motion-based features' in [113]). – A spatial peak in a frame is a local maximum with a value higher than 0.75 of the maximum pressure (see feature 2). The following features were derived from the local peaks; the mean (37) and variance (38) over time of the number of spatial peaks per frame. Also, the mean over all spatial peaks and time of the distance of the spatial peak to the center of mass is a feature (39). The last feature based on spatial peaks is the mean over time and spatial peaks of the change in distance of each peak w.r.t. the center of mass (40).

- *Derivatives* were calculated as the mean absolute pressure differences within the rows and columns between frames. Features were derived from the mean over time and rows or columns of the above values (41-42). The mean absolute pressure difference for all channels was also calculated. The last feature was based on the mean over time and channels (43).

- *Variance* over channels and time (44).

- *Direction of movement* indicated the angle in which the center of mass was moving between frames. These angle values were divided into

quadrants of 90° each. For example, if the hand moves from the middle of the sensor grid to the upper right corner of the sensor grid, the center of mass moves at a 45° angle which falls within the upper right quadrant (i.e., the first quadrant). To deal with vectors that were close to the edge of two quadrants two points around the vector were evaluated, each weighting 0.5. A histogram represented the percentage of frames that fell into each quadrant (45-48).

- *Magnitude of movement* indicated the amount of movement of the center of mass. Statistics on the magnitude were calculated per gesture consisting of the mean, standard deviation, sum, and the range (49-52).

- *Periodicity* was the frequency with the highest amplitude in the frequency spectrum of the movement of the center of mass in the row and column direction, respectively (53-54).

2.4.2 *Classification experiments*

The extracted features were used for classification in MATLAB® (release 2013a). We performed two classification experiments: (1) classification of the touch gestures from the total dataset based on the gestures' class, thereby disregarding the variant (e.g. 'gentle grab' and 'normal grab' both belong to the same class: 'grab'); (2) classification of the touch gestures within each variant, splitting the dataset into 3 subsets: normal, gentle and rough. Due to their more pronounced nature, rough gesture variants were expected to have a more favorable signal-to-noise ratio compared to the softer variants.

For both classification experiments the dataset was split into a train/ validation set and test set using leave-one-subject-out cross-validation (31 folds) to train a user-independent model (i.e., dataset from each subject was only part of either the train set or the test set). Hyperparameters were optimized on the train/ validation set using leave-one-subject-out cross-validation (30 folds). Classification results were evaluated using the best performing hyperparameters found from the 30 folds (i.e., training/validation set only) to classify the test set. This procedure was repeated for all 31 folds. Note that each fold can have different optimized parameter values. The baseline of classifying a sample into the correct class based on random guessing is $1/14 \approx 7\%$ for both experiments. We will discuss details of each of the classifiers individually.

2.4.2.1 *Bayesian classifier*

The Gaussian Bayesian classifier has no hyperparameters to optimize. The mean and covariance for the features per class were calculated from the training dataset. These parameters for the multivariate normal distribution were used to calculate the posterior probability of a test sample belonging to the given class. Samples were assigned to the class with the maximum posterior probability.

2.4.2.2 *Decision tree*

Decision trees were trained using the CART learning algorithm with Gini's diversity index as splitting method. First a full tree was grown after which the tree was pruned. A parameter search for the optimal pruning level, using cross-validation as described above, was performed using a range of 5 to 30 in increments of 5.

2.4.2.3 *Support Vector Machine*

SVMs were trained using the LIBSVM software library [24], both with a linear kernel (hyperparameter C) and with a Radial Basis Function (RBF) kernel (hyperparameters C and γ). We chose to test two kernels due to their different approaches, the linear kernel separates the classes globally while the RBF kernel allows for a local division of two classes. The hyperparameters were optimized, using cross-validation as described above. A (grid) search was conducted for optimal parameters by growing the sequences of the parameter values exponentially ($C = 2^{-5}, 2^{-3}, ..., 2^{15}; \gamma = 2^{-15}, 2^{-13}, ..., 2^{3}$) as proposed by [52]. Before training, features were rescaled to the range of [0,1] by subtracting the minimum feature value from all feature values and dividing the result by the range of the feature values. Scaling prevents features with greater numeric ranges from dominating those with smaller numeric ranges [52].

2.4.2.4 *Neural Network*

A feedforward neural network was trained using Levenberg-Marquardt optimization. Stopping criteria were set to a maximum of 1,000 training iterations or 6 subsequent increases of the error on the validation set. The neural network toolbox in MATLAB automatically maps the range of the original input features to the range of [-1,1]. Because of memory constraints the architecture was set to two layers of 54 and 27 neurons, respectively to get results in a timely fashion. Leave-one-subject-out cross-validation was used, and, the dataset from the remaining 30 subjects was

split into a train set (70%) and a validation set (30%). The best performing network on the validation set of five runs was used to evaluate the test set (i.e., the samples of the left-out subject).

2.4.3   *Results*

Table 6 provides an overview of the overall accuracies for the whole dataset and per variant for different classifiers. Classification of 14 gesture classes independent of variants resulted in an overall accuracy of up to 60% using SVMs with the RBF kernel, which is more than 8 times higher than classification by random guessing ($\approx$ 7%). SVMs with the RBF kernel performed slightly better than the Bayesian classifiers, the SVMs with the linear kernel and the neural networks. Decision trees performed worse than the other classifiers.

Classification within each gesture variant showed that the accuracies for the rough variants (up to 62%) were higher than for the normal variants (up to 60%), which were higher than those for the gentle variants (up to 54%). The exception was the Bayesian classifier. In this case the normal variants performed slightly better than the rough variants. The SVM classifiers (both kernels) performed slightly better than the Bayesian classifier and neural network. Again, decision trees performed worse than the other classifiers.

Table 6: Overall accuracies of leave-one-subject-out cross-validation for the variants per classifier, standard deviations in parentheses

Variant	All	Normal	Gentle	Rough
Bayesian	.57(.11)	.59(.13)	.52(.14)	.58(.12)
Decision tree	.48(.10)	.49(.13)	.43(.10)	.52(.10)
SVM linear	.59(.11)	.60(.11)	.54(.13)	.62(.13)
SVM RBF	.60(.11)	.60(.11)	.54(.13)	.62(.12)
Neural network	.59(.12)	.58(.13)	.52(.13)	.59(.13)

2.5   DISCUSSION

In this section we will discuss the touch gesture recognition results in depth, looking into accuracy differences between subjects and between different classifiers, the interaction between gestures and the different vari-

ants and confusions between touch gestures. Also, we will reflect critically on the collection of the touch gesture data.

2.5.1 *Classification results and touch gesture confusion*

From the classification results in Table 6 it can be seen that the more complex classifiers (i.e., SVM and neural network) performed better than the simpler decision tree. However, the performance of the simpler Bayesian classifier was only slightly lower than those of the SVM and neural network. This indicates that recognition rates are reasonably robust across different classification methods. Moreover, the accuracy reported in this work (i.e., 60% for SVMs with the RBF kernel) was higher than the accuracy of 53% that was previously reported for the CoST dataset using Bayesian classifiers [60]. This indicates that the additional features and the use of more complex classification methods with hyperparameter optimization have improved the accuracy.

The subject independent model generalized well for some subjects but not for others as shown by the large individual differences in accuracy for the total dataset in Table 7. Differences in accuracy between subjects ranged from 44% for the Bayesian classifiers and the decision trees to 50% for the linear SVMs and neural networks. These individual differences make it harder to build a reliable subject independent model for touch gesture recognition. Depending on the application a trade-off can be made to build subject-dependent models which could increase accuracy at the expense of the need for training sessions. Between classifiers, results per subject differed on average 13%. These differences were largely due to the overall lower decision tree results, per subject the accuracies for the other four classifiers differed on average 6%. As expected, gentle gestures were considerably harder to classify which can be due to the lower pressure levels used for this gesture variant (see Table 3), resulting in a lower signal-to-noise ratio.

To gain insight into the interaction between gestures and their variants, we classified the gesture variants (i.e., 3 classes) and the combination of gestures and their variants (i.e., 42 classes) using the Bayesian classifier as a baseline due to its simplicity. Classification of the gesture variants using leave-one-subject-out cross-validation yielded accuracies ranging from 39% to 64% ($M = 50\%, SD = 6\%$). Over participants the correct rate for the classification of all gestures dependent on variant ranged from 15% to 47% ($M = 32\%, SD = 9\%$). Misclassification was most common between the gestures' variants which is in line with the low accuracy for the classification of the gesture variants. Confusions between gestures were similar

Table 7: Accuracy per participant for all data for the different classifiers. *Legend* – accuracy: ⩾ 50% , ⩾ 70%

Participant	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
Bayesian	.42	.58	.78	.70	.71	.45	.58	.66	.56	.75	.66	.54	.60	.37	.47	.41
Decision tree	.39	.48	.65	.63	.58	.43	.52	.55	.52	.67	.54	.47	.48	.30	.39	.29
SVM linear	.52	.58	.83	.73	.76	.52	.60	.65	.58	.78	.71	.63	.59	.41	.48	.44
SVM RBF	.51	.58	.82	.72	.76	.50	.61	.68	.56	.80	.72	.63	.62	.42	.42	.47
Neural network	.48	.59	.79	.74	.73	.47	.63	.66	.59	.82	.71	.61	.62	.44	.40	.47

Participant	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
Bayesian	.65	.58	.50	.51	.56	.58	.34	.66	.41	.69	.60	.54	.55	.58	.58
Decision tree	.62	.48	.45	.41	.50	.52	.23	.56	.36	.54	.48	.42	.55	.50	.51
SVM linear	.58	.65	.57	.52	.56	.67	.33	.65	.44	.70	.62	.44	.55	.63	.63
SVM RBF	.65	.65	.60	.52	.56	.66	.37	.65	.44	.72	.65	.46	.58	.65	.65
Neural network	.61	.63	.57	.49	.56	.65	.32	.67	.46	.74	.64	.49	.54	.63	.62

to those found for classification independent of gesture variant. For example: 'gentle grab' was correctly classified in 36% of the samples and was most often misclassified as 'normal grab' (24%), 'gentle squeeze' (16%), 'normal squeeze' (8%) and 'rough grab' (5%).

Misclassification was mostly due to confusions between similar touch gestures. Table 8 shows the confusion matrix for the SVM with the RFB kernel of the whole dataset as this classifier yielded the best results. The five most frequently confused gesture pairs were: grab and squeeze (sum of 294 confused samples); pat and tap (280); rub and stroke (223); scratch and tickle (219); hit and slap (154). Within gesture variants the rankings of most confused pairs were similar to those of the combined variants. Also, confusions between touch gestures depicted in Table 8 largely matched the touch gesture pairs that were reported to be difficult for the participants in Section 2.3.6. However, some small differences were observed: although 'pinch vs. squeeze' was in the top 5 most often reported difficulties in the confusion matrix this was not one of the most frequently confused gesture pairs (sum of 104 confused samples). Conversely, 'scratch vs. tickle' was one of the five most confused gesture pairs but was not among the most often mentioned difficulties (mentioned by 3 participants).

Recognizing a large set of different touch gestures can reduce the classification accuracy, especially when gestures show many overlapping characteristics. Therefore, it is important to find the right balance for each application. To illustrate this trade-off we composed a subset of gestures by starting with the original 14 gestures and removing one of the gestures for each of the five most commonly confused gesture pairs, the subset consisted of nine gestures: grab, massage, pinch, poke, press, slap, stroke, tap and tickle. Classification of this gesture subset independent of variant using a Bayesian classifier with leave-one-subject-out cross-validation yielded accuracies ranging from 45% to 94% ($M = 75\%, SD = 12\%$). The average performance increased by 18% for the recognition of nine touch gestures compared to the results with fourteen touch gestures using the same classifier. However, at the cost of the ability to distinguish between more classes.

To get an indication of the most important features, the top 5 features for each optimized decision tree using leave-one-subject-out cross-validation were listed (i.e., the first five splits). Table 9 shows the top features ranked on frequency. While it is possible for features to appear multiple times in the top 5 with different cut-off values this was not the case for the features displayed here. Therefore, the maximum frequency for the features listed is equal to the number of cross-validation folds (=31). These five highest frequency features were among the most important features for

Table 8: Confusion matrix of leave-one-subject-out cross-validation using SVMs with RBF kernel for all data [Overall accuracy = 60%].

Legend – classification of touch gesture captures into a class: ⩾ 10% , ⩾ 50%

		Grab	Hit	Massage	Pat	Pinch	Poke	Press	Rub	Scratch	Slap	Squeeze	Stroke	Tap	Tickle
	Grab	397	0	17	1	11	1	31	4	4	0	177	2	0	0
	Hit	1	317	0	45	1	15	1	0	0	77	3	1	45	0
	Massage	4	0	386	2	1	1	0	63	26	1	14	11	1	26
	Pat	8	58	1	268	1	2	4	1	22	59	0	12	149	18
	Pinch	3	4	6	1	398	27	25	8	8	0	66	1	6	3
	Poke	1	27	0	11	68	438	50	0	2	4	3	1	40	5
Predicted class	Press	19	4	0	7	30	25	374	17	1	7	23	6	8	2
	Rub	0	0	78	2	1	0	8	239	56	0	2	98	0	40
	Scratch	6	0	7	5	0	0	2	50	274	0	0	12	0	92
	Slap	0	77	0	70	2	0	0	2	1	358	0	14	44	1
	Squeeze	117	0	15	0	38	0	50	1	2	0	268	1	0	0
	Stroke	0	1	28	8	1	0	2	125	34	3	0	383	4	15
	Tap	0	68	0	131	6	46	11	2	1	48	0	2	248	16
	Tickle	2	2	19	6	0	3	0	45	127	1	1	12	13	339
	Sum	558	558	557	557	558	558	558	557	558	558	557	556	558	557

Actual class

Table 9: Features that were most frequently ranked within the top 5 for decision tree classification using 31-fold cross-validation

Feature (no.)	Frequency
Mean pressure of the 7th sensor row (10)	31
Absolute distance traveled in the column direction (27)	30
Average spatial peak distance to the center off mass (39)	30
Overall mean absolute pressure difference between frames (43)	30
Size of the contact area of the frame with the highest mean pressure (21)	27

most trained decision trees indicating that these features are reasonably robust. Mean pressure of the 7th sensor row was found to be an important feature for all trees. The 7th sensor row was positioned on the side of the mannequin arm facing away from the participant. When the participant's hand was (partially) folded around the arm it is supposed that the fingers pressed down on this sensor area. A possible explanation for the importance of this feature is that the level of pressure in this sensor area can indicate whether the hand is folded around the arm as would be expected for gestures such as grab and squeeze.

No gender differences were observed based on basic touch gesture characteristics (see Table 5). To look for more subtle differences we classified the touch gestures based on gender using a Bayesian classifier and 10-fold cross-validation. Accuracies ranged from 75% to 78% ($M = 76\%, SD = 0.01\%$), which is similar to the baseline accuracy when classifying every sample as 'male' ($24/31 \approx 77\%$). Based on our findings we have no reason to assume that gender differences play a significant role in touch gesture classification. However, it should be noted that our sample size does not allow us to rule out possible differences.

2.5.2 *Considerations regarding the data collection*

The instructions during the data collection were given in English to include non-native Dutch speakers. However, this could have resulted in translation discrepancies between the English language and the participants' native language. Silvera-Tawil et al. [102] gave the example of the back-translation of the word 'pat' from Spanish to English which can be either translated to 'pat' or 'tap'. Based on observations in a pilot test we opted to include visual examples of the different touch gestures rather than providing participants with the definitions in Table 2 to reduce the language barrier. The use of visual examples instead of giving text-based

definitions could however have reduced the interpersonal differences as participants might have tried to mimic the examples.

To minimize the influence on the participants' natural touch behavior we opted for not restricting the time taken for each touch gesture. Also, there were no constraints on the number of instances of a touch gesture that could be part of a single capture. A consequence of this decision is that a single tap and three taps are both treated as a single touch gesture. This raises the question whether a single tap has a different meaning than three consecutive taps. Furthermore, as features were calculated from the segmented data, segmentation has an influence on features that cover gesture duration (e.g. gesture duration in frames).

The sensor data was labeled according to the instructions (i.e., if the participant was instructed to perform a 'gentle grab', the corresponding sensor data was labeled as such). During segmentation some touch gesture captures were filtered out based on minimal change in gesture intensity, successfully removing skipped touch gestures. However, this procedure does not control for all possible mistakes, which makes it probable that the dataset contains incorrect labels. Manual annotation of the video recordings or outlier detection could help filter out mistakes such as cases where a touch gesture was performed that was different from the one that was instructed.

The inclusion of touch gesture variants seemed to have increased the diversity of the ways in which the touch gestures were performed. Descriptive statistics confirmed that participants used pressure to distinguish between the gesture variants, using less than normal pressure for the gentle variants and more than normal pressure for the rough variants. The definitions of the gentle variant and the rough variant given by the participants also indicated that the amount of pressure is an important way to distinguish between the two for example by the use of the keywords 'soft' and 'hard'. Although speed is also used to differentiate between gentle and rough as indicated by the use of the keywords 'slow' and 'fast', respectively. The downside is that the reliance on pressure to distinguish between both the gestures and the different variants of the same gestures has probably increased the difficulty of the touch gesture recognition. Notably, in the definitions from Table 2 the use of words such as 'forcible', 'gently' and 'firmly' again point to the importance of force/ pressure and also temporal components are mentioned (e.g. 'quickly', 'repeatedly'). As these characteristics seem to be inherent to some of the touch gestures, one may argue that a roughly performed pat, which should generally be 'gentle and quick', would resemble more of a slap, which should generally be 'quick and sharp'. The proposition that some gestures do not lend

themselves as easily for variants is further supported by the self reports of the participants.

2.6 CONCLUSION

To study automatic touch recognition we collected CoST, a dataset containing 7,805 gesture captures of 14 different touch gestures. All touch gestures were performed in three variants: gentle, normal and rough on a pressure sensor grid wrapped around a mannequin arm. We compared the performance of different classifiers: a Bayesian classifier, a decision tree, SVMs (with linear and RBF kernel) and a neural network to establish a baseline for touch gesture recognition within CoST.

The touch data showed similarities between gestures and large differences in the way these gestures were performed which was increased by the inclusion of the different gesture variants. From the different classifiers that were compared, the best results were obtained using SVMs with the RBF kernel while the decision tree yielded the worst performance. Classification of the 14 touch gestures independent of the gesture's variant yielded an average accuracy of up to 60%. The subject independent model generalized well for some individuals but not for others. Moreover, gentle gesture variants proved to be harder to classify than the normal and rough variants which can be due to the less favorable signal-to-noise ratio of these softer gestures. Additionally, misclassifications were found to be most common between touch gestures with similar characteristics such as grab and squeeze.

3

TOUCH CHALLENGE '15: RECOGNIZING SOCIAL TOUCH GESTURES

The organization of the challenge described in this chapter was a collaboration between Merel Jung and Mannes Poel from the University of Twente and Laura Cang and Karon MacLean from the University of British Columbia. Merel Jung and Laura Cang took the lead in the organization of the challenge and in the writing of the publication[1] of the challenge outcomes on which this chapter is based. The content of this chapter is identical to that of the published paper with some minor textual adaptations to embed the content into this dissertation. The future work described in the original paper has been moved to Chapter 7 of the dissertation.

Research into the automatic recognition of social touch has not received much attention yet. To spark interest into this relatively new field we organized a machine learning challenge on the recognition of social touch gestures. For this challenge two touch datasets were publicized: the Corpus of Social Touch (CoST; see Chapter 2) and the Human-Animal Affective Robot Touch gesture set (HAART) [23]. In this chapter we will describe the challenge protocol and summarize the results of the challenge which was hosted in conjunction with the 2015 ACM International Conference on Multimodal Interaction (ICMI).

3.1   INTRODUCTION

Advances in the field of touch recognition could open up applications for touch-based interaction in areas such as human-robot interaction [4, 112]. Unfortunately, the recognition of touch behavior has received far less research attention than recognition of behaviors in the visual and auditory

---

1 Jung, M. M., Cang, X. L., Poel, M., and MacLean, K. E., Touch challenge '15: Recognizing social touch gestures, in Proceedings of the International Conference on Multimodal Interaction (ICMI), (Seattle, WA), 2015, pp. 387–390.

Table 10: The attributes of the two datasets which were provided for the challenge.

Attribute	CoST	HAART
# of touch gestures	14	7
Sensor grid size	8×8	8×8[a]
Sensor sample rate	135 Hz	54 Hz
Sensor values	0–1023	0–1023
Gesture duration	variable	8s[b]
Touch surface	mannequin arm	dependent on condition
Conditions	gentle and normal variations	substrates and covers
# of subjects	31	10
Train/test split	21/10 subjects	7/3 subjects
# of gesture captures	5,203	829

[a]trimmed from collected 10×10 grid, [b]trimmed from collected 10s capture

modalities (e.g. see [126]). For this reason we aimed to spark interest in this relatively new field by organizing a touch challenge.

In previous work, touch data has been collected from subjects performing different sets of touch gestures on different surfaces/ embodiments (e.g. [27, 37, 102, 106], see also Table 1). Although the authors have published their classification results, the data itself has not been shared with the research community. In contrast, we opted to publish two distinct touch datasets for this machine learning challenge to allow researchers with expertise in other sensory modalities to try out their processing techniques on our touch data. Appropriating methods developed in more mature fields such as speech recognition and video analysis could be beneficial for moving touch recognition forward.

The remainder of the chapter is organized as follows. In Section 3.2 we will describe the two touch datasets provided. Next, the challenge protocol will be highlighted in Section 3.3. In Section 3.4 we will provide an overview of the results and discussions of the test set label submissions. Lastly, we will conclude with high-level findings in Section 3.5.

3.2   TOUCH DATASETS

For the challenge, two datasets were made available containing labeled pressure sensor data of social touch gestures. Table 10 summarizes the datasets' attributes.

3.2.1   CoST: Corpus of Social Touch

CoST [60, 62, 65, 113] (see also Chapter 2) contains 14 touch gestures: grab, hit, massage, pat, pinch, poke, press, rub, scratch, slap, squeeze, stroke, tap, and tickle. These gestures were registered on an 8×8 pressure sensor grid which was wrapped around a mannequin arm. This corpus consists of the data from 31 subjects performing the 14 touch gestures in 3 variations: gentle, normal, and rough. Subjects were restricted neither in the amount of time taken for performing each gesture nor the number of gesture repetitions performed in each capture. The data provided for this challenge consisted of the gentle gesture variation (2,601 captures) and the normal gesture variation (2,602 captures). This dataset was provided in both CSV and MATLAB file formats and included segmented gesture captures of varying length, sampled at 135 Hz, and containing pressure values of the 64 channels ranging from 0 to 1023 (i.e., 10 bits). Labels consisted of touch gesture type, gesture variation, and subject number.

3.2.2   HAART: Human-Animal Affective Robot Touch

HAART [23] contains 7 touch gestures: pat, constant contact without movement (press), rub, scratch, stroke, tickle, and 'no touch'. These gestures were found to be the most often used of those in Yohanan et al's Touch Dictionary [124], gathered to communicate emotion in human-animal interaction. For the HAART dataset (collected from 10 subjects), each touch action was performed on a 10×10 pressure sensor [37] for 10 seconds. To assess feature robustness under realistic operating conditions when installed on a robotic animal, each subject contributed gestures with the sensor mounted on all permutations of 3 substrate conditions (firm and flat; foam and flat; foam and curve) and 4 fabric cover conditions (none; short minkee[2]; long minkee; synthetic fur). The resulting dataset includes 829 gesture captures (12 conditions $\times$ 7 gestures $\times$ 10 subjects minus 11 erroneous capture instances). Each capture is 10 seconds of a continuously repeated gesture, sampled at 54 Hz and trimmed to the middle 8s (432 frames); there are generally 10–15 gesture instances per capture. This dataset was provided as a CSV file and included the center 8×8 frame (trimmed for consistency with CoST) with pressure values ranging from 0 to 1023. Labels consisted of touch gesture type, condition set, and subject number.

---

2 Minkee (or minky) is a chenille-like fabric commonly used for baby blankets and stuffed toys.

Table 11: Results for the CoST dataset.

Paper	Classifier	Accuracy
Ta et al. [108]	random forest[a]	61.3%
Ta et al. [108]	random forest[b]	60.8%
Ta et al. [108]	SVM[b]	60.5%
Ta et al. [108]	SVM[a]	59.9%
Gaus et al. [41]	random forest	58.7%
Gaus et al. [41]	multiboost	58.2%
Hughes et al. [53]	logistic regression	47.2%
Balli Altugl et al. [7]	random forest	26.0%

[a]trained on filtered data, [b]train on all data

3.3 CHALLENGE PROTOCOL

The aim of this challenge was to develop relevant features and apply classification methods for recognizing social touch gestures. Gesture classification was independent of the condition (i.e., gesture variant, substrate and cover) for example, 'gentle stroke' and 'normal stroke' were considered to be part of the same class. Participants had the choice of working on one of the datasets or on both. For the train/test sets, subjects were randomly split into 21 train, 10 test subjects for CoST and 7 train, 3 test subjects for HAART. This split ensured that for each of the two datasets, any one subject's touch data belonged to either the training set or the test set.

The training data for both the CoST and HAART datasets was made available to the registrants of the challenge. We provided the test datasets without class labels a month after initial publication. Participants were given 2 weeks to process the test data. Any number of test label submissions could be made up to a deadline (see Tables 11 and 12); once this date had passed, we released the true test labels as well as a summary of their results to the challenge participants.

3.4 CHALLENGE RESULTS AND DISCUSSION

Results of the test label submissions were reported in the form of a confusion matrix and accuracy was used to measure overall performance (see Tables 11 and 12).

Table 12: Results for the HAART dataset.

Paper	Classifier	Accuracy
Ta et al. [108]	random forest	70.9%
Ta et al. [108]	SVM	68.5%
Hughes et al. [53]	logistic regression	67.7%
Gaus et al. [41]	random forest	66.5%
Gaus et al. [41]	multiboost	64.5%
Balli Altugl et al. [7]	random forest	61.0%

3.4.1 *Data pre-processing*

There is not much standardization in the extraction of feature sets for touch data processing. This section discusses the data pre-processing steps that were taken by the challenge participants consisting of data filtering, feature extraction, and feature selection.

The data that was provided for the challenge was previously filtered for erroneous entries and segmented into gesture captures. However, gesture captures of the CoST dataset are of variable duration and can contain a single gesture instance or multiple repetitions. This increased the difficulty of automatic segmentation based on pressure differences over time. Ta et al. explored additional techniques for automatic segmentation to further reduce the number of excess frames [108]. However, these methods for automatic segmentation did not improve classification. The gesture captures from the training set were then manually segmented based on the shape and duration which offered little to negative improvement (see Table 11) suggesting that classifiers are fairly robust to imprecise segmentation.

For the challenge, many interesting features were extracted but we describe only a couple of notable approaches here. In previous literature (e.g. [2, 37, 62, 106]) as well as for this challenge [7, 41, 53, 108] statistics were calculated from the pressure sensor data such as the mean pressure over time. Also, feature extraction methods were borrowed from other domains: speech applications, human action recognition, and image analysis. Ta et al., for example, applied the Sobel operator, an image processing technique used for edge detection [108]. By sharpening the contrast, a second set of data frames was constructed, garnering new values using the same feature extraction procedures. Most features were extracted by feature engineering however, Hughes et al. also included deep autoencoders for automatic feature extraction using dimension reduction, these features

were then used to train Hidden Markov Models (HMMs) [53]. The CoST dataset was used to determine HMM likelihoods for class membership; these values were included in the feature sets for both CoST and HAART data to examine the viability of applying learned features from one dataset to the other.

Feature selection was performed by evaluating the performance of different features or feature sets as a whole on the training set. Relevant features selected using random forest [108] or sequential floating forward search [7] were found to improve the accuracy on the test set. Others compared the accuracies of different feature sets as a whole. Using this approach, the combination of all feature sets yielded the best results [41, 53]. Identifying a small number of highly discriminating features can benefit applications in which computational power is costly, such as onboard real-time touch recognition. Balli Altuglu and Altun [7] showed that a small feature set (number of features < 10) could perform well on the HAART dataset.

3.4.2 *Social Touch Classification*

Random forest was found to be the most popular classification method [7, 41, 108] and has been used in previous work on touch gesture recognition [2, 37]. Other classification methods that were explored were Support Vector Machines (SVMs) [108], also used by [27, 62], multiboost [41] (a different boosting algorithm was used by [102]), and simple logistic regression [53].

Accuracies reported for the challenge ranged from 26.0%–61.3% for the CoST dataset and from 61.0%–70.9% for the HAART dataset (see Tables 11 and 12). Previously reported accuracies for the whole CoST dataset, independent and without knowledge of the gesture variant, yielded accuracies up to 60% using leave-one-subject-out cross-validation [62] (see also Chapter 2). Additionally, accuracies of up to 64.6% were reported for only the rough gesture variants when using 10-fold cross-validation [113]. For the HAART dataset, accuracies up to 90.3% were reported using 20-fold cross-validation when subject and condition labels were included as features [23]. However, direct comparisons between the accuracies reported for the challenge and those reported by the authors of the datasets are not meaningful because of the differences in data division, use of different subsets of the CoST dataset and use of condition and/or subject information as labels.

As accuracy rates alone provide little information we looked at the confusion matrices for notable patterns. Frequent confusions between touch

gestures for the CoST dataset reported by the challenge participants were: 'grab-squeeze', 'hit-pat-slap-tap', 'rub-stroke', and 'scratch-tickle' [7, 41, 53, 108]. The touch gestures were difficult to distinguish across approaches for data pre-processing and classification algorithms. Previous work on the CoST dataset, although using different parts and splits of the dataset, found similar confusions [60, 65, 62, 113] (see also Chapter 2). For the HAART data, rub and tickle were the hardest to correctly classify across challenge participant approaches [7, 41, 53, 108]. Often misclassified was rub as scratch or stroke and tickle as scratch, while the reverse (e.g. misclassification of scratch as rub) was less common. Cang et al. also found that rub and tickle were hardest to classify correctly when using the data from the untrimmed 10×10 sensor grid [23]. Compared to the challenge results, their confusion matrices showed more symmetry, indicating there were frequent confusions between certain gesture pairs. Rub was also one of the most difficult to correctly classify for the CoST dataset [7, 41, 53, 60, 65, 62, 108, 113] (see also Chapter 2).

Based on observations from the recording of the HAART dataset, similarities were observed in how subjects performed the touch gestures which may help to explain certain confusions. Scratch and tickle both followed a similar motion trajectory and tended to have fluttery finger movements. Rub and stroke again have analogous motions where the flat of the hand exerts pressure along a roughly linear path. Confusions between touch gestures on the CoST dataset could also be explained by gestures showing similarities on characteristics such as duration, contact area, repetition probability, and frequency of direction changes [65, 62] (see also Chapter 2).

3.5 CONCLUSION

We extended this machine learning challenge to the research community working on multimodal interaction with the goal of sparking interest in the touch modality and to promote exploration of the use of data processing techniques from other more mature modalities for touch recognition. Two datasets were made available containing labeled pressure sensor data of social touch gestures that were performed by touching a touch-sensitive surface with the hand.

The outcomes of the challenge are encouraging, participants' various approaches open up further avenues for exploring data processing of social touch. The findings of the challenge have provided insights into how techniques for feature extraction that are prominent for other modalities may be applied to touch data. Interestingly enough, many of these techniques

were reasonably transferable to touch gesture data without much modification. Furthermore, despite the use of different data pre-processing techniques and classification algorithms, we observed consistent classification confusions between specific gesture pairs. It is yet unclear whether these classification difficulties can be resolved by finer-grained feature extraction or if the problem is actually our discretization of touch gestures. For instance, scratch and tickle could be regarded as the same gesture class. To conclude, this challenge has allowed for the field of social touch recognition to 'pick up a few tips and tricks' from data processing techniques used for more mature modalities, presenting an opportunity for customizing these methods to meet the particular needs of touch sensor data.

Part III

PUTTING TOUCH IN SOCIAL CONTEXT: SOCIAL TOUCH IN HUMAN-ROBOT INTERACTION

In this section we will present three studies on the use of social touch in interaction with robot pets. The first study will be about the context specific interpretation of social touch in interaction with a robot pet companion. Next, a study will be presented on the ways in which robot pets can express themselves in social interactions using affective breathing behaviors. Lastly, a study on the potential benefits of a robot pet companion with more advanced touch interaction capabilities for health care applications will be presented.

4

UNDERSTANDING SOCIAL TOUCH WITHIN CONTEXT
FOR NATURALISTIC HUMAN-ROBOT INTERACTION

The following chapter[1] covers research which was carried out by Merel
Jung under the supervision of Mannes Poel and Dirk Heylen and with
the help of Dennis Reidsma. The content of this chapter is identical to that
of the published paper with some minor textual adaptations to embed the
content into this dissertation. The future work described in the paper has
been moved to Chapter 7 of the dissertation.

In Part II we focused on the sensing and recognition of different social
touch gestures. However, to respond in an appropriate manner, a social
robot should also be able to interpret these touch gestures within con-
text. In the previous chapters we studied touch gestures that were col-
lected in a controlled lab setting in which participants were given specific
instructions on which gestures they should perform. This controlled ap-
proach lacks social context which could help to recognize touch gestures
and their inferred meaning. In contrast, in this chapter we will present a
study towards the interpretation of social touch in a contextualized lab
setting in which participants acted as if they were coming home in differ-
ent emotional states. Human interaction with a robot pet companion will
be observed by looking at touch behavior as well as other social behaviors
such as speech. Moreover, to further close the interaction loop we will also
investigate appropriate robot responses.

4.1 INTRODUCTION

To behave socially intelligently a robot should not only be able to sense
and recognize touch gestures but should also be able to interpret those

---

1 Based on Jung, M. M., Poel, M., Reidsma, D., and Heylen, D. K. J., A first step towards the
automatic understanding of social touch for naturalistic human-robot interaction, Frontiers
in ICT, vol. 4, no. 3, 2017.

touch gestures in order to respond in a socially appropriate manner as illustrate in Figure 1. Perhaps robot seal Paro is the most famous example of a social robot that responds to touch [117]. Paro is equipped with touch sensors with which it distinguishes between soft touches (which are always interpreted to be positive) and rough touches (which are always interpreted to be negative) [117]. This interpretation of touch is oversimplified as the complexity of the human tactile system allows for touch behaviors to vary not only depending on the intensity but also based on movement, velocity, abruptness, temperature, location and duration [48]. Moreover, the meaning of touch can often not be inferred from the type of touch alone but is also dependent on other factors such as concurrent verbal and nonverbal behavior, the type of interpersonal relationship [51, 107] and the situation in which the touch takes place [59]. Although previous research [48, 49, 51] indicated that there is no one-to-one mapping of touch gestures to a specific meaning of touch, touch can have a clear meaning in a specific context [59].

The focus of this part of the dissertation is on (touch) interaction with a robot pet companion. According to Veevers, a pet companion can fulfill different roles in the life of humans. A pet can facilitate interpersonal interaction or can even serve as a surrogate for interpersonal interaction and expensive and/or exotic pets can be owned as a status symbol [114]. Furthermore, interaction with a pet is associated with health benefits and more recent studies indicate that these effects also extend to interaction with robot pets [8, 32] (see also Chapter 6). Although touch is a natural way to interact with real pets, currently commercially available robot pets such as Paro [117], Hasbro's companion pets[2] and JustoCat[3] are equipped with only a few touch sensors and do not interpret different types of touch within context.

We argue that the recognition and interpretation of touch consists of three levels: 1) low-level touch parameters such as intensity, duration and contact area; 2) mid-level touch gestures such as pat, stroke and tickle; 3) high-level social messages such as affection, greeting and play. To automatically understand social touch, research focuses on investigating the connection between these levels such as which touch gestures are used to communicate different emotions (e.g. [124]). Current studies in the domain of social touch for human-robot interaction have been mainly focused on highly controlled settings in which users were requested to perform different touch behaviors, one at a time, according to predefined labels (e.g. [23, 27, 61, 102, 103, 124], see also Chapter 2). In this study we

---

2 http://joyforall.hasbro.com
3 http://justocat.com

will focus on the latter two levels (i.e., touch gestures and social messages) as we are interested in the meaning of touch behaviors. To gain more insight into the factors that are relevant to interpret touch behaviors within social context we opted to elicit touch behaviors by letting participants act out four affective scenarios in which they interacted with a robot pet companion. Moreover, in contrast to most previous studies, participants could freely act out the given scenarios with the robot pet within the confined space of a living room setting.

In this chapter we present contributions in two areas. Firstly, we explore the use of touch behaviors as well as the expressed social messages and expected robot pet responses in different affective scenarios. Secondly, we reflect upon the challenges of the segmentation and labeling of touch behaviors in a less controlled setting in which no specific instructions are given on the kinds of (touch) behaviors that should be displayed. We address the first contribution with the following three research questions. RQ1) What kinds of touch gestures are used to communicate with a robot pet in the different affective scenarios? RQ2) Which social messages are communicated and what is the expected response in the different affective scenarios? RQ3) Which other social signals can aid the interpretation of touch behaviors? Furthermore, we reflect upon our effort to segment and label touch behaviors in a less controlled setting with the forth research question. RQ4) How well do annotation schemes work in a contextualized lab situation?

The remainder of the chapter is structured as follows. Related work on the meaning of social touch in both interpersonal and human-robot interaction will be discussed in Section 4.2 followed by the description of the material and methods for the study presented in Section 4.3. Then, the results will be provided and discussed in Section 4.4 and Section 4.5, respectively. Conclusions will be drawn in Section 4.6.

4.2   RELATED WORK

Previous studies have looked into the meaning of touch in both interpersonal interaction and human-robot interaction. In the latter case researchers have looked at the automatic recognition of emotions and other social messages based on human touch.

In a diary study on the use of interpersonal touch, different meanings of touch were categorized based on the participants' verbal translations of the touch interactions [59]. Seven main categories were distinguished: positive affect touches (e.g. support), playful touches (e.g. playful affection), control touches (e.g. attention-getting), ritualistic touches (e.g. greeting),

hybrid touches (e.g. greeting/ affection), task-related touches (instrumental intrinsic) and accidental touches. Interestingly, there was a lack of reports on negative interpersonal touch interaction. Within these categories, common contextual factors were identified such as the type of touch, any accompanying verbal statement and the situation in which the touch took place. It was found that depending on the context, a specific form of touch can have multiple meanings and that different forms of touch can have a similar meaning. Furthermore, touch was found to be often preceded, accompanied or followed by a verbal statement.

In a study on human-robot interaction, participants were asked to indicate which touch gestures they were likely to use to communicate emotional states to a cat-sized robot animal [124]. Gestures that were judged to be likely to be used were performed sequentially on the robot. Participants expected that the robot's emotional response would be either similar or sympathetic to the emotional state that was communicated. The nature of the touch behavior was found to be friendly as no aggressive gestures (e.g. slap or hit) were used even when negative emotions were communicated. Five categories of intent were distinguished based on touch gesture characteristics which could be mapped to emotional states: affectionate, comforting, playful, protective and restful. Also, video segments of the touch gestures were annotated to characterize the gestures based on the point of contact, intensity and duration revealing differences between touch gestures and their use in different emotional states. In follow-up research the touch sensor data recorded in this study (i.e., [124]) was used to classify 26 touch gestures and 9 emotional states using random forests [2]. Between-subjects emotion recognition of 9 emotional states yielded an accuracy of 36% while within-subjects the accuracy was 48%. Between-subjects touch gesture recognition of 26 gestures yielded an accuracy of 33%. Furthermore, the authors' results indicated that accurate touch gesture recognition could improve affect recognition.

In other work, Kim et al. instructed participants to use four different touch gestures to give positive or negative feedback to a humanoid robot while playing a game [71]. A model was trained to infer whether a touch gesture was meant as a positive or a negative reward for the robot. It was found that participants used 'pat' and 'rub' to give positive feedback and 'hit' to give negative feedback while 'push' could be used for both although the touch gesture was mostly used for negative feedback. Knight et al. argued for the importance of body location as contextual factor to infer the meaning of touch [73]. The authors made the distinction between what they called *symbolic gestures* which have social significance based on

the involved body location(s) (e.g. footrub and hug) and body location independent *touch subgestures* (e.g. pat and poke).

Although previous studies indicate that there is a link between touch gestures and the higher level social meaning of touch, Silvera-Tawil et al. argued that the meaning of touch (i.e., level 3 as described in Section 4.1) could also be recognized directly based on characteristics from touch sensor data (i.e., level 1) and other factors such as the context and the touch location [103]. In their effort to automatically recognize emotions and social messages directly from sensor data without first recognizing the used touch gestures, participants were asked to perform six basic emotions: anger, disgust, fear, happiness, sadness, and surprise on both a mannequin arm with an artificial skin and a human arm. In addition, six social messages were communicated: acceptance, affection, animosity, attention-getting, greeting and rejection. Recognition rates for the emotions were 47% for the algorithm and 52% for human classification. The recognition rates for the social messages were found to be slightly higher, yielding accuracies of 50% and 62% for the algorithm and human classification, respectively. These results show that humans currently outperform the automatic recognition of these forms of social touch.

Some attempts have been made to study touch interaction in a less controlled setting, for example Noda et al. elicited touch during the interaction with a humanoid robot by designing a scenario in which participants used different touch gestures to communicate a particular social message such as greeting the robot by shaking hands, playing together by tickling the robot and hugging the robot to say goodbye at the end of the interaction [87]. Results showed an accuracy of over 60% for the recognition of the different touch behaviors that were performed within the scenario. In another study on the use of touch in multimodal human-robot interaction, participants were given various reasons to interact with a small humanoid robot such as giving reassurance, getting attention and giving approval [28]. The robot was capable of recognizing touch, speech and visual cues and participants were free to use different modalities. Also, participants rated videos in which a confederate interacted with the robot using different modalities. Results showed that touch was often used to communicate with the robot and that touch was especially important for expressing affection. Furthermore, playing with the robot and expressing loneliness were deemed more suitable than displaying negative emotions.

To summarize, previous studies illustrate that touch can be used to express and communicate different kinds of affective and social messages [28, 59, 103, 124]. Moreover, touch gestures that were used to communicate were often positive in nature and their meaning was dependent on the con-

text such as one's emotional state [28, 59, 124]. These findings confirm that currently available robot pet companions such as Paro, which only distinguishes between positive and negative touch, are not sufficiently capable of understanding and responding to people in a socially appropriate way. Furthermore, there are indications that other modalities might be helpful in interpreting the social messages as touch behavior generally does not occur in isolation [28, 59]. For the reasons outlined above we opted to study humans interacting with a robot pet companion, in different emotional states, and, in a contextualized lab situation.

4.3   MATERIAL AND METHODS

In this study we elicited interactions between a human and a robot pet companion in a lab-built living room setting. We were specifically interested in the use of touch behaviors as well as the expressed social messages and expected robot pet responses in different affective scenarios. Participants of the study were instructed to act as if they had come home in different emotional states (i.e., stressed, depressed, relaxed and excited). These four emotional states were chosen as they span opposite ends of the valence and arousal scale (see Figure 8): stressed (negative valence, high arousal), depressed (negative valence, low arousal), relaxed (positive valence, low arousal) and excited (positive valence, high arousal) [93]. Furthermore, similar emotional states have been used in a more controlled research setting before and the results from this study indicate that emotional state influences touch behavior as well as the expected robot response [124]. To gain more insight into the factors that are relevant to interpret touch within a social context we annotated touch behaviors from video footage of the interactions. Also, a questionnaire was administered and interviews were conducted to interpret the high-level meaning behind the interactions and to gain insight into the responses that would be expected from the robot pet.

4.3.1   *Participants*

In total 31 participants (20 male, 11 female) volunteered to take part in the study. The age of the participants ranged from 22 to 64 years (M = 34.3; SD = 12.8) and 28 were right-handed, 2 left-handed and 1 ambidextrous. All studied or worked at the University of Twente in the Netherlands. Most (21) had the Dutch nationality, others were Belgian, Ecuadorean, English, German (2x), Greek, Indian, Iranian, Italian and South Korean. This study
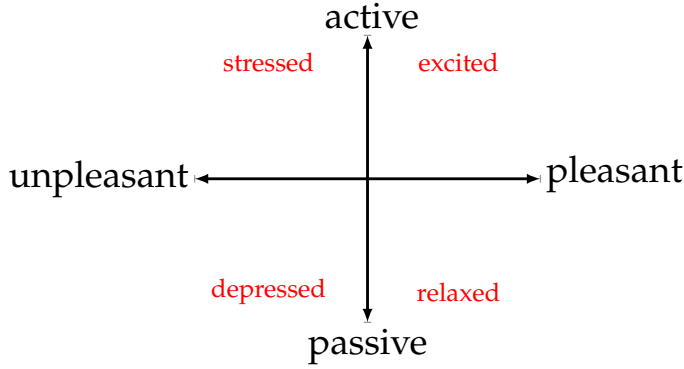
Figure 8: Mapping of emotional stated based on associated valence and arousal levels, model adapted from [93].

was approved by the ethics committee of the Faculty of Electrical Engineering, Mathematics and Computer Science of the University of Twente.

4.3.2 *Apparatus/ Materials*

4.3.2.1 *Living room setting*

The living room setting consisted of a space of approximately 23 m^2 containing a small couch, a coffee table and two plants (see Figure 9, top). Two camcorders were positioned facing the couch at an approximately 45 degree angle to record the interactions (50 fps, 1080p). To allow participants to interact freely with the robot pet (i.e., no wires) and have a controlled interaction (i.e., no unpredictable robot behavior) a stuffed toy dog was used as a proxy for a robot pet (see Figure 9). The robot pet (35 cm; in a laying position) was positioned on the couch at the far end from the door facing the table.

4.3.2.2 *Questionnaire*

The questionnaire was divided into two parts. Part one was completed before the interview was conducted and part two after the interview. Part one consisted of demographics: gender, age, nationality, occupation and handedness followed by six questions about the reenactment of the scenarios rated on a 4-point Likert scale ranging from 1 (strongly disagree) to 4 (strongly agree). Four questions were about the participants' ability to imagine themselves in the scenarios: 'I was able to imagine myself coming home feeling stressed/ depressed/ relaxed/ excited'. The other two

Figure 9: The living room setting with the robot pet on the couch (top) and the pet up-close (bottom).

questions were about the robot pet: 'I was able to imagine that the pet was a functional robot' and 'I based my interaction with the robot pet on how I interact with a real animal'.

The second part consisted of a questionnaire about the expectations of living with a robot pet which was based on the 11-item Comfort from Companion Animals Scale (CCAS) [125]. Participants were asked to imagine that they would get a robot pet like the one in the study as a gift. This robot pet can react to touch and verbal commands. Participants were asked to answer the questions about the role they expected the robot pet would play in their life. The questions from the CCAS were adjusted to fit the purpose of the study, for example the item 'my pet provides me with companionship' was changed to 'I expect my robot pet to provide me with companionship'. Items were rated on a 4-point Likert scale ranging from 1 (strongly disagree) to 4 (strongly agree), as all items were phrased positively a higher score indicates greater expected comfort from the robot pet.

4.3.2.3 *Interview*

A semi-structured interview was conducted between the first and the second part of the questionnaire. The video footage of their reenactment of the scenarios was shown to the participants and they were asked to answer the following questions after watching each of the four scenario fragments: 1) 'what message did you want to communicate to the robot?', 2) 'what response would you expect from the robot?' and 3) 'how could the robot express this?'. The participant, the interviewer and the computer screen were recorded during the interview using a camcorder.

4.3.3 *Procedure*

Upon entering the room in which the study took place, the participant was welcomed and was asked to read the instructions and sign an informed consent form. Then, participants were taken into the hallway where they received the instructions for the example scenario in which they were asked to act out coming home in a neutral mood. If the instructions were clear to the participants, they were asked to interact with a robot pet by acting out four different scenarios, one by one, in which they would come home in a particular emotional state, feeling either: stressed, depressed, relaxed or excited. The study had a within-subject design, instructions for each of the scenarios were given to each of the participants in random order. In each scenario the participant was instructed to enter the 'living room', sit down on the couch and act out the scenario as he/ she saw

fit. Participants were instructed to focus on the initial interaction as the robot pet would not respond ($\approx$ 30 seconds were given as a guideline), however the duration of the interaction was up to the participant who was instructed to return to the hallway when he or she had finished an interaction. When the participant had returned to the hallway at the end of an interaction the next scenario was provided.

After the last scenario the participant was asked to fill out a questionnaire asking for demographic information and about acting out the scenarios. Then the video footage of their reenactment of the scenarios was shown to the participant and an interview was conducted on these interactions. After the interview the participants completed the second part of the questionnaire about their expectations if they were to own a functional robot pet. The entire procedure took approximately 20 minutes for each participant. At the end of the study participants were offered a candy bar to thank them for participating.

4.3.4   *Data analysis*

4.3.4.1   *Questionnaire*

The questionnaire data was analyzed using IBM SPSS Statistics version 22. The median values and the 25th and 75th percentiles (i.e., Q1 and Q3, respectively) were calculated for the questions about the reenactment of the scenarios. The ratings on the items of the expected comfort from the robot pet scale were summed before calculating these descriptive statistics. Additionally, a Friedman test was conducted to check whether there was a statistical difference between the perceived ability of the participants to imagine themselves in the different scenarios. The significance threshold was set at .05 and the exact p-value is reported for a two-tailed test.

4.3.4.2   *Annotation of scenario videos*

The video footage from the two cameras were synced and put together in a split screen video before annotation. Videos were coded by two annotators[4], which included one of the members of the researcher team henceforth 'the first coder', using the ELAN[5] annotation software.

For the segmentation of touch behaviors we followed a method that is commonly used to segment signs and co-speech gestures into movement units which in the simplest form consist of three phases: a preparation

---

4 We would like to acknowledge Josca Snippe for participating in the annotation process.
5 Max Planck Institute for Psycholinguistics, The Language Archive, Nijmegen, The Netherlands; http://tla.mpi.nl/tools/tla-tools/elan
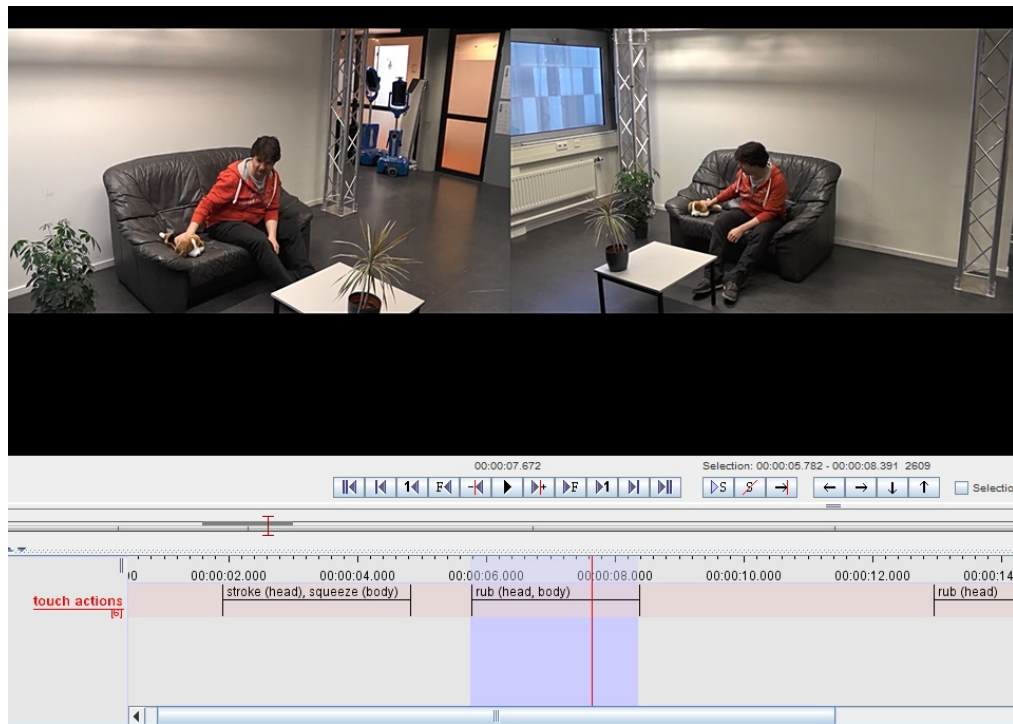
Figure 10: Screenshot of the annotation process showing the tier in which the touch gestures and the body location on the robot pet are annotated.

phase, an expressive phase, and a retraction phase [69, 72]. The onset of a movement unit is defined at the first indication of the initiation of a movement which is usually preceded by the departure of the hand's resting position. The end of a movement unit is defined as the moment when the hand makes first contact with a resting surface such as the lap or an arm rest. Similarly, the touch actions were segmented by the first coder from the moment that the participant reached out to the robot pet to make physical contact until the contact with the pet was ended and the hands of the participants returned to the resting position. Per touch action segment the following information was coded by the two annotators in a single annotation tier: the performed sequence of touch gestures and the robot pet's body part(s) on which each touch gesture was performed (see Figure 10). The touch gesture categories consisted of the 30 touch gestures plus their definitions from the touch dictionary of Yohanan and MacLean [124] which is a list of plausible touch gestures for interaction with a robot pet. Furthermore, based on observations we added an additional category for *puppeteering* which was defined as 'participant puppeteers the robot pet to pretend that it is moving on its own' and to reduce forced-choice we added *other* which was defined as 'the touch gesture performed cannot be described by any of the previous categories'. The robot pet's body parts were divided into six categories: head (i.e., back, top and sides of the head and ears), face (i.e., forehead, eyes, nose, mouth, cheeks and chin), body (i.e., neck, back and sides), belly, legs and tail.

Coding the touch behaviors that were performed during each touch segment proved to be difficult. Both annotators were often unsure when to define the start of a new touch gesture as gestures were often followed-up in quick succession. Furthermore, hybrid forms of several touch gestures were often observed resulting in difficulties to categorize the touch behavior into one of the categories. In Table 13 some of the touch gestures are listed that were frequently observed but that were also difficult to distinguish based on their dictionary definitions. These touch gestures are all of relatively long duration compared to quick gestures such as pat and slap and all include movement across the contact area. The distinguishing features are based on the gesture's intensity, human contact point (e.g. whole hand vs. fingernails) and the movement pattern (e.g. back and forth or seemingly random). An example of commonly encountered confusion was in cases where the hand was moved repeatedly back and forth on the fur of the robot pet which indicated the use of a rub gesture while the use of gentle pressure seemed to indicate a stroke-like gesture. Furthermore, the use of video footage to code touch gestures made it difficult to determine the exact part of the hand that was used to perform the gesture

which is the only differentiating feature to distinguish between a rub and a scratch gesture based on these definitions.

Table 13: Example touch gesture categories with definitions, adapted from Yohanan and MacLean [124].

Gesture label	Gesture definition
Massage	Rub or knead the robot pet with your hands.
Rub	Move your hand repeatedly back and forth on the fur of the robot pet with firm pressure.
Scratch	Rub the robot pet with your fingernails.
Stroke	Move your hand with gentle pressure over the robot pet's fur, often repeatedly.
Tickle	Touch the robot pet with light finger movements.

Even after several iterations of revisiting the codebook in order to clarify what the distinguishing features of several touch gestures are it was still not possible to reach an acceptable level of agreement. Difficulties were caused by a mixture of touch events that were hard to observe on video and differences in interpretation by the annotators, which included both the segmentation of individual touch gestures (i.e., within the larger predefined segments) and the assignment of labels, despite the commonly developed annotation scheme. Furthermore, as one touch segment could consist of a sequence of touch gestures it was difficult to calculate the inter-rater reliability (i.e., Cohen's kappa) as the number of touch gestures could differ per coder. The location of the touch gestures on the robot pet's body was related to the coding of the touch gestures themselves and therefore it was also not possible to reach an acceptable agreement on this part.

Due to the difficulties described above we decided instead to coarsely describe the interactions in the results section based on the modalities that the participants used to communicate to the robot pet. Also, a Friedman test was conducted to check whether there was a statistical difference between the duration of the interactions in the different scenarios. The significance threshold was set at .05 and the exact p-value is reported for a two-tailed test. The implications of the findings from the annotation process will be explicated in the discussion section.

4.3.4.3 *Interview*

The interview answers were grouped per scenario based on common themes. The data was split into two parts. 1) Information on the social

messages (and possible behaviors to express those) that were communicated by the participant to the robot pet. 2) Information on the expected messages and behaviors that were expected to be communicated by the robot pet. Themes were labeled and the number of participants that mentioned the specific topic was counted for each scenario. Furthermore, the communicated social messages for each scenario were mapped to the expected responses from the robot pet to look for frequently occurring patterns.

4.4   RESULTS

4.4.1   *Questionnaire*

Participants' rating of their ability to imagine themselves in the four different scenarios on a scale ranging from 1 (strongly disagree) to 4 (strongly agree)) were the following: stressed $(Mdn(Q1, Q3) = 3(3, 3))$, depressed $(Mdn(Q1, Q3) = 3(2, 3))$, relaxed $(Mdn(Q1, Q3) = 3(3, 3))$, excited $(Mdn(Q1, Q3) = 3(2, 4))$. There was no statistically significant difference between the ratings of the scenarios $(\chi^2(3) = 3.297, p = .352)$. Median $(Q1, Q3)$ perceived ability to imagine the pet as a functional robot was $2(2, 2)$ and the statement 'I based my interaction with the robot pet on how I interact with a real animal' was rated at $3(2, 4)$. The total scores for the expected comfort from the robot pet ranged from 18 to 37 $(Mdn(Q1, Q3) = 30(26, 35))$, possible total scores ranged from 11 to 44 where a higher score indicated greater expected comfort.

Table 14: Number of participants that engaged in different levels of interaction with the robot pet per scenario.

	Emotional state			
Interaction type	Stressed	Depressed	Relaxed	Excited
No interaction	3	3	0	0
Speech only	2	0	0	0
Touch only	8	12	13	7
Touch + speech	18	16	18	24
Sum	31	31	31	31

4.4.2 *Observations from the scenario videos*

Between the different scenarios there were some differences in the level of interaction with the robot pet (see Table 14). Participants often used both touch and speech to communicate with the robot pet. Almost all participants used at least the touch modality to communicate, few exceptions occurred in the negative valence scenarios (i.e., stressed and depressed). Examples of touch behaviors that were observed were participants sitting next to the robot pet on the couch while touching it using stroking-like gestures, hugging the pet and having the robot pet sit on their laps while resting a hand on top of it. Speech was most prevalent in the excited scenario while it was least prevalent in the depressed scenario. Observed behaviors included participants using speech to greet the robot pet when entering, to talk about their day, to express their emotional state and to show interest in the pet. Some instances of pet-directed speech were observed as well. Another notable observation was that participants oriented the robot pet to face them indicating that they wanted to make eye contact. Furthermore, some participants incorporated the use of their mobile phone in the scenarios for example to indicate that they would be preoccupied with their own activities (e.g. sending text messages to friends), to take a picture of the robot pet or to watch online videos together. Others engaged in fake activities with imaginary objects such as playing catch or watching TV together.

The duration of an interaction was measured as the time in seconds between the start of the interaction (i.e., opening the door to enter the living room) and the end (i.e., closing the door after leaving the room). Overall, the duration of the interactions ranged between 17 and 112 seconds. There was a statistically significant difference in the duration of interaction between the four scenarios ($\chi^2(3) = 16.347$, p = .001). A post hoc analysis with Wilcoxon signed-rank tests was conducted with a Bonferroni correction applied, resulting in a significance level set at p < .008. The median (Q1, Q3) duration in seconds for each of the scenarios was: stressed 41(29, 55), depressed 42(32, 55), relaxed 42(32, 53) and excited 35(28, 45). The duration of interaction in the excited was significantly shorter compared to the other scenarios: stressed (Z = −2.968, p = .002), depressed (Z = −3.875, p < .001) and relaxed (Z = 3.316, p = .001). The other scenarios did not differ significantly (all p's > .008).

4.4.3   *Interview*

In general, most participants watched the whole scenario before answering the questions while others commented on their behavior right away. Also, some participants mentioned at the beginning that they felt a bit awkward watching themselves on video. The social messages that were communicated by the participant to the robot pet and messages that were expected to be communicated by the robot pet are listed for each scenario in Table 15. Table 16 shows the mapping between the two most frequently communicated social messages for each scenario and the most common expected responses from the robot pet to these messages. We will discuss the interview results further based on these mappings.

4.4.3.1   *Stressed*

In the stressed scenario most participants wanted to communicate that they were stressed by indicating to the robot pet that they had lots of things to do or that they were preoccupied with something (n = 17). Notably, some of these participants involved the robot pet as a way to regulate their emotions by touching the pet as a means of distraction. In response some of these participants wanted company from the robot pet by staying close and through physical interaction (n = 6). Importantly, the pet's behavior should be calm and the robot should not be too demanding. Other participants wanted support from the robot pet by calming them down and providing comfort (n = 6).

In contrast, some participants did not want to interact with the robot pet at all as they preferred to be alone in this situation or did not wanted to be distracted by the pet (n = 6). In response most participants wanted that the robot pet showed its understanding of the situation by keeping its distance (n = 3). Others mentioned that the robot pet should have its own personality and should behave accordingly, which might result in the robot pet asking for attention even if this behavior is undesirable in this situation or that the pet would mind its own business (n = 2).

4.4.3.2   *Depressed*

In the depressed scenario participants often communicated to the robot pet that they were looking for comfort in order to feel less depressed (n = 11). In response these participants often wanted comfort from the robot pet (n = 6). They wanted the pet to do this by sitting on their lap or right next to them and making sounds. Also, participants specified that the robot pet should not approach them too enthusiastically. Others

Table 15: Social messages that were communicated to the robot pet and its expected responses for each scenario. The number of participants is in parentheses.

Emotional state			
Stressed	Depressed	Relaxed	Excited
Social messages that were communicated to the robot pet			
Express emotional state (17)	Seek emotional support (11)	Enjoy company (14)	Express emotional state (15)
Do not want to interact (6)	Express emotional state (8)	Express emotional state (8)	Actively seek interaction (11)
Acknowledge (3)	Do not want to interact (6)	Acknowledge (7)	Enjoy company (4)
Seek emotional support (3)	Want to interact (3)	No expectations (2)	Do not want to interact (1)
Actively seek interaction (2)	Acknowledge (1)		
	Enjoy company (1)		
	No expectations (1)		
Social messages that the robot pet is expected to communicate			
Keep company (8)	Keep company (12)	Keep company (12)	Pick up the mood (24)
Provide emotional support (7)	Provide emotional support (11)	Pick up the mood (7)	Engage in interaction (6)
Focus on own needs (6)	Engage in interaction (4)	Engage in interaction (5)	Show appreciation (1)
Understand the situation (5)	Focus on own needs (2)	Focus on own needs (5)	
Engage in interaction (3)	Ask for attention (1)	No interaction (1)	
Do not understand (2)	Show appreciation (1)	Do not understand (1)	

indicated that the robot pet should keep them company (n = 3) by staying close and showing its understanding of the situation.

Other participants just wanted to express how they felt (n = 8), for example by telling the pet why they were feeling depressed. In response most of these participants also expected that the robot pet would either provide emotional support (n = 4) or would keep them company (n = 3).

Table 16: Breakdown of the most frequently communicated messages to the robot pet and the expected responses for each scenario. The number of participants is in parentheses.

Emotional state	Communicated message		Expected robot pet response
Stressed	Express emotional state (17)	→	Keep company (6)
			Provide emotional support (6)
	Do not want to interact (6)	→	Understand the situation (3)
			Focus on own needs (2)
Depressed	Seek emotional support (11)	→	Provide emotional support (6)
			Keep company (3)
	Express emotional state (8)	→	Provide emotional support (4)
			Keep company (3)
Relaxed	Enjoy company (14)	→	Keep company (10)
	Express emotional state (8)	→	Pick up the mood (5)
Excited	Express emotional state (15)	→	Pick up the mood (15)
	Actively seek interaction (11)	→	Pick up the mood (6)
			Engage in interaction (5)

4.4.3.3 *Relaxed*

In the relaxed scenario participants often wanted to communicate that they enjoyed the pet's company (n = 14), for example by having the pet sit on their lap or right next to them, touching the robot and talking to it. In response these participants often wanted company from the robot pet (n = 10), for example by staying close, listening and engaging in physical interaction. Furthermore, the pet's behavior should be calm and should reflect that it enjoys being together with the human (e.g. wagging tail or purring).

Other participants mentioned that they wanted to express that they were feeling relaxed (n = 8) such as by telling the pet about their day and that everything was alright. In response most of these participants wanted

the robot pet to pick up on their mood by displaying relaxed behavior as well such as by lying down (n = 5).

4.4.3.4 *Excited*

In the excited scenario participants often wanted to communicate their excitement to the robot pet (n = 15), for example by touching and talking to the robot. In response all these participants wanted the robot pet to pick up on their mood by becoming excited as well (n = 15). The robot pet could show its excitement by actively moving around, wagging its tail and making positive sounds.

Other participants wanted to actively interact with the robot pet (n = 11) by playing with it or going out for a walk together. In response most of these participants wanted the robot pet to pick up on their mood as well (n = 24) or preferred that the robot pet would actively engage them in play behavior (n = 6).

4.5 DISCUSSION

4.5.1 *Categorization of touch behaviors*

In this study we observed participants that interacted freely with a robot pet companion. As a consequence we observed an interesting but complex mix of touch behaviors such as the use of multiple touch gestures that were alternated, hybrid forms of prototypical touch gestures and combinations of simultaneously performed touch gestures (e.g. stroking while hugging). A previous attempt to annotate touch behaviors by Yohanan and MacLean [124] was limited to the coding of characteristics of touch gestures that were performed sequentially which completely eliminated difficulties regarding segmentation and labeling that were encountered in this study. Segmentation and labeling of individual touch gestures based on a method borrowed from previous work on air gestures proved not to be straightforward. Although air gestures and touch gestures both rely on the same modality (i.e., movements of the hand(s)) their communicative functions are different. Air gestures, especially sign language, are a more explicit form of communication compared to communication through touch in which there is no one-to-one mapping between touch gestures and their meaning. Furthermore, in less controlled interactions it proved to be difficult to categorize touch behaviors into discrete touch gesture categories based on dictionary definitions, such as the gestures defined in Table 13. These results indicate that this approach might not be suitable to capture the nature of touch behavior in less controlled settings.

In accordance with previous findings from Yohanan and MacLean [124] we frequently observed the use of massage, rub, scratch, stroke and tickle-like gestures to communicate to the robot pet. As a result valuable information would be lost if these gestures were to be collapsed into a single category to bypass the difficulties to clearly distinguish between these gestures. Some of the difficulties were due to the use of video footage to observe touch behavior. For example, the intensity level can only be roughly estimated from video (see also [124]) and some details such as the precise point of contact were lost because of occlusion. However, confusions in identifying touch gestures with similar characteristics were also observed in studies where touch behaviors were captured by pressure sensors and algorithms were trained to automatically recognize different gestures (e.g. [61, 62, 102], see also Chapters 2 and 3). Moreover, segmentation and categorization of touch behavior based on touch sensor data would still remain challenging.

As the segmentation and categorization of touch behaviors into touch gestures might not be that straightforward in less controlled settings it might be more sensible to recognize and interpret social messages directly from touch sensor data as was previously suggested by Silvera-Tawil et al. [103]. Moreover, feature extraction methods from other modalities such as image processing (e.g. edge detection [108]), speech (e.g. dynamic time warping [108]) and action recognition (e.g. motion history histogram [41]) proved to be transferable to touch gesture recognition [61] (see also Chapter 3. Therefore, the existing body of literature on the transition towards automatic behavior analysis of these modalities in naturalistic settings might provide valuable insights for the understanding of touch behavior as well (e.g. see [45, 66, 86]).

4.5.2   *Observed multimodal behaviors*

The following coarse descriptions of interactions with the robot pet from two different participants illustrate the use of multimodal cues in the depressed and excited scenario, respectively.

> "Participant walks into the living room and sits down on the couch next to the robot pet. Immediately she picks up the pet and holds it against her body using a hug-like gesture. While holding the pet she tells the pet that she had a bad day while she makes eye contact from time to time. Then she sits quietly while still holding the pet and making eye contact. Finally, she puts the pet back on the couch and gets up to leave the room."

"Participant runs into the living room and slides in front of the couch. He picks up the robot pet from the couch and then sits down on the couch with the pet resting on his leg. Then he talks to the pet using pet-directed speech: 'How are you? How are you? Yes! You're a good dog! Good doggy!'. Meanwhile he touches the pet using stroke-like gestures and looks at it. He then puts the robot pet back on the couch again while he still touches the pet using stroke-like gestures. Finally, he gets up from the couch and leaves the room."

As illustrated in the descriptions above, participants often talked to the robot pet while touching it (see also Table 14) indicating that the combination of speech (emotion) recognition and touch recognition might aid the understanding of touch behavior. Although we observed forms of speech that had characteristics of pet-directed speech (e.g. short sentences, repetition and higher pitched voice) it should be noted that no analysis of the prosodic features of the speech was performed. However, the use of pet-directed speech has been observed previously, for example, Batliner et al. found that children used pet-directed speech when interacting with Sony's pet robot dog AIBO [11]. A limitation of the current setup is that it did not allow for a detailed analysis on the added value of other social cues such as facial expression, posture and gaze behavior for the interpretation of touch behavior.

By allowing the participants to interact freely with the robot pet within the confined space of a living room setting we were able to observe behavior that might otherwise not be observed. Social interaction involving objects such as taking pictures of the robot with a mobile phone were also observed by Cooney et al. who argued that these factors should be investigated to enable rich social interaction with robots [28]. However, it is important to keep in mind that although participants in this study were able to interact freely within the given context, the results are confined to the given interaction scenarios. Furthermore, as the study relied on acted behaviors participants might have displayed prototypical behaviors to clearly differentiate between the scenarios. However, although participants indicated in the questionnaire and during the interview that they had some difficulties acting out the scenarios with a stuffed animal, social behaviors such as making eye contact while talking were observed (see also the descriptions above) indicating that at least most participants treated the pet as a social agent. Additionally, it should be noted that touch was not only used to communicate to the robot pet but was also often used to move/ puppeteer the robot pet as it was unable to move on its own.

Surprisingly, interactions in the excited scenario were shorter despite the fact that all participants engaged in some form of interaction with the robot pet (see Table 14). A possible explanation is that participants often only quickly wanted to convey their excitement compared to other scenarios where they were seeking comfort or quietly sat down together with the robot pet to enjoy each other's company (see Table 15). Furthermore, previous studies indicate that some emotions are more straightforward than others, for example, anger was found to be easier to express through touch than sadness [48]. Similarly, excitement might have been easier to convey than the other emotional states in this study.

4.5.3   *Communicated social messages and expected robot pet responses*

The interview results showed that the communicated messages and expected robot pet responses differed depending on the affective scenario and individual preference (see Table 15). Moreover, Table 16 shows that there is no one-to-one relation between communicated messages and expected responses. For example, variation in expectations from the robot pet in the stressed scenario ranged from actively providing support to staying out of the way meaning that in order to respond in a socially appropriate manner, a robot pet should be able to judge whether the user wants to be left alone and when to engage in interaction. From the interviews it became clear that this is not always clear-cut, in the depressed and stressed scenarios some participants indicated that they did not want to initiate interaction but that they might be open to the robot pet approaching them (sometimes after a while). Participants often wanted to communicate their emotional state to the robot pet, especially in the high arousal scenarios (see Table 15) However, it should be noted that the focus on emotional states in the scenarios provided in the study might have biased participants towards expressing this emotional state.

Whether a robot pet should completely adapt its behavior to the user is dependent on the role of the pet. In this study the nature of the bond between the participant and his/ her robot pet was not specified. Some participants argued that a robot pet should mimic a real pet with its own personality and needs which might conflict with the current needs of the user. In contrast, other participants proposed that the robot pet could take the role of therapist/ coach which would focus on the user's needs. According to these participants, such a robot pet should be able to cheer you up, provide comfort, talk about feelings and communicate motivational messages. In the role of a friend the robot pet should also take the user's needs into account, albeit to a lesser extent.

In this study we observed how various people, in this case males and females from the working-age population, interacted with a robot pet companion. However, it should be noted that individual factors such as previous experience with animals, personality, gender, age and nationality might play an important role in these interactions. Interestingly, even though the robot pet's embodiment clearly resembled a dog some participants treated the robot pet as a cat. Whether participants treated the robot pet as a dog or a cat seemed to depend on their preference and history with real pets. Additionally, it should be noted that the participants studied or worked in the computer science department and that all were at least to some extent familiar with social robots. As a result some participants took the current state of technology into account when suggesting possible robot behaviors, for example one participant mentioned that it is nontrivial to build a robot dog that would be able to jump on the couch. The use of a stuffed animal dog as a proxy for a functioning robot pet allowed for a more controlled setup. However, the lack of response from the robot pet resulted in less realistic interactions as the participant had to puppeteer the pet or imagine its response. Furthermore, it is important to note that participants were asked to act as if they were coming home in a particular emotional state. Although this is a common approach in studies on touch behavior (e.g. see [48, 49, 103, 124]) it is unclear whether the same results would have been found if the emotional states were induced in the participants. Despite the above mentioned considerations we observed an interesting range of interactions and were able to find patterns in the social messages that were communicated and the responses that were expected from the robot pet.

4.6 CONCLUSION

To gain more insight into the factors that are relevant to interpret touch within a social context we studied interactions between humans and a robot pet companion in different affective scenarios. The study took place in a contextualized lab setting in which participants acted as if they were coming home in different emotional states (i.e., stressed, depressed, relaxed and excited) without being given specific instructions on the kinds of behaviors that they should display.

Results showed that depending on the emotional state of the user, different social messages were communicated to the robot pet such as expressing one's emotional state, seeking emotional support or enjoying the pet's company. The expected response from the robot pet to these social messages also varied based on the emotional state. Examples of expected

responses were keeping the user company, providing emotional support or picking up on the user's mood. Additionally, the expected response from the robot pet was dependent on the different roles that were envisioned such as a robot that mimics a real pet with its own personality or a robot companion that serves as a therapist/ coach offering emotional support. These findings can help to inform the design of a behavioral model for robot pet companions

Findings from the video observations showed the use of multimodal cues to communicate with the robot pet. Participants often talked to the robot pet while touching it and making eye contact confirming previous findings on the importance of studying touch in multimodal interaction. Furthermore, in contrast to controlled studies in which touch gestures are performed sequentially guided by specific instructions we observed the simultaneous use of multiple touch gestures, touch gestures that were quickly alternated and hybrid forms of prototypical touch gestures. Due to the complexity of the observed interactions, segmentation and labeling of touch gestures proved to be difficult. This finding indicates that the categorization of touch behaviors into discrete touch gesture categories based on dictionary definitions is not a suitable approach to capture the nature of touch behavior in less controlled settings.

5

AFFECTIVE BREATHING BEHAVIOR FOR ROBOT PETS

The following chapter covers collaborative work carried out at the Sensory Perception & Interaction Research Group at the University of British Columbia[1] under the supervision of Prof. Dr. Karon MacLean. The CuddleBit project was led by Paul Bucci and Laura Cang who took the lead in the publication[2] on which this chapter is based. Paul Bucci is the lead designer of the CuddleBit robots. Laura Cang and Merel Jung took the lead in the design and execution of the user study described in this chapter. Jussi Rantala and Merel Jung took the lead in the data analysis of the user study. This chapter covers study 1 as described in Bucci et al. [21]. The content of this part of the paper is restructured to focus on the user evaluation and textual adaptations have been made to embed the study into this dissertation.

A social robot should be able to respond to human touch. The findings of Chapter 4 show that a robot pet might express itself through sounds (e.g. barking or purring), body movements (e.g. tail wagging), certain actions (e.g. laying down to express a relaxed state) or even by talking. In this chapter we will focus on robot pets that express themselves through body movement in the form of different breathing patterns. We will investigate human recognition of different emotional states that are communicated via breathing behavior. In addition, we will explore the influence of robot materiality on the interpretation of these breathing patterns by displaying these behaviors on two distinct robot forms.

5.1 INTRODUCTION

Robot pets can use different social cues to express themselves. For example, elephant-like robot Probo uses facial expressions to communicate different emotions and robotic seal Paro expresses itself using body movements and sounds [95, 100]. In contrast, in this chapter we will focus on the expressive capabilities of less complex 1-DOF robot pets using breathing behavior. As tactile interactions such as holding and stroking are a common way to interact with pets these breathing patterns can be experienced both visually and haptically[3]. In this work we will explore the use of breathing patterns to express different emotional states. Previous research shows that emotions influence a person's respiration pattern [18, 90]. Parameters involved are the depth and rate of breathing and the morphology of the breathing curve such as its regularity [18]. In this study we used these parameters to design breathing behaviors that reflect different emotional states.

In this chapter we will evaluate the expressive potential of one degree of freedom (1-DOF) zoomorphic robot pets collectively called the 'CuddleBits'. We will explore the design and recognition of rendered affective breathing-like behaviors representing four emotional states: stressed, depressed, relaxed and excited. Furthermore, two distinct robot forms will be compared: a rigid wood-based form resembling a rib cage called 'RibBit' and a flexible, plastic-based form resembling a ball of fur called 'FlexiBit'. The CuddleBits were designed using a rapid prototyping approach that allowed quick design and iteration on these robots and their behaviors [21].

To evaluate the emotional expression capabilities of two distinct CuddleBit forms we address the following research questions.

- RQ1: Can 1-DOF robot movements be perceived as communicating different emotional states? *Hypothesis:* Different levels of arousal will be interpreted more accurately than different levels of valence.

- RQ2: How is interpretation of emotional content influenced by robot materiality such as a soft furry texture? *Hypothesis:* FlexiBit's behavior will be perceived as conveying more positive valence than RibBit's.

The remainder of the chapter is structured as follows. Related work on affective robot pets and the emotional expression of breathing behavior will be discussed in Section 5.2 followed by the methods for the presented

---

3 The behavior might also be experienced auditorily depending on the technical implementation of the simulated breathing mechanism

study in Section 5.3. Then, the results will be provided and discussed in Section 5.4 and Section 5.5, respectively. Conclusions will be drawn in Section 5.6.

5.2 RELATED WORK

The expression of emotions through breathing have previously been studied in visual (and haptic) displays. A study on the expression of affect through respiration in virtual humans has shown that the addition of breathing behavior can increase the expressive capabilities of a virtual agent [31]. Emotions such as excitement, fear and anger could be better expressed when using simulated breathing in addition to facial expressions and body posture compared to when only the latter two social cues were used. In other work, Dawson et al. explored lifelike behaviors for a mobile phone by using breathing, ear movements and vibration to express different levels of valence and arousal [30]. Results showed that three levels of arousal could be expressed using different breathing speeds and ear movements for the higher arousal states. In contrast, valence proved to be harder to communicate. The combination of ear wiggling with fast symmetrical breathing could successfully convey an excited state. However, the level of valence could not be communicated reliably for the other emotional states (i.e., angry, neutral, depressed and relaxed).

Moreover, there have been some first attempts to explore the ability of robot pets to express themselves using simulated breathing, examples include the Haptic Creature [97, 123] and the CuddleBot [23]. In a study on the affective display for the Haptic Creature robot, parameters for three levels of arousal and valence were defined to express nine emotional states using simulated breathing, stiffening of the ears and purring [123]. Results showed that arousal could be communicated through breathing rate and ear stiffness whereas valence could not be communicated reliably. Interviews suggested that symmetry and depth of breathing might help to convey different levels of valence. In this study we will focus specifically on breathing behavior by further exploring the expressive space of breathing patterns for a 1-DOF robot pet.

In a controlled follow-up study on the potential calming effect of the Haptic Creature's simulated breathing mechanism, participants were asked to stroke the robot with two hands while it rested on their lap with the robot's breathing either turned on or turned off [97]. Participants' heart and respiration rates were found to decrease significantly as a result of stroking the breathing robot compared to when the robot was not breathing and self reports indicated that participants also felt calmer and

Figure 11: The two CuddleBit forms: FlexiBit (top) and RibBit (bottom) explore distinct form factors.

happier. These findings point to the potential health benefits of equipping robot pets with a breathing mechanism, especially in health-care settings (see also Chapter 6).

5.3 METHODS

We evaluated the expressive capabilities of two distinct CuddleBit forms in a study with a within-subject design. For each robot, participants were asked to rate how well the different breathing behaviors reflected each of four affective states: stressed, depressed, relaxed and excited. As in Chapter 4, these emotional states were chosen because they span opposite ends

of the valence and arousal scale: stressed (negative valence, high arousal), depressed (negative valence, low arousal), relaxed (positive valence, low arousal) and excited (positive valence, high arousal) [93].

5.3.1 *Participants*

In total 20 participants (11 male, 8 female, 1 identified otherwise) took part in the study. The age of the participants ranged from 22 to 36 (M = 25.4; SD = 3.6) and their cultural backgrounds were from North America, Europe, Southeast Asia, Middle East and Africa. All participants had completed at least an undergraduate degree and were compensated 5 CAD for their time. The study was conducted under UBC Ethics H15-02611.

5.3.2 *Apparatus/ Materials*

5.3.2.1 *CuddleBit robots*

In this study the emotional expression capabilities of two distinct CuddleBit forms were compared (see Figure 11). The FlexiBit is made out of plastic slices fixed to a base at the bottom and joined together at the top with machine screws. Plastic flexibility, volume and curvature provide passive compliance and the robot is actuated by a servomotor controlled by an Arduino Uno microcontroller board[4]. The shape of the robot can be adjusted by varying the sizes of the slices and/or the base. The outside is covered with faux fur and the robot is often compared to a Tribble (a fuzzy alien species from the Star Trek universe).

The RibBit is a wooden rib cage on a stand and its parts are assembled using glue and BBQ skewers as pins and rods. Internal springs provide compliance to the rigid actuation by an Arduino-controlled servomotor. The shape of the robot can be adjusted by modifying and laser cutting digital patterns. Contrary to FlexiBit, RibBit is rigid due to its wooden construction and its mechanics are fully exposed. For more details on the design of the CuddleBits the reader is referred to Bucci et al. [21] and Cang et al. [22].

5.3.2.2 *Design and selection of breathing behaviors*

Breathing behaviors were designed by extending a keyframe-based editor for vibrotactile sensations (i.e., Macaron [96]) to robot position control, the result was named 'MacaronBit'. To design different breathing patterns a user starts with a pure sine wave and can adjust its parameters (i.e.,

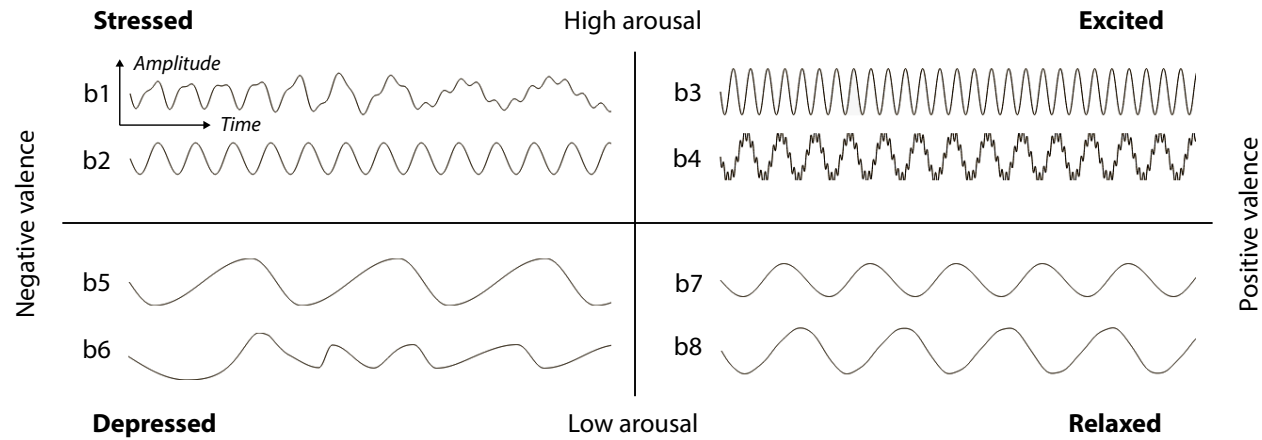---

4 http://arduino.cc/en/Main/ArduinoBoardUno

Figure 12: Researcher-designed waveforms for expressing four emotional states at the opposite ends of the valence and arousal scale [93].

frequency, amplitude, bias and amplitude/frequency variability). The editor supports immediate playback, key-framing (parameter interpolation between key points), waveform generation and click-and-drag editing.

Using MacaronBit, three members of the research team individually created a total of 18 behaviors for the FlexiBit robot to express the four different emotional states. These behaviors were then matched to one of the emotional states by five team members. The two most agreed-upon behaviors for each of the emotional states were selected resulting in a final selection of eight waveforms (see Figure 12). The same eight behaviors were implemented for the RibBit robot with amplitude scaled by range of motion.

5.3.3   *Procedure*

After giving informed consent participants were seated and invited to inspect the inanimate robots. Then participants filled out their demographic information. In the main part of the study participants were given the task of rating each behavior on a 5-point semantic differential scale ranging from '−2 Mismatch' to '+2 Match' for two different robots displaying four emotions: stressed, depressed, relaxed and excited. For instance, for 'FlexiBit feels stressed', a participant would play each behavior and rate how well it matched the robot portraying stress. During playback and rating, participants kept one hand on the robot, and moused with the other. As a result motion was experienced largely haptically. Noise-canceling headphones played pink noise to mask mechanical noises. Instructions were communicated by microphone. Ratings for each robot were performed separately. Robot block order was counterbalanced, with a 2 minutes break in between. For each block, all four emotions were presented on the same screen so participants could compare globally. Behaviors (15 s clips) could be played at will during the block. Order of behaviors and emotion was randomized between participants. To reduce cognitive load, participants saw the same behavior/emotion order for the second block. In total, each participant performed 64 ratings (8 behaviors × 4 emotions × 2 robots). In addition to the behavior ratings, we also recorded the time it took to complete the ratings in each of the emotional states which was estimated by duration of mouse focus within each emotion rating sub task. Afterwards, a semi-structured interview was conducted about possible roles that a CuddleBit could play in the life of the participants and how their behaviors were perceived. Each session took approximately 30 minutes.
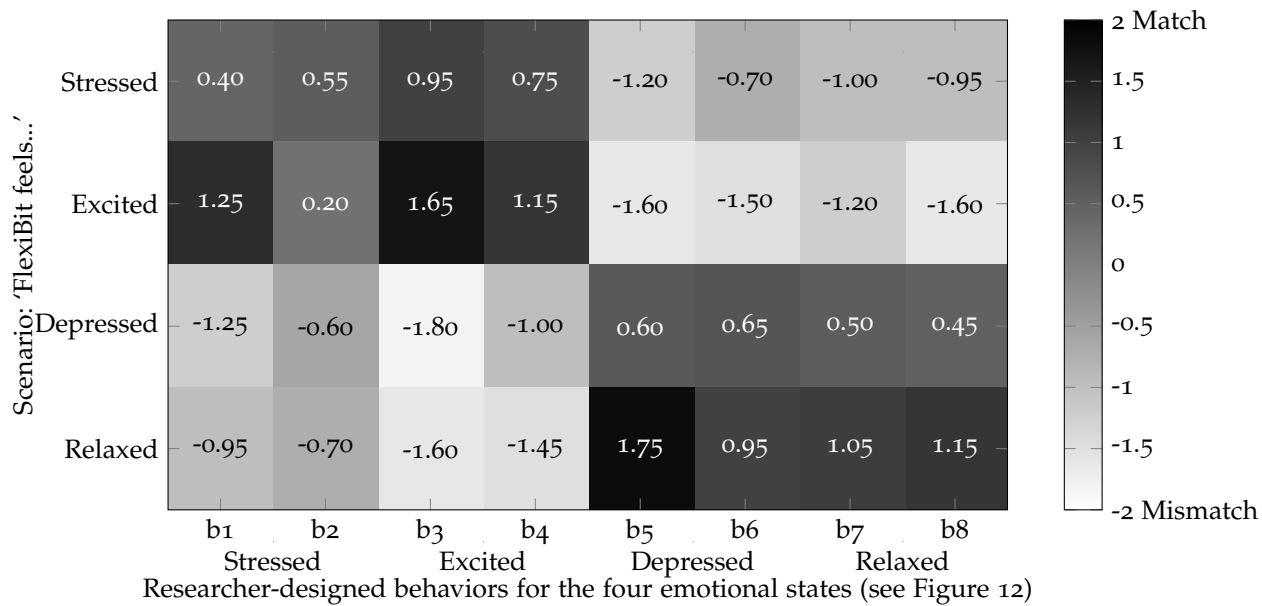
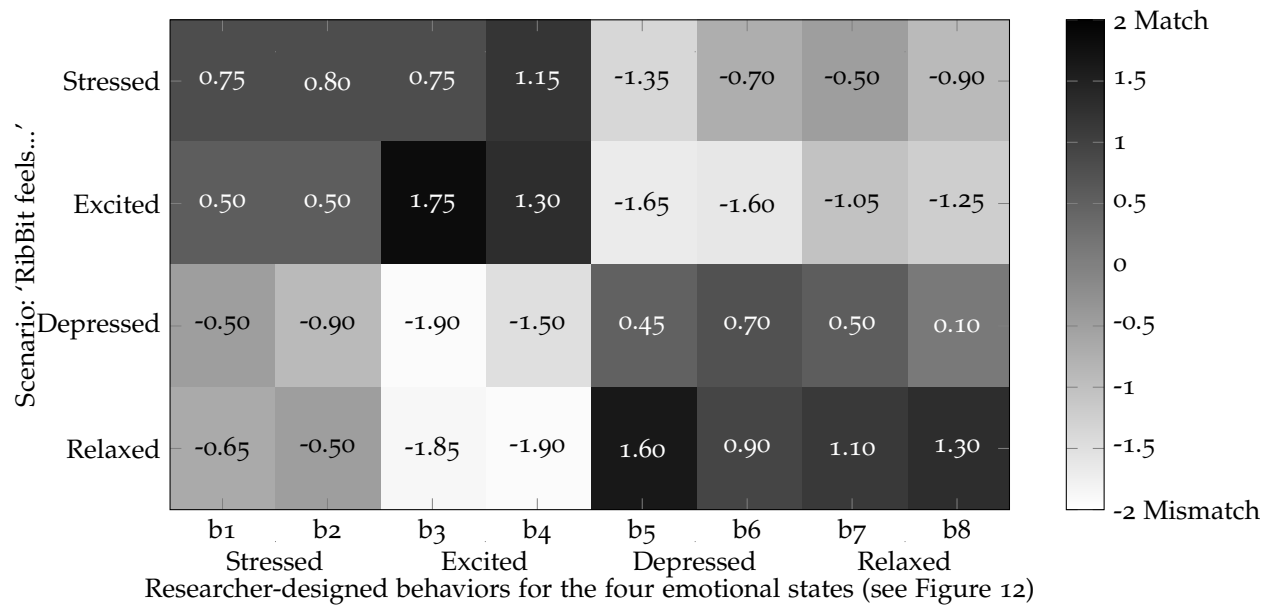Figure 13: Mean behavior ratings for FlexiBit per scenario.

Figure 14: Mean behavior ratings for RibBit per scenario.

5.4   RESULTS

Visual inspection of the mean behavior ratings across all participants showed that the behaviors matched relatively well with the scenarios that they were designed for (see Figure 13 and 14). Interestingly, the behaviors also matched the scenario in which the robot's emotional state reflected the same level of arousal but a different level of valence (i.e., stressed-excited and depressed-relaxed).

We compared the ratings of each pair of behaviors designed for the same emotional state for both CuddleBits (i.e., 4 behavior pairs × 2 robots) using a pairwise Wilcoxon signed-rank tests with Bonferroni correction ($\alpha = .05/8 = .006$). Ratings for the two designed behaviors for the same emotion state were not significantly different (all p's $\geqslant$ .059) Therefore, we averaged the ratings of each pair (e.g. b1 and b2 in Figure 12) for further analysis.

Next, we explored the effect of the scenario that was given to participants (e.g. 'RibBit feels excited') on the ratings of the designed behaviors. Friedman's test on behavior ratings showed significant differences between the behavior ratings per scenario for each robot (all p's $<$ .001). Post hoc analyses using Wilcoxon signed-rank tests with a Bonferroni correction ($\alpha = .05/6 = .008$) were conducted to further analyze the effect of scenario on researcher-designed behaviors (see Table 17). In the stressed, excited and relaxed scenarios significant differences were found between high and low arousal behaviors for both robots (i.e., stressed-depressed, stressed-relaxed, excited-depressed and excited-relaxed, all p's $\leqslant$ .002). No significant differences were found between behaviors with the same arousal level but different valence content (all p's $\geqslant$ .017). The results for the depressed scenario show similar patterns with some exceptions. No significant differences were found with either robot when comparing ratings of relaxed-stressed (p's $\geqslant$ .014). In addition, with RibBit, no significant differences were found for the ratings of depressed-stressed (p = .012). However, comparison of excited-stressed revealed a significant difference (p = .007).

Furthermore, the behavior ratings for FlexiBit were compared to those for RibBit for each scenario (i.e., 4 behavior ratings × 4 scenarios) using a Wilcoxon signed-rank tests with Bonferroni correction. No statistically significant differences were found between ratings of emotions displayed on the two distinct robot forms ($\alpha = .05/16 = .003$; all p's $\geqslant$ .026).

Additionally, we approximated task difficulty using the time spent to evaluate the behaviors. Spending more time would suggest challenge in aligning robot behaviors with emotion. A two-way (2 robots × 4 scenarios)

repeated measures ANOVA showed no significant main and interaction effects for the time spent on rating behaviors ($\alpha = .05$; all p's $\geqslant .198$).

Table 17: Pairwise comparison p-values (Wilcoxon) of behavior ratings (rows) for different scenarios (columns) per robot. Significant differences (i.e., p $\leqslant$ .008) are marked gray.

| | Scenario: robot feels... | | | |
	stressed	excited	depressed	relaxed
FlexiBit				
S-E	.391	.142	.159	.076
S-D	.001	.000	.004	.000
S-R	.001	.000	.014	.001
E-D	.000	.000	.000	.000
E-R	.000	.000	.000	.000
D-R	1.000	.713	.668	.501
RibBit				
S-E	.759	.037	.007	.017
S-D	.000	.000	.012	.000
S-R	.001	.002	.032	.001
E-D	.000	.000	.000	.000
E-R	.000	.000	.000	.000
D-R	.582	.270	.713	.668

5.5 DISCUSSION

5.5.1 *Recognition of the robot's emotional state*

Generally, participants were able to perceive differences in behaviors designed to convey high or low arousal. Speed or frequency was mostly mentioned as having communicated arousal variation: low arousal from low frequency and high arousal from high frequency. In contrast, the interpretation of the robot's expressed level of valence was less straightforward. Possible reasons are that breathing might not naturally convey valence variations and/or additional DOFs are needed to disambiguate them. It is less likely that materiality played a role because of the consistency in the findings between the two distinct prototypes. These results confirm that these 1-DOF robots were able to reproduce earlier findings regard-

ing both affective dimensions [30, 123]. Moreover, these findings support our hypothesis that different levels of arousal are easier to interpret than different levels of valence.

Unexpectedly, the behavior ratings for the depressed scenario diverged significantly from those in the other scenarios. Interviews suggest two reasons. Depressed was being confused with stressed that is, participants reported experiencing both emotions in concert or as a result of the other. Furthermore, breathing (by RibBit in particular) did not have the ability to express depression for some participants. Suggestions to improve the believability and differentiability for expressing a stressed state included sighing and avoidance actions like retreating or turning away. It should be noted that the breathing behaviors were designed based on the perception and imagination of three computer science researchers which may not broadly generalize as an effective emotional display.

The results from this study are consistent with previous studies on the recognition of human emotions based on other social cues such as facial expressions. As an example, a soccer player yells after a goal: without knowing which side the soccer player is on, it is difficult to visually distinguish between a yell of anguish or victory [10]. Similarly, for low arousal states, it might be difficult to tell the difference between someone who is relaxed or depressed. Observers may always need to rely on context to interpret the level of valence, either through extended interaction or through external environmental and situational cues. The design of behavior to express the level of arousal might then be more important, especially if interaction and contextual cues are stronger than any inherent behavior features. Moreover, a robot pet with multiple degrees of freedom might be more capable to express valence through other forms of actuation. For example, a robot pet could raise its head or wag its tail to express positive valence.

5.5.2   *Effect of robot form*

There was no difference in how participants perceived the behaviors on the two distinct robot forms. In post-study interviews participants reported that the movement expressed by the two robot forms were sensorially different: FlexiBit's fur felt nicer to touch but RibBit's motion was perceived to be more precise. Notably, RibBit's movements were interpreted as breathing or a heartbeat (i.e., biological functions) despite the exposed inner workings which emphasize the mechanical aspects of the robot. These findings indicate that movement rather than materiality dominated how participants interpreted the expression of the emotions.

Therefore, our hypothesis that FlexiBit's behavior will be perceived as conveying more positive valence than RibBit's is not supported. The observed low impact of robot form on the expression of emotional state indicates that behaviors might be reusable on physically distinct robots with similar abstract abilities such as breathing.

5.6 CONCLUSION

In this chapter we evaluated the expressive potential of breathing behaviors for 1-DOF zoomorphic robots. We investigated the extent to which researcher-designed emotional breathing behaviors could communicate four different affective states. Additionally, we were interested in the influence of robot form on the interpretation of these breathing behaviors. For this reason two distinct robot forms were compared: a rigid wood-based form resembling a rib cage called 'RibBit' and a flexible, plastic-based form resembling a ball of fur called 'FlexiBit'. In the study participants rated for each robot how well the different breathing behaviors reflected each of four affective states: stressed, depressed, relaxed and excited.

We started out with two hypotheses. Firstly, we hypothesized that different levels of arousal would be communicated more accurately than different levels of valence. Secondly, we presumed that robot materiality would influence how the breathing patterns were perceived. To this end we hypothesized that FlexiBit's behavior would be perceived as conveying more positive valence than RibBit's. In accordance with our first hypothesis the results showed that both robot forms were able to express high and low arousal states through breathing behavior, whereas valence could not be expressed reliably. Low arousal states could be communicated by low frequency breathing behavior and higher frequency breathing conveyed high arousal. In contrast, context might play a more important role in the interpretation of different levels of valence. Contrary to our second hypothesis robot form did not influence the perception of the behavior that was expressed. These findings can help to inform future design of affective behavior for robot pet companions.

6

TOUCH INTERACTION WITH ROBOT PETS IN A HEALTH-CARE SETTING

The following chapter covers collaborative work which is based on the master's thesis of Lisa van der Leij [111]. The study was designed by Lisa van der Leij with assistance from Merel Jung and Saskia Kelders. Lisa van der Leij collected the data which was analyzed by Lisa van der Leij and Merel Jung with assistance from Saskia Kelders. Merel Jung took the lead in the writing of a paper based on the work described in the master's thesis with contributions from Lisa van der Leij and Saskia Kelders. The content of this chapter is identical to that of the submitted paper[1] with some minor textual adaptations to embed the content into this dissertation. The future work described in the paper has been moved to Chapter 7 of the dissertation.

In the previous two chapters we focused on how a robot pet could understand and respond to human touch (Chapter 4) and how robot pets could express their emotions through breathing behavior (Chapter 5). As a next step we will explore the use of touch in interaction with a fully functioning robot pet outside of the lab. Robot pet companions such as robotic seal Paro are increasingly used in care for the elderly due to the positive effects that interaction with these robots can have on the well-being of patients with dementia. As touch is one of the most important interaction modalities for patients with dementia this can be a natural way to interact with these robots. In this chapter we will present a study that aims to inform the development of robot pets that can understand and respond to human touch for health care applications. In the study we administered questionnaires and conducted interviews with health care providers in dementia care to explore in what ways people with dementia could benefit from interaction with a robot pet with more advanced touch recognition

---

1 Jung, M. M., van der Leij, L., and Kelders, S. M., An exploration of the benefits of an animallike robot companion with more advanced touch interaction capabilities for dementia care, Manuscript submitted for publication

capabilities, what types of (tactile) behaviors patients with dementia do/ would use in their interactions with Paro and which other target groups could benefit from interaction with Paro or a more advanced robot pet.

6.1   INTRODUCTION

Social robots such as robot pet companions are increasingly being used in health care (e.g. [12, 15, 67, 100, 117]). Research shows that interaction with robot pets can have positive effects similar to those from interaction with real animals (e.g. [8, 16, 32, 92, 97, 117]). Moreover, robots have advantages compared to real animals: their behavior can be better controlled; they do not carry diseases; and they need less care [32, 35, 100]. These advantages are especially important when working with a vulnerable population such as the elderly. When interacting with real animals, touch is one of the most used forms of interaction. Furthermore, recent studies point to the benefits of touch interaction in robot therapy [92, 97]. However, currently commercially available companion robots such as robot seal Paro [100] and Sony's robotic dog AIBO [39] do not focus specifically on touch interaction, which seems like a missed opportunity.

Paro is one of the most researched robotic animals and one that is quite often implemented in regular care [12, 67, 92, 100]. It was developed in Japan for robot therapy and is mostly used with children and people with dementia [13, 14, 100, 117, 118, 119]. Paro is equipped with touch sensors but does not recognize or interpret different touch gestures [117]. Instead, the robot distinguishes between positive (i.e., soft) and negative (i.e., rough) touches. However, research indicates that people use mostly positive forms of touch when interacting with another human [59] or a robot pet [124]. Furthermore, these positive forms of touch can have different meanings depending on the context, for example, the intent of a touch could be affectionate, comforting/supportive or playful [59, 64, 124]. In addition, the results from the study presented in Chapter 4 showed that the social message that was communicated to a robot pet varied based on the user's emotional state which also affected the response that was expected from the pet [64]. Based on these findings it is therefore likely that recognizing and acting on these different meanings will benefit interaction with robot pets.

The focus of this chapter will be on the use of touch in the interaction with robot pet companions. In this study we will consult health care providers to explore whether people with dementia would benefit from interaction with a robot pet with more advanced touch recognition capabilities. Dementia is a global, progressive and chronic condition, in which

there are severe impairments in a person's ability to think, reason and remember [121]. Dementia has been recognized by the WHO as a public health care priority because of the increase in the number of people living with dementia due to the aging population, and the enormous burden this places on the health care system [122]. To keep the costs of dementia care manageable and the quality of care high, innovative solutions are needed. The use of (animallike) robots is one of these solutions [12, 67]. Additionally, enabling these robot pets to understand and respond to human touch might be a way to further improve the effects of these robots in dementia care as touch is one of the most important interaction modalities in patients with dementia. As the disease progresses, verbal communication becomes harder for these patients and nonverbal interaction, especially touch, plays a prominent role in care, for instance, for communicating messages of comfort and safety [9, 89].

In this study we administered questionnaires and conducted interviews with two groups of health care providers in dementia care: a group that worked with Paro and a group that had no experience with the use of robot pet companions. We will address the following research questions to inform the development of robot pet companions that can understand and respond to human touch for health applications. (1) In what ways could people with dementia benefit from interaction with a robot pet with more advanced touch recognition capabilities? (2) What types of (tactile) behaviors do/ would patients with dementia use in their interactions with Paro? (3) Which other target groups could benefit from interaction with Paro or a more advanced robot pet?

The remainder of this chapter is structured as follows. Related work on the effectiveness of robot pet companions in care for the elderly and touch interaction with robot pets will be discussed in Section 6.2. Then, the material and methods for the presented study will be described in Section 6.3. The results will be presented and discussed in Section 6.4 and Section 6.5, respectively. Conclusions will be drawn in Section 6.6.

6.2 RELATED WORK

6.2.1 *Effectiveness of robot pet companions in care for the elderly*

There are potential health benefits to be gained from interacting with a robot pet companion. For example, a study by Banks et al. found that the company of robotic dog AIBO could be as effective as a real dog in reducing loneliness in elderly patients living in a long-term care facility [8]. Furthermore, there are indications that stroking and interacting with

Paro could lower blood pressure in elderly people which is similar to the effects found for interaction with real animals [32, 92]. In another body of work, elderly people (including people in various stages of dementia) in different care facilities interacted freely with Paro for either a couple of weeks [117, 119] or up to one year [118]. The results of these studies indicated that interaction with the robot pet could improve mood, make people more active, lower stress and promote social contact with the robot as well as with peers and nursing staff [117, 118, 119]. Furthermore, the use of Paro seemed to ease the burden on the nursing staff as their reported stress levels decreased after the robot's introduction [119]. In addition, the results showed that Paro is interesting enough for long-term interaction and proved to be durable and safe enough for long-term use [118].

Systematic reviews into the effectiveness of socially assistive robots in care for the elderly found that these robots (most studies investigated robots that were animallike) have the potential to improve psychological and physiological outcomes, but the methodological quality of the existing studies is low [12, 67]. Bemelmans et al. stress the need for structured interventions, similar to those used in animal-assisted therapy [35], with measurable outcomes as without proof of the added value of robots for therapy their image of being mere entertaining gadgets might remain and reimbursement could be problematic [12, 13, 15]. Furthermore, end-users (i.e., patients and care providers) should be included in the process to ensure successful adoption of robot therapy [12, 15, 67].

In an effort to study the effects of robot pets in health care in a more structured manner, several interventions for Paro in dementia care were developed together with care providers such as nurses, activity coordinators, therapists and medical doctors [15]. These interventions could be divided into three types: (1) therapeutic applications to stimulate perception, psychological functioning, psychosocial well-being and social behavior; (2) facilitation of daily care activities by providing comfort and distraction during stressful activities; (3) support of social visits by having a shared focus point as Paro can attract attention. Based on these three types, individualized interventions were defined and tested with patients suffering from dementia in two studies [13, 14]. 69 therapeutic interventions and 17 care facilitating interventions were conducted in a within-subject quasi-experimental time series ABAB study lasting four months which was completed by 71 participants [13]. Overall, the interventions showed a significantly positive effect indicating that Paro can be a valuable tool in dementia care. In the other study 23 participants interacted with Paro once or twice a week for a period of 3 weeks [14]. Of the 35 conducted interventions 19 were therapeutic interventions, 7 were to facilitate care and

9 were to support social visits. Care providers considered the majority of the interventions to be feasible (26 out of 35) and of added value (22 out of 35). Both studies showed that Paro was most suitable for therapeutic interventions [13, 14].

6.2.2 *Touch interaction with robot pet companions*

Compared to hard-shelled robot pets such as AIBO [39], soft fur-covered robots such as Paro and the elephant-like robot Probo are more pleasant to touch and can evoke affective behaviors such as stroking and hugging [92, 95, 100, 117]. For a more natural interaction these robot pets should be able to understand and respond to these different types of touch. Recently, research labs have started to develop robot pets that focus specifically on touch interaction. For instance, the Huggable robotic teddy bear has its own full-body somatosensory system in order to sense and process human touch [73, 106]. Another example is the Haptic Creature which is a zoomorphic lap pet that can sense human touch all over its body and expresses itself by purring, simulated breathing and stiffening of the ears [97, 123, 124]. The use of simulated breathing for robot pets has been explored further in the CuddleBit robots [21, 22] (see also Chapter 5). These small zoomorphic robots react to human touch by expressing emotions through different breathing patterns resulting in a haptic affective display.

The study presented in this chapter will build upon previous research on the effectiveness of existing robot pets such as Paro in health care facilities and exploratory lab research on the development of new robots that can engage in tactile interaction. In contrast to most previous research, the purpose of this study is not to investigate the effectiveness of Paro in dementia care per se. Instead, our focus is on how people with dementia could benefit from interaction with an animallike robot companion that is able to understand and respond to different types of touch.

6.3 MATERIAL AND METHODS

6.3.1 *Study design*

For this study, we opted to recruit health care providers who worked in geriatric psychiatry departments. As these health care providers work with patients suffering from dementia on a daily basis we expected that they would have a lot of insight into the needs of people with dementia and in what ways these patients might benefit from interaction with a robot pet with more advanced touch recognition capabilities. The perspec-

tives and expectations of a sample of health care providers on the use of a robot pet with more advanced interaction capabilities in dementia care and health care in general were explored through interviews. In addition, in a questionnaire the health care providers were asked to assess the likelihood that people with dementia would use different touch gestures to interact with a robot pet. Paro was used as the main example of a robot pet because the robot seal is already used in Dutch care facilities.

6.3.2   Participants

In total 9 health care providers from two care facilities in the eastern part of the Netherlands were recruited to participate in the study. One group of health care providers did not have any experience with the use of robot pets (the layman group, $n = 4$). These participants were recruited by one of the members of the research team during coffee breaks. The other group of health care providers worked at a different care facility and these participants did have experience using Paro (the expert group, $n = 5$). In this case participants were recruited via a contact person at the care facility. At the time of the interviews, Paro had been available to the expert group for one year. Their amount of experience with Paro differed because they cared for different patients with different needs. That is, some of the experts used Paro every day (e.g. as part of a bedtime routine for a patient who has difficulty sleeping) and others used Paro more incidentally (e.g. as a means to calm down a patient who is displaying vocally disruptive behaviour). All experts had experience in using Paro for multiple goals and with multiple patients. All participants were females who had completed secondary vocational training. The age of the participants was between 30 and 52 ($M = 44, SD = 10$) for the layman group and between 24 and 67 ($M = 46, SD = 15$) for the expert group. Participants from the layman group had more years of experience working in a geriatric psychiatry department ($M = 13, SD = 7$) compared to the expert group ($M = 5, SD = 2$). The study was approved by the ethics committee of the Faculty of Behavioural, Management and Social Sciences of the University of Twente.

6.3.3   Materials

6.3.3.1   Interviews

Semi-structured interviews were conducted based on a predefined framework. Two versions of the interview were prepared to adjust the questions

to ask about either expectations (layman group) or experiences (expert group). In order to compare the expectations of the layman group with the experiences of the expert group, the layman group was introduced to Paro by means of a short video fragment[2] at the beginning of the interview. Four main topics were explored within both versions of the interview:

1. Vision on the use of robot pets in health care: what are suitable target groups and for what kind of interventions can robot pets be used? In addition, the expert group was also asked about Paro's advantages and disadvantages.

2. Expectations of (tactile) interaction capabilities of robot pets: which emotions or other social messages should a robot pet communicate to a person with dementia and how should a robot communicate these? Additionally, the expert group was asked which social cues Paro conveys in reaction to being touched by people with dementia.

3. Added value of a robot pet with more advanced capabilities to understand and respond to social touch for patients suffering from dementia: would a more socially intelligent robot pet be more effective and would the advancements in the robot's capabilities affect how the robot could be used? In addition, the expert group was also asked whether Paro's interaction capabilities were perceived to be sufficient.

4. Other contexts in which robot pets with more advanced capabilities to understand and respond to social touch could be used: which other target groups could benefit from interaction with a more socially intelligent robot pet and how could the robot be used for these new target groups?

6.3.3.2 *Questionnaire*

In the questionnaire, participants were asked to rate the likelihood that people with dementia would use different touch gestures in their interaction with a robot pet. 30 different touch gestures were rated on a 5-point Likert scale from 1 (unlikely) to 5 (likely). The list of different touch gestures was based on the 'touch dictionary' from Yohanan and MacLean which consists of 30 different touch gesture labels and their definitions [124]. For this study the touch gestures and definitions were translated to Dutch by the members of the research team.

---

2 http://youtube.com/watch?v=-fkxdFwu8yE

6.3.4   *Procedure*

The participants were welcomed by the interviewer in a quiet office within the care facility where the participants worked. For practical reasons the interviews with the participants from the expert group were conducted in two groups consisting of either two or three participants while the interviews with the participants from the layman group were conducted individually. Participants were informed about the nature of the interview and the questionnaire before signing the informed consent form. Then, the interview was conducted according to the predefined framework. The interview started by collecting demographic information followed by a video fragment of Paro which was only shown to the layman group. During the interview follow-up questions were asked where necessary. After the interview participants completed the questionnaire. The interviewer remained present in case clarification was needed on the questionnaire. However, none of the participants needed any help. The total duration of each session was approximately 30 minutes. Afterwards the participants received a small token of appreciation.

6.3.5   *Data analysis*

6.3.5.1   *Interviews*

The interviews were analyzed following the guidelines from Baarda et al. [5]. Transcriptions of the interviews were grouped based on topics using an inductive approach. The interview data was divided into fragments and each fragment was assigned a label that described the content. Multiple labels could be assigned to a fragment if the fragment could not be split without losing context. For each label the number of participants per group that mentioned the specific topic was counted and these numbers are indicated in the results. In the results section we will highlight the differences between groups or combine the results of both groups in cases where results were similar. The interview results are presented within themes that do not necessarily overlap with the four main topics of interview question listed in Section 6.3.3.1.

6.3.5.2   *Questionnaire*

The questionnaire data was analyzed using IBM SPSS Statistics version 22. The focus of the questionnaire was on the likelihood of the use of different touch gestures to interact with Paro. However, due to the inclusion of two groups of participants, the ratings of the layman group and the expert

group were first compared by conducting Mann-Whitney U tests using Bonferroni adjusted alpha levels of .0017 per test (.05/30). Exact p-values are reported for two-tailed tests. Touch gestures with a median rating of > 3 were considered likely to be used, gestures with a rating of 3 were considered neutral while those with a rating of < 3 were considered unlikely to be used. Additionally, within each of these three categories the touch gestures were ranked based on the summed likelihood ratings of all participants.

6.4 RESULTS

6.4.1 *Interviews*

6.4.1.1 *Usages of Paro in dementia care*

Participants mentioned six different goals for which Paro can/ could be used (see also Table 18).

1. All Participants perceived Paro to be especially suitable as a means to distract patients who are restless or sad (layman: $n = 4$; experts: $n = 5$). Restless behavior was reported to often occur during the evening (layman: $n = 2$; experts: $n = 3$).

2. About half of the participants mentioned that Paro can be used to interrupt problematic behaviors (layman: $n = 3$; experts: $n = 2$), especially in cases of vocally disruptive behavior (layman: $n = 2$; experts: $n = 1$). A participant from the expert group explained that Paro can provide stimulation to confirm to the patient that he/she is still alive.

3. Five of the participants mentioned that Paro can be used to make contact with patients with severe dementia by stimulating their senses through interaction (layman: $n = 3$; experts: $n = 2$).

4. About half of the participants stated that Paro can stimulate communication both between patients and between patients and health care providers (layman: $n = 3$; experts: $n = 2$).

5. Some of the participants mentioned that Paro can be used to support health care providers (layman: $n = 2$; experts: $n = 2$). For example, participants from the expert group mentioned that the soothing effect of Paro can facilitate care moments ($n = 1$) and that Paro can provide one-on-one contact at moments when the care staff does not have the time or the manpower available to provide this individual

level of contact (n = 1). The layman group also expected that Paro would support them in their work but they attributed this to the reduced need for supervision (n = 2). In contrast, participants from the expert group argued that interaction with Paro should be supervised to prevent patients from damaging Paro (n = 2) and to observe the outcome of the intervention (n = 3).

6. A few of the participants argued that Paro can be used as company to relieve feelings of loneliness (layman: n = 1; experts: n = 1). However, another participant from the layman group (n = 1) argued that she would not use Paro in this case as a robot could not replace human contact.

Table 18: Goals for which Paro can/ could be used according to the health care providers. The number of participants that mentioned each goal are listed per group.

Goal	No. of experts	No. of layman
Distraction	4	5
Interrupt problematic behavior	3	2
Make contact	3	2
Stimulate communication	3	2
Support care providers	2	2
Relieve feelings of loneliness	1	1(+)/1(-)

6.4.1.2 *(Touch) interaction between Paro and people with dementia*

Some participants mentioned that the mere presence of Paro can already have a positive effect on people with dementia (experts: n = 2). However, the others argued that feedback from Paro is essential to achieve the desirable effect (layman: n = 4; experts: n = 3). Furthermore, according to some participants Paro's response should be adapted to the patient's touch (layman: n = 3). Participants indicated that Paro's soft fur (layman: n = 2; experts: n = 2), big eyes (layman: n = 2; experts: n = 2) and auditory signals (layman: n = 1; experts: n = 2) seem to elicit care responses in the patients who in return express their love to Paro and comfort the robot (layman: n = 1; experts: n = 2). Patients display their affectionate behavior in the form of touch gestures such as stroking and hugging (layman:

n = 2; experts: n = 4) and verbal expressions (experts: n = 2). According to the participants, Paro mirrors the positive interaction by conveying safety and security (layman: n = 2; experts: n = 5) and love (layman: n = 1; experts: n = 1) in the form of an auditory response (layman: n = 1; experts: n = 4) and by moving its head and fins (layman: n = 2; experts: n = 3). One of the participants from the layman group mentioned that Paro's response subsequently seems to have a positive effect on the mood of the patients.

A simulated heartbeat was suggested by some of the participants as a valuable additional communication channel for robot pets which could at the same time also have a soothing effect on the patients (layman: n = 1; experts: n = 1). Currently, Paro's tactile feedback is mostly conveyed through body movements as described by one of the participants in the expert group:

> "I notice when they [the patients] hold him [Paro] against their neck, he lifts his head and as a result the whiskers move along their faces which is a very sensitive area for these people, they can feel it clearly".

Additionally, two of the participants mentioned that patients are able to recognize Paro's negative response to aggressive touch gestures such as rough grabbing (experts: n = 2). For example, one of these experts recalled that one of her patients had exclaimed: "oh he [Paro] doesn't like that" in response to Paro's negative reaction. In response to these comments, some other participants complemented that the use of rough touch occurs rarely (experts: n = 2).

6.4.1.3 *Suitability of Paro for people with dementia*

In general, all health care providers were positive about the use of robot pets in dementia care. Paro was described by participants of the layman group as a promising (n = 3) and easy-to-use intervention (n = 1). Furthermore, few participants of the layman group argued that Paro could be a low maintenance (n = 1), more robust (n = 1) and more interactive (n = 1) alternative to real animals The expert group was particularly positive about the effects that Paro has on people with dementia (n = 3). However, it was mentioned that robot pets might not be a suitable solution for every person with dementia. Seven of the participants indicated that such a robot should fit within the patients' perception of their environment and whether the patients will interact with a robot pet is dependent on their affinity for animals (layman: n = 3; experts: n = 4). Moreover, some of the participants argued that the use of a robot intervention should be

discussed with the patient's family (layman: $n = 2$; experts: $n = 1$) and health care providers should be trained to use these interventions in an effective and respectful manner (experts: $n = 2$).

Paro's shortcomings also came to light during the interviews. Some of the participants in the layman group argued that a robotic cat or dog might be more suitable than Paro's seal appearance as patients might be more familiar with these animals and would feel more safe ($n = 2$). Also, it was mentioned by one participant that these familiar animals might elicit more reactions from the patients (layman: $n = 1$). Interestingly, one of the participants in the expert group stated that the form of the robot pet does not matter as long as the robot has similar functionality as Paro. Moreover, two participants from the expert group were in general positive about Paro's endearing appearance although, at first, they were skeptical about the seal-like design as well. Additionally, some remarks were made about the auditory responses of Paro which do not seem to be the best fit for this target group. Two participants from the expert group described the responses as repetitive, irritating, too loud and too high-pitched which can overstimulate the patients, especially in a group setting ($n = 2$). Also, the feel of Paro was discussed. Paro was perceived to be not pliable enough by some of the participants (experts: $n = 3$) especially on its underside as the hard internal structure can be felt through the fur. Moreover, Paro's body was described as being too rigid.

However, in spite of the aforementioned shortcomings Paro was described as being a sufficiently effective intervention for people suffering from dementia. When asked, almost all participants answered that Paro does not need more advanced tactile interaction capabilities (layman: $n = 3$; experts: $n = 5$). Some participants argued that additional stimuli might overstimulate the patients (layman: $n = 1$; experts: $n = 2$).

6.4.1.4   *Suitability of Paro for other target groups*

The most frequently mentioned alternative target groups for which Paro could be a suitable intervention were people with an intellectual disability (layman: $n = 3$; experts: $n = 2$) and children (layman: $n = 2$; experts: $n = 2$). Subgroups that were specifically mentioned by some of the participants were autistic children (experts: $n = 1$) and those in hospitals (layman: $n = 1$; experts: $n = 2$). In the latter case it was argued that Paro could provide comfort. Notably, these groups either have reduced cognitive capabilities or their cognitive abilities are still under development. This commonality is understandable as Paro was originally designed for people with dementia. Indeed, almost all participants stated that Paro might be too childish and too simple for adults with normal cognitive

health (layman: $n = 4$; experts: $n = 4$). Even within the target group of people with dementia, Paro was perceived by most participants to be especially suitable for patients with severe dementia (layman: $n = 2$; experts: $n = 4$). In addition, some participants from the expert group reported that Paro seems to be more effective for these patients compared to patients with mild dementia ($n = 2$). In contrast, one participant from the expert group had positive experience with the usage of Paro for a patient with a chronic physical illness who had normal cognitive abilities. Additionally, some participants argued that Paro might be effective for patients with psychiatric problems as well, for example, to calm down restless patients or to reduce aggressive behavior (experts: $n = 2$).

Participants indicated that robot pets should be more technologically advanced in order to be suitable for most other target groups. Examples of target groups for more advanced robot pets that were given by individual participants were robot pets for people with reduced mobility who are unable to care for a real animal (layman: $n = 1$) or in rehabilitation where people could exercise together with a robotic dog (layman: $n = 1$). Another potential target group that was mentioned were healthy elderly people that still live independently: in this case a robot could help to prevent loneliness (experts: $n = 1$). However, this participant mentioned that the current price of Paro is a limiting factor for the wide adoption of robot companions for personal use. Moreover, it was argued that a more advanced robot could improve mental health by starting conversations to help people open up about repressed emotions (layman: $n = 1$). Two participants mentioned that in general, as with all target groups, it will depend on the person whether he/she will benefit from interaction with a robot companion (layman: $n = 2$).

6.4.2 *Questionnaire*

The likelihood ratings for the use of different touch gestures did not differ significantly between the expert and layman groups (all p's $\geqslant$ .048). The data of both groups was therefore combined and divided into three categories based on the median ratings of the likelihood that people with dementia would use the touch gestures in their interaction with a robot pet (see Table 19). The touch gestures are ranked within each column based on the summed ratings of all participants which ranged from 9 to 45 as 9 participants rated each touch gesture on a scale from 1 to 5. Notably, the touch gestures that were rated to be most likely (i.e., stroke, cradle and hold) were positively natured while the least likely gestures (i.e., slap, hit and pick) were of a more negative nature.

Table 19: Likelihood that people with dementia would use different touch gestures in their interaction with a robot pet according to health care providers. Summed ratings are listed within parentheses.

Likely (med > 3)	Neutral (med = 3)	Unlikely (med < 3)
Stroke (44)	Massage (33)	Finger idly (27)
Cradle (44)	Press (29)	Squeeze (25)
Hold (43)	Push (28)	Shake (25)
Rub (40)	Pinch (25)	Grab (24)
Contact (39)		Poke (24)
Pull (39)		Tap (23)
Hug (38)		Swing (23)
Tickle (38)		Tremble (23)
Kiss (37)		Toss (22)
Nuzzle (37)		Pick (21)
Rock (36)		Hit (19)
Pat (34)		Slap (16)
Lift (33)		
Scratch (33)		

6.5 DISCUSSION

6.5.1 *Usages for robot pets in dementia care*

The health care providers perceived that interaction with Paro could increase the well-being of people with dementia. According to Keyes, mental health is more than the absence of mental illness, it comprises of three components: emotional, psychological and social well-being [70]. Emotional well-being is the presence of positive emotions and the absence of negative emotions [70]. The interviews indicate that interaction with Paro could induce a positive affective state such as calmness and could reduce negative emotions such as sadness. Additionally, the results show that Paro could stimulate communication between patients as well as between patients and health care providers. As positive relations with others contribute to psychological well-being [70], Paro is perceived to improve well-being in this area as well. Both findings are in agreement with the results from literature reviews on the effects of social robots in care for

the elderly [12, 67]. However, from this study it is unclear to what extent Paro could improve the social well-being of people with dementia which refers to how a person functions within society. Improving social well-being seems to be especially important for elderly people who are still aware of their participation within society and/or those who still live independently. Robot pets such as Paro could help elderly people to stay connected by updating them about recent events. An existing example is Nabaztag, a rabbit-shaped robot with an internet connection through which users can have access to news and social media [67].

During the interviews health care providers mentioned that they felt that Paro would support(s) them in their work due to its positive behavioral and psychological effects. In Bemelmans et al. [15] similar expectations were mentioned and Wada et al. [119] reported a reduction in the stress levels of the health care providers due to the reduced need for supervision when the elderly people interacted with Paro. Some of the participants from the layman group in this study also expected this reduction in the need for supervision. In contrast, participants from the expert group indicated that they often intentionally stayed with the patient to observe and evaluate the effect of Paro. Additionally, some of participants in the expert group stated that due to the costliness of the Paro robot and the unpredictable behavior of people with dementia, the possibility of unsupervised interaction with Paro was assessed on an individual basis.

The use of robot pet companions to reduce loneliness was only mentioned by a few participants while literature reviews indicate that this is a goal for which both real animals [16] and robot pets are frequently used [12, 67]. Meanwhile, the use of robots in care for the elderly has sparked a debate about several ethical issues including whether robots will reduce the need for human contact or could even replace humans in the future. Although the participants were not specifically asked about these ethical issues, some of the care providers made comments related to this topic during the interviews. For example, one of the participants stated that she was against the use of Paro to relieve feelings of loneliness as a robot cannot replace human contact. Similarly, Sparrow and Sparrow [105] argued that robots can merely simulate affection, concern and friendship as these robots are unable to engage in genuine social interaction. For this reason they stated that it is a form of deception to let elderly people believe that they are loved and cared for by a robot. Additionally, participants mentioned that adults with normal cognitive health might not feel that they are being taken seriously if they were given the opportunity to interact with a robot pet. Nevertheless, it was also reported that a patient with a chronic physical illness, who had normal cognitive abilities, liked to in-

teract with Paro in spite of knowing that it was not a real animal. In this case there was no form of deception as described by Sparrow and Sparrow [105]. Instead, this case seems in agreement with the view of Sharkey and Sharkey [98] who argued that there are more explanations for why someone would interact with a robot besides being the victim of deceit as people are often willing to interact with an object as if it was a living creature. This argument is supported by the theory that states that people generally treat computers as social actors [84]. With respect to the use of Paro for people with dementia, the participants of this study seem to accept the fact that these patients might see Paro as a living animal and might form a connection with it. From the interviews it became clear that the expert group mostly used Paro in targeted interventions and less often as a means to just keep elderly people company. Despite the ethical considerations, the health care providers in this study perceived Paro as a valuable tool to improve the well-being of people with dementia. Moreover, it was mentioned that, due to the sensitivity of the matter, the use of Paro should always be discussed in advance with the patient's relatives.

6.5.2 *Types of (tactile) interactions between people with dementia and a robot pet*

When asked during the interviews, the health care providers indicated that Paro, with its current capabilities, is a sufficiently effective tool for interventions in dementia care. Nevertheless, during the interviews some shortcomings of Paro came to light and even a few additional functionalities were proposed. This discrepancy could be due to the lack of experience with other social robots and insufficient knowledge of available technological options. We will discuss the insights gained from this indirect information that was provided during the interviews by reviewing how people with dementia (could) interact with Paro according to the health care providers.

The participants indicated that Paro is/ would be able to initiate interaction due to its tactile (softness), visual and auditory cues which can elicit caring responses from people with dementia. Patients are reported/ expected to comfort Paro and show affection by touching the robot seal using positive and affective forms of touch and by talking to it. The predominant use of friendly touch gestures in interactions with a robotic pet was also found in a study with a healthy user group by Yohanan and MacLean [124]. It was indicated by some of the care providers that the use of negative touch gestures by the patients with dementia was often accidental, for example grabbing hold of Paro a bit too roughly. While

Paro is able to distinguish between positive and negative forms of touch, the robot does not distinguish between different forms of positive touch which is the predominant method of interaction. Indeed, Yohanan and MacLean [124] showed that positively natured touch can be used to convey different intentions to a robot pet, that is: protective, comforting, restful, affectionate and playful. Taking these various intentions into account could avoid mistakes such as a robot reacting negatively to a playful, yet slightly rough, touch interaction (e.g. tossing). If Paro was able to better understand the intention behind an interaction the robot could adapt its response to the patient's need (see also Chapter 4). Direct observations of elderly people interacting with an animallike robot will help to further inform to which different types of touch the robot should be able to respond. Moreover, as people with dementia are also reported/ expected to use speech in their interaction with Paro, speech (emotion) recognition might also help to interpret their intentions [33].

Paro's auditory responses were criticized by some of the experts as they can overstimulate the patients and can cause annoyance to the health care providers. In [91] similar critique was expressed by relatives and health care providers. Paro has a built-in volume control function which can be operated using a tiny screwdriver. However, it seems that the participants from the expert group were not aware of this functionality. In addition, adjusting the volume has to be done by the health care providers and requires the use of a special tool which increases the burden of the care providers when working with Paro. The health care providers will have to control for overstimulation, especially when working in groups, which is not desirable. It might be more desirable if Paro's auditory response could be automatically adjusted to the social context, for example by using sensors. Furthermore, not only the volume of the auditory response seems to contribute to the overstimulation and annoyance mentioned by some of the experts. The high pitch and repetitiveness were also part of the critique on the auditory response of Paro.

Tactile responses such as a simulated heartbeat, as suggested by some of the participants in this study, or breathing (e.g. see [21] and Chapter 5) might serve as a more suitable alternative. Paro's current use of tactile responses consists of (accidental) physical contact during its movements. Simulated breathing could be a valuable additional response as it has been found to have a calming effect on the person who is interacting with a robot pet [97]. Furthermore, tactile responses might be less intrusive for others compared to auditory responses.

The way that patients interact with Paro might be influenced by their previous interactions with pets. Some participants argued that a dog or

cat design might therefore be more suitable as these are more familiar animals. However, the developers of Paro deliberately opted for an appearance of a less familiar animal to reduce the chances that the robot could not live up to the user's expectations [100]. For similar reasons others opted for zoomorphic robots with a minimalistic appearance [21, 124]. Currently it is unclear which design would be the most suitable for elderly people with dementia.

6.5.3   *Other target groups that could benefit from interaction with robot pets*

In its current form Paro is perceived to be suitable for children and people with an intellectual disability. Indeed, Paro has previously been used for robot therapy with hospitalized children [99]. Most participants argued that companion robots should have more advanced capabilities to be suitable for a more general audience. As the world's population is aging rapidly and elderly people tend to live longer independently, social robots can be used to assist these people [20, 55, 67]. Service robots such as the Care-O-bot [43] can support elderly people in their everyday activities (e.g. eating and taking medication) and companion robots can be used to enhance their well-being [20, 67]. In order to develop robot pet companions that are suitable for healthy elderly people that still live in their own home, more research into their needs will be necessary. Matching the needs of a target group is important to increase the chance that users will accept the technology [20]. For example, previous research shows that social robots should not negatively affect the self-image of elderly people (e.g. making them feel disabled and dependent) and that the robot's appearance should be serious instead of toy-like [20, 55].

6.5.4   *Considerations regarding the study*

The sample size of the study is small and consists of female participants who share a similar educational background and occupation. Therefore, the results should be interpreted with caution as the findings of this study might not be generalizable. We deliberately recruited health care providers because of their vast experience with the daily care of people with dementia. However, other stakeholders such as patients and their relatives might have different experiences and visions regarding the use of robot pet companions and desirable forms of interaction with these robots. Additionally, it should be noted that there might be a bias in our sample of recruited health care providers. All participants were generally positive

about Paro which might have influenced their decision to agree to take part in our study.

In this study we opted to use Paro as the main example of a robot pet companion for both groups because the expert group already had experience with the robot seal. None of the participants had any previous experience with other social robots. Although the interviews were also set up to explore robot pets in general, it is unclear to which degree the results are applicable to other robots. Moreover, participants that had no experience with Paro (i.e., the layman group) were shown a video of the robot. As a result it might have been more difficult for these participants to judge the capabilities of Paro without seeing and interacting with the robot. In future studies it might therefore be valuable to give health care providers a broader overview of existing robot technology by introducing them to a range of different robots and prototypes similar to the approach taken by [55].

6.6 CONCLUSION

The aim of this study was to inform the development of robot pet companions that can understand and respond to human touch. Such robots might be able to better suit the needs of people with dementia for which touch is an especially important interaction modality. In addition, robots with more advanced interaction capabilities might also be more suitable for other target groups. For this study two groups of health care providers in dementia care were recruited. One group worked with robot seal Paro and the other group did not have any experience with the use of robot pets. Through interviews and a questionnaire we explored in what ways people with dementia could benefit from interaction with an animallike robot with more advanced touch recognition capabilities, how people with dementia (would) interact with Paro (using touch gestures) and which other target groups could benefit from interaction with Paro or an animallike robot with more advanced interaction capabilities.

Interaction with Paro was perceived by the health care providers as an effective intervention that can help to improve the well-being of people with dementia. For example, interaction with Paro could provide distraction, could interrupt problematic behaviors and could stimulate communication. Furthermore, our findings indicated that people with dementia were reported/ expected to mostly use positive forms of touch and speech to interact with Paro. However, Paro's ability to recognize and interpret different types of touch is limited and the social context is not taken into account. Responding to different touch gestures that were reported to be

important during interaction such as stroke, cradle and hug might already result in more effective communication with Paro (see Part II). Moreover, Paro's auditory responses were perceived to be unsuitable for patients with dementia because of the risk of overstimulation. Therefore, more subtle haptic responses such as breathing patterns or a heartbeat might be a valuable addition to Paro's interaction repertoire. Additionally, Paro was perceived to be most suitable for specific target groups such as people with dementia and young children due to its limited interaction abilities. Robot pet companions with more advanced social capabilities such as the ability to have a conversation might better fit the needs of other target groups such as healthy elderly people that still live independently.

Part IV

REFLECTION

7

CONCLUSION

The aim of this dissertation was to work towards socially intelligent robots that can understand and respond to human touch. We have argued in Chapter 1 that a social robot should be able to *sense*, *classify* and *interpret* human touch and *respond* to this in a socially appropriate manner. To this end we have presented work that addressed different parts of the interaction cycle illustrated in Figure 1. In this last chapter we will reflect on our research efforts and we will provide directions for further research.

7.1 RECOGNITION OF SOCIAL TOUCH GESTURES

To work towards more reliable recognition of social touch gestures we have collected and disseminated CoST: Corpus of Social Touch. This dataset contains 7,805 gesture captures of 14 different touch gestures which were performed in 3 variants: gentle, normal and rough. Our results showed that classification of the 14 touch gestures in the CoST dataset independent of the gesture's variant yielded an average accuracy of 60% using SVMs with the RBF kernel (see Chapter 2). Furthermore, gentle gesture variants proved to be harder to classify than the normal and rough variants. Misclassifications were most common between touch gestures with similar characteristics such as grab and squeeze. Additionally, the results showed that there are individual difference in how gestures are performed which makes it difficult to train a generalizable model.

The results from the touch machine learning challenge showed how pre-processing techniques and classification algorithms that are prominent for other modalities can be applied to touch data (see Chapter 3). Many of these methods proved to be reasonably transferable to touch gesture data without much modification. Accuracies up to 61% were reported for the CoST dataset using random forest, which is similar to our results presented in Chapter 2. However, it should be noted that the machine learning challenge comprised a subset of CoST (i.e., gentle and normal variants)

and that the train and test data division was different from the leave-one-subject-out cross-validation results reported in Chapter 2.

Despite the use of different methods, consistent classification confusions between specific gesture pairs were reported by the different challenge participants. Moreover, these misclassifications were also similar to our own findings presented in Chapter 2. The findings presented in Chapters 2 and 3 indicate that the discretization of certain gesture pairs such as rub-stroke and scratch-tickle might be problematic. Additionally, the segmentation and labeling of touch gestures proved to be difficult when participants interacted with a robot pet companion without being given specific instructions on the kinds of behaviors that they should display (see Chapter 4). In contrast to controlled studies in which touch gestures are performed sequentially guided by specific instructions we observed more complex behaviors. Examples include the simultaneous use of multiple touch gestures, touch gestures that were quickly alternated and hybrid forms of prototypical touch gestures. The findings of this study indicated that the categorization of touch behaviors into discrete touch gesture categories based on dictionary definitions is not a suitable approach to capture the complex nature of touch behavior in less controlled settings. This might be problematic when moving towards the automatic understanding of social touch in more naturalistic interactions.

7.2 SOCIAL TOUCH IN THE CONTEXT OF HUMAN-ROBOT INTERACTION

To gain more insight into the factors that are relevant to interpret touch within a social context we studied interactions between humans and a robot pet companion in different affective scenarios. Our results showed that depending on the emotional state of the user, different social messages were communicated to a robot pet such as expressing one's emotional state, seeking emotional support or enjoying the pet's company (see Chapter 4). The expected response from the robot pet to these social messages also varied based on the emotional state. Examples of expected responses were keeping the user company, providing emotional support or picking up on the user's mood. Additionally, the expected response from the robot pet was dependent on the different roles that were envisioned such as a robot that mimics a real pet with its own personality or a robot companion that serves as a therapist/ coach offering emotional support. Furthermore, findings from video observations showed the use of multimodal cues to communicate with the robot pet such as talking to the pet while touching it and making eye contact.

In another study the use of breathing patterns to convey the emotional state of a 1-DOF robot pet was explored (see Chapter 5). The results showed that it is possible to communicate different levels of arousal while the communication of different levels of valence proved to be more difficult indicating that context might be an important factor in the interpretation of valence. Furthermore, the movement rather than materiality seemed to have dominated how participants interpreted emotional expression. This finding indicates that behaviors might be reusable on physically distinct robots with similar abstract abilities such as breathing.

The results of both these studies (i.e., those described in Chapters 4 and 5) confirmed previous findings in the literature on the importance of context in the interpretation of social and affective behavior. Unfortunately, currently commercially available robot companions such as robot seal Paro, which is specifically designed for robot therapy, do not interpret touch within social context (see Chapter 6). As a result, social interaction with Paro is limited and the use of the robot is mostly restricted to specific target groups such as people with dementia and young children.

Similarly to the findings reported in Chapter 4, the results from Chapter 6 showed that people with dementia express affection to Paro by using touch and talking to it. These findings point to the need to study touch within multimodal interaction. Moreover, our findings show that Paro's auditory response can be perceived to be repetitive, irritating, too loud and too high-pitched and can as a result overstimulate patients with dementia. The use of a haptic response using various breathing behaviors as described in Chapter 5 might well avoid such situations. Despite the fact that Paro's interaction abilities could be improved, our results show that the robot is perceived as an effective intervention in dementia care. However, a robot pet that can interpret human touch into context and responds in a socially appropriate manner might suit the needs of patients with dementia even better as touch is an especially important interaction modality for this target group. In addition, robots with more advanced social capabilities might also be more suitable for other target groups.

7.3   CHALLENGES AND OPPORTUNITIES

Comparing the touch gesture recognition findings from different approaches used in the touch machine learning challenge described in Chapter 3 as well as in our own findings presented in Chapter 2 has provided us with the opportunity to pinpoint the difficulties that need to be addressed to increase the reliability of touch recognition. While we have seen commonalities in feature sets used for touch recognition (e.g. calcu-

lating statistics), developing a standard would help ease the feature engineering process. For example, common feature extraction approaches for automatic video classification based on text, audio and video are already documented [19]. Moreover, further research into highly discriminating features using feature selection or dimension reduction can be beneficial for applications that require (onboard) real-time touch classification in which computational power is costly. Furthermore, at this moment it is still unclear how many and which touch gestures should be distinguished and what the minimum requirements are regarding touch recognition performance in order to have a meaningful touch interaction. As the requirements will most likely depend on the application it is important to first assess the use of touch within the intended context. Consequently, social context plays an important role in the classification step as described in Figure 1 in order to determine which types of touch should be recognized. Additionally, new sensor technologies might help to better differentiate between similar touch gestures.

There are still a lot of open questions regarding the understanding of social touch as there is a lot more to determining touch semantics and intent than performing touch gesture recognition. Further research will be necessary to determine whether direct recognition and interpretation of higher level social messages from touch sensor data would be a viable option as this might be a more suitable approach for touch recognition in more natural interaction. In this case, the step of touch gestures recognition before inferring the higher level social meaning of the touch would be omitted. Moreover, multimodal cues could add to the contextual understanding of touch data. Our findings presented in this dissertation indicate that the use of verbal behavior that coincides with touch interaction is an interesting direction for future studies into the automatic understanding of social touch.

Currently, most studies, including the work presented in this dissertation, focuses on parts of the interaction. However, future work should eventually tie together the full interaction cycle. For example, as the results of Chapter 4 are based on acted scenarios it is important to verify whether similar behaviors occur in a naturalistic setting in which people would interact with a fully functioning robot pet in their own home. A first step could be to induce emotions in participants and observe their interactions with a responding robot pet in a lab setting.

With regard to robot responses, future work should explore the expression of affect beyond the four emotional states that were studied in Chapter 5. Additionally, future studies could focus on examining longer interactions in which the robot would adapt its behavior to the users actions.

In this scenario it might be interesting to study the effects of the robot's breathing behavior on the user's emotional state. For example, could a happy robot pet cheer up a sad person?

In the health care domain, additional research will be necessary to determine both the short and long term effects of interaction with robot pets on different aspects of well-being. For example, the effects of robot pets can be assessed through questionnaires such as the Mental Health Continuum-Short Form (MHC-SF) [76]. Furthermore, the development of robot pets with more advanced social capabilities such as touch and speech recognition might result in more intelligent interactions which could help to better adapt to the needs of people with dementia and could make interactions more interesting for a broader audience. Moreover, to further improve interactions with robot pet companions future studies should explore which forms of robot response modalities (e.g. sounds, movements or breathing patterns) best suit a target group. Additionally, the advantage that a familiar robot appearance can have by relating to previous experiences versus the risk of unmet expectations could be the focus of future research.

LIST OF PUBLICATIONS

THIS DISSERTATION IS BASED ON THE FOLLOWING PUBLICATIONS

Bucci, P, Cang, X. L., Valair, A, Marino, D, Tseng, L, Jung, M. M., Rantala, J, Schneider, O. S., and MacLean, K. E., "Sketching cuddlebits: Coupled prototyping of body and behaviour for an affective robot pet," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI), (Denver, CO)*, 2017, pp. 3681–3692, (partially included).

Jung, M. M., Cang, X. L., Poel, M., and MacLean, K. E., "Touch challenge '15: Recognizing social touch gestures," in *Proceedings of the International Conference on Multimodal Interaction (ICMI), (Seattle, WA)*, 2015, pp. 387–390.

Jung, M. M., Poel, M., Poppe, R., and Heylen, D. K. J., "Automatic recognition of touch gestures in the corpus of social touch," *Journal on Multimodal User Interfaces*, vol. 11, no. 1, pp. 81–96, 2016.

Jung, M. M., Poel, M., Reidsma, D., and Heylen, D. K. J., "A first step towards the automatic understanding of social touch for naturalistic human-robot interaction," *Frontiers in ICT*, vol. 4, no. 3, 2017.

Jung, M. M., Van Der Leij, L., and Kelders, S. M., "Touch interaction with animallike robot companions: The use of Paro in dementia care," Manuscript submitted for publication.

THIS DISSERTATION WAS ALSO INFLUENCED BY THE FOLLOWING PUBLICATIONS THAT RESULTED FROM WORK THAT WAS CARRIED OUT BY THE AUTHOR DURING THE PHD PROGRAM

Huisman, G., Bruijnes, M., Kolkmeier, J., Jung, M. M., Darriba-Frederiks, A., and Rybarczyk, Y., "Touching virtual agents: Embodiment and mind," in *Innovative and Creative Developments in Multimodal Interaction Systems*, Springer, 2014, pp. 114–138.

Jung, M. M., "Towards social touch intelligence: Developing a robust system for automatic touch recognition," in *Proceedings of the International Conference on Multimodal Interaction (ICMI), (Istanbul, Turkey)*, 2014, pp. 344–348.

Jung, M. M., Boensma, R. W., Huisman, G., and Van Dijk, B., "Touched by the storyteller: The influence of remote touch in the context of storytelling," in *Proceedings of the Humaine Association Conference on Affective Computing and Intelligent Interaction (ACII), (Geneva, Switzerland)*, 2013, pp. 792–797.

Jung, M. M., Poel, M., Poppe, R., and Heylen, D. K. J., "Corpus of Social Touch (CoST)," *University of Twente*, 2016, `http://dx.doi.org/10.4121/ uuid:5ef62345-3b3e-479c-8e1d-c922748c9b29`.

Jung, M. M., Poppe, R., Poel, M., and Heylen, D. K. J., "Touching the void – introducing CoST: Corpus of Social Touch," in *Proceedings of the International Conference on Multimodal Interaction (ICMI), (Istanbul, Turkey)*, 2014, pp. 120–127.

Van Wingerden, S., Uebbing, T. J., Jung, M. M., and Poel, M., "A neural network based approach to social touch classification," in *Proceedings of Workshop on Emotion Representation and Modelling in Human-Computer-Interaction-Systems (ERM4HCI), (Istanbul, Turkey)*, 2014, pp. 7–12.

Willemse, C. J.A. M., Huisman, G., Jung, M. M., Van Erp, J. B. F., and Heylen, D. K. J., "Observing touch from video: The influence of social cues on pleasantness perceptions," in *International Conference on Human Haptic Sensing and Touch Enabled Computer Applications (EuroHaptics), (London, UK)*, 2016, pp. 196–205.

BIBLIOGRAPHY

[1] Ackerley, R., Wasling, H. B., Liljencrantz, J., Olausson, H., Johnson, R. D., and Wessberg, J., "Human c-tactile afferents are tuned to the temperature of a skin-stroking caress," *Journal of Neuroscience*, vol. 34, no. 8, pp. 2879–2883, 2014.

[2] Altun, K. and MacLean, K. E., "Recognizing affect in human touch of a robot," *Pattern Recognition Letters*, vol. 66, pp. 31–40, 2015.

[3] App, B., McIntosh, D. N., Reed, C. L., and Hertenstein, M. J., "Nonverbal channel use in communication of emotion: How may depend on why.," *Emotion*, vol. 11, no. 3, p. 603, 2011.

[4] Argall, B. D. and Billard, A. G., "A survey of tactile human–robot interactions," *Robotics and Autonomous Systems*, vol. 58, no. 10, pp. 1159–1176, 2010.

[5] Baarda, D. B., De Goede, M. P. M., and Teunissen, J, "Basisboek kwalitatief onderzoek. handleiding voor het opzetten en uitvoeren van kwalitatief onderzoek," *Groningen: Wolters Noordhoff*, 2005.

[6] Bailenson, J. N., Yee, N., Brave, S., Merget, D., and Koslow, D., "Virtual interpersonal touch: Expressing and recognizing emotions through haptic devices," *Human–Computer Interaction*, vol. 22, no. 3, pp. 325–353, 2007.

[7] Balli Altuglu, T. and Altun, K., "Recognizing touch gestures for social human-robot interaction," in *Proceedings of the International Conference on Multimodal Interaction (ICMI), (Seattle, WA)*, 2015, pp. 407–413.

[8] Banks, M. R., Willoughby, L. M., and Banks, W. A., "Animal-assisted therapy and loneliness in nursing homes: Use of robotic versus living dogs," *Journal of the American Medical Directors Association*, vol. 9, no. 3, pp. 173–177, 2008.

[9] Barnett, K. E., *The development of a theoretical construct of the concepts of touch as they relate to nursing. final report.* `http://eric.ed.gov/?id=ED044521`, Accessed: 2017-03-28, 1970.

[10] Barrett, L. F., Mesquita, B., and Gendron, M., "Context in emotion perception," *Current Directions in Psychological Science*, vol. 20, no. 5, pp. 286–290, 2011.

[11] Batliner, A., Biersack, S, and Steidl, S, "The prosody of pet robot directed speech: Evidence from children," in *Proceedings of Speech Prosody (Dresden, Germany)*, 2006, pp. 1–4.

[12] Bemelmans, R., Gelderblom, G. J., Jonker, P., and De Witte, L., "Socially assistive robots in elderly care: A systematic review into effects and effectiveness," *Journal of the American Medical Directors Association*, vol. 13, no. 2, pp. 114–120, 2012.

[13] Bemelmans, R., Gelderblom, G. J., Jonker, P., and Witte, L. de, "Effectiveness of robot Paro in intramural psychogeriatric care: A multicenter quasi-experimental study," *Journal of the American Medical Directors Association*, vol. 16, no. 11, pp. 946–950, 2015.

[14] Bemelmans, R., Gelderblom, G. J., Jonker, P., and Witte, L. de, "How to use robot interventions in intramural psychogeriatric care; a feasibility study," *Applied Nursing Research*, vol. 30, pp. 154–157, 2016.

[15] Bemelmans, R., Gelderblom, G. J., Spierts, N., Jonker, P., and Witte, L. de, "Development of robot interventions for intramural psychogeriatric care.," *GeroPsych: The Journal of Gerontopsychology and Geriatric Psychiatry*, vol. 26, no. 2, pp. 113–120, 2013.

[16] Bernabei, V, De Ronchi, D, La Ferla, T, Moretti, F, Tonelli, L, Ferrari, B, Forlani, M, and Atti, A. R., "Animal-assisted interventions for elderly patients affected by dementia or psychiatric disorders: A review," *Journal of psychiatric research*, vol. 47, no. 6, pp. 762–773, 2013.

[17] Bickmore, T. W., Fernando, R., Ring, L., and Schulman, D., "Empathic touch by relational agents," *Transactions on Affective Computing*, vol. 1, no. 1, pp. 60–71, 2010.

[18] Boiten, F. A., Frijda, N. H., and Wientjes, C. J. E., "Emotions and respiratory patterns: Review and critical analysis," *International Journal of Psychophysiology*, vol. 17, no. 2, pp. 103–128, 1994.

[19] Brezeale, D. and Cook, D. J., "Automatic video classification: A survey of the literature," *Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, vol. 38, no. 3, pp. 416–430, 2008.

[20] Broadbent, E., Stafford, R., and MacDonald, B., "Acceptance of healthcare robots for the older population: Review and future directions," *International Journal of Social Robotics*, vol. 1, no. 4, pp. 319–330, 2009.

[21] Bucci, P, Cang, X. L., Valair, A, Marino, D, Tseng, L, Jung, M. M., Rantala, J, Schneider, O. S., and MacLean, K. E., "Sketching cuddlebits: Coupled prototyping of body and behaviour for an affective robot pet," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI), (Denver, CO)*, 2017, pp. 3681–3692.

[22] Cang, L., Bucci, P., and MacLean, K. E., "Cuddlebits: Friendly, low-cost furballs that respond to touch," in *Proceedings of the International Conference on Multimodal Interaction (ICMI), (Seattle, WA)*, 2015, pp. 365–366.

[23] Cang, X. L., Bucci, P., Strang, A., Allen, J., MacLean, K. E., and Liu, H., "Different strokes and different folks: Economical dynamic surface sensing and affect-related touch recognition," in *Proceedings of the International Conference on Multimodal Interaction (Seattle, WA)*, 2015, pp. 147–154.

[24] Chang, C.-C. and Lin, C.-J., "LIBSVM: A library for support vector machines," *Transactions on Intelligent Systems and Technology*, vol. 2, 27:1–27:27, 3 2011.

[25] Chang, J., MacLean, K. E., and Yohanan, S., "Gesture recognition in the haptic creature," in *Proceedings of the International Conference EuroHaptics, (Amsterdam, The Netherlands)*, 2010, pp. 385–391.

[26] Cooney, M. D., Becker-Asano, C., Kanda, T., Alissandrakis, A., and Ishiguro, H., "Full-body gesture recognition using inertial sensors for playful interaction with small humanoid robot," in *Proceedings of International Conference on Intelligent Robots and Systems (IROS), (Taipei, Taiwan)*, 2010, pp. 2276–2282.

[27] Cooney, M. D., Nishio, S., and Ishiguro, H., "Recognizing affection for a touch-based interaction with a humanoid robot," in *Proceedings of the International Conference on Intelligent Robots and Systems (IROS), (Vilamoura, Portugal)*, 2012, pp. 1420–1427.

[28] Cooney, M. D., Nishio, S., and Ishiguro, H., "Importance of touch for conveying affection in a multimodal interaction with a small humanoid robot," *International Journal of Humanoid Robotics*, vol. 12, no. 01, p. 1 550 002, 2015.

[29] Dahiya, R. S., Metta, G., Valle, M., and Sandini, G., "Tactile sensing – from humans to humanoids," *Transactions on Robotics*, vol. 26, no. 1, pp. 1–20, 2010.

[30]   Dawson, J. Q., Schneider, O. S., Ferstay, J., Toker, D., Link, J., Haddad, S., and MacLean, K., "It's alive!: Exploring the design space of a gesturing phone," in *Proceedings of Graphics Interface (Regina, Sascatchewan, Canada)*, 2013, pp. 205–212.

[31]   De Melo, C. M., Kenny, P., and Gratch, J., "Real-time expression of affect through respiration," *Computer Animation and Virtual Worlds*, vol. 21, no. 3-4, pp. 225–234, 2010.

[32]   Eachus, P., "Pets, people and robots: The role of companion animals and robopets in the promotion of health and well-being," *International Journal of Health Promotion and Education*, vol. 39, no. 1, pp. 7–13, 2001.

[33]   El Ayadi, M., Kamel, M. S., and Karray, F., "Survey on speech emotion recognition: Features, classification schemes, and databases," *Pattern Recognition*, vol. 44, no. 3, pp. 572–587, 2011.

[34]   Field, T., "Touch for socioemotional and physical well-being: A review," *Developmental Review*, vol. 30, no. 4, pp. 367–383, 2010.

[35]   Filan, S. L. and Llewellyn-Jones, R. H., "Animal-assisted therapy for dementia: A review of the literature," *International Psychogeriatrics*, vol. 18, no. 4, pp. 597–612, 2006.

[36]   Fisher, J. D., Rytting, M., and Heslin, R., "Hands touching hands: Affective and evaluative effects of an interpersonal touch," *Sociometry*, pp. 416–421, 1976.

[37]   Flagg, A. and MacLean, K. E., "Affective touch gesture recognition for a furry zoomorphic machine," in *Proceedings of the International Conference on Tangible, Embedded and Embodied Interaction (TEI), (Barcelona, Spain)*, 2013, pp. 25–32.

[38]   Flagg, A., Tam, D., MacLean, K. E., and Flagg, R., "Conductive fur sensing for a gesture-aware furry robot," in *Proceedings of the Haptics Symposium (HAPTICS), (Vancouver, Canada)*, 2012, pp. 99–104.

[39]   Fujita, M., "Aibo: Toward the era of digital creatures," *The International Journal of Robotics Research*, vol. 20, no. 10, pp. 781–794, 2001.

[40]   Gallace, A. and Spence, C., "The science of interpersonal touch: An overview," *Neuroscience & Biobehavioral Reviews*, vol. 34, no. 2, pp. 246–259, 2010.

[41]   Gaus, Y. F. A., Olugbade, T., Jan, A., Qin, R., Liu, J., Zhang, F., *et al.*, "Social touch gesture recognition using random forest and boosting on distinct feature sets," in *Proceedings of the International Conference on Multimodal Interaction (ICMI), (Seattle, WA)*, 2015, pp. 399–406.

[42] Gazzola, V., Spezio, M. L., Etzel, J. A., Castelli, F., Adolphs, R., and Keysers, C., "Primary somatosensory cortex discriminates affective significance in social touch," *Proceedings of the National Academy of Sciences*, vol. 109, no. 25, E1657–E1666, 2012.

[43] Graf, B., Reiser, U., Hägele, M., Mauz, K., and Klein, P., "Robotic home assistant care-o-bot® 3-product vision and innovation platform," in *Proceedings of the Workshop on Advanced Robotics and its Social Impacts (ARSO), (Tokyo, Japan)*, 2009, pp. 139–144.

[44] Guéguen, N., "Touch, awareness of touch, and compliance with a request," *Perceptual and motor skills*, vol. 95, no. 2, pp. 355–360, 2002.

[45] Gunes, H. and Schuller, B., "Categorical and dimensional affect analysis in continuous input: Current trends and future directions," *Image and Vision Computing*, vol. 31, no. 2, pp. 120–136, 2013.

[46] Haans, A. and IJsselsteijn, W., "Mediated social touch: A review of current research and future directions," *Virtual Reality*, vol. 9, no. 2-3, pp. 149–159, 2006.

[47] Henricson, M., Ersson, A., Määttä, S., Segesten, K., and Berglund, A.-L., "The outcome of tactile touch on stress parameters in intensive care: A randomized controlled trial," *Complementary therapies in clinical practice*, vol. 14, no. 4, pp. 244–254, 2008.

[48] Hertenstein, M. J., Holmes, R., McCullough, M., and Keltner, D., "The communication of emotion via touch.," *Emotion*, vol. 9, no. 4, pp. 566–573, 2009.

[49] Hertenstein, M. J., Keltner, D., App, B., Bulleit, B. A., and Jaskolka, A. R., "Touch communicates distinct emotions.," *Emotion*, vol. 6, no. 3, pp. 528–533, 2006.

[50] Hertenstein, M. J., Verkamp, J. M., Kerestes, A. M., and Holmes, R. M., "The communicative functions of touch in humans, nonhuman primates, and rats: A review and synthesis of the empirical research," *Genetic, Social, and General Psychology Monographs*, vol. 132, no. 1, pp. 5–94, 2006.

[51] Heslin, R., Nguyen, T. D., and Nguyen, M. L., "Meaning of touch: The case of touch from a stranger or same sex person," *Journal of Nonverbal Behavior*, vol. 7, no. 3, pp. 147–157, 1983.

[52] Hsu, C.-W., Chang, C.-C., Lin, C.-J., *et al.*, *A practical guide to support vector classification*, (Date last accessed June 10, 2015), 2003. [Online]. Available: http://www.csie.ntu.edu.tw/~cjlin/papers/guide/guide.pdf.

[53] Hughes, D., Farrow, N., Profita, H., and Correll, N., "Detecting and identifying tactile gestures using deep autoencoders, geometric moments and gesture level features," in *Proceedings of the International Conference on Multimodal Interaction (ICMI), (Seattle, WA)*, 2015, pp. 415–422.

[54] Huisman, G., Bruijnes, M., Kolkmeier, J., Jung, M. M., Darriba-Frederiks, A., and Rybarczyk, Y., "Touching virtual agents: Embodiment and mind," in *Innovative and Creative Developments in Multimodal Interaction Systems*, Springer, 2014, pp. 114–138.

[55] Hutson, S., Lim, S. L., Bentley, P. J., Bianchi-Berthouze, N., and Bowling, A., "Investigating the suitability of social robots for the wellbeing of the elderly," in *Proceedings of the International Conference on Affective Computing and Intelligent Interaction (ACII), (Memphis, TN)*, 2011, pp. 578–587.

[56] Jeong, S., Logan, D., Goodwin, M., Graca, S., O'Connell, B., Goodenough, H., *et al.*, "A social robot to mitigate stress, anxiety, and pain in hospital pediatric care," in *Proceedings of the International Conference on Human-Robot Interaction (HRI) Extended Abstracts, (Portland, OR)*, 2015, pp. 103–104.

[57] Ji, Z., Amirabdollahian, F., Polani, D., and Dautenhahn, K., "Histogram based classification of tactile patterns on periodically distributed skin sensors for a humanoid robot," in *Proceedings of the International Symposium on Robot and Human Interactive Communication (ROMAN), (Atlanta, GA)*, 2011, pp. 433–440.

[58] Johnson, K. O., Yoshioka, T., and Vega-Bermudez, F., "Tactile functions of mechanoreceptive afferents innervating the hand," *Journal of Clinical Neurophysiology*, vol. 17, no. 6, pp. 539–558, 2000.

[59] Jones, S. E. and Yarbrough, A. E., "A naturalistic study of the meanings of touch," *Communications Monographs*, vol. 52, no. 1, pp. 19–56, 1985.

[60] Jung, M. M., "Towards social touch intelligence: Developing a robust system for automatic touch recognition," in *Proceedings of the International Conference on Multimodal Interaction (ICMI), (Istanbul, Turkey)*, 2014, pp. 344–348.

[61] Jung, M. M., Cang, X. L., Poel, M., and MacLean, K. E., "Touch challenge '15: Recognizing social touch gestures," in *Proceedings of the International Conference on Multimodal Interaction (ICMI), (Seattle, WA)*, 2015, pp. 387–390.

[62] Jung, M. M., Poel, M., Poppe, R., and Heylen, D. K. J., "Automatic recognition of touch gestures in the corpus of social touch," *Journal on Multimodal User Interfaces*, vol. 11, no. 1, pp. 81–96, 2016.

[63] Jung, M. M., Poel, M., Poppe, R., and Heylen, D. K. J., "Corpus of Social Touch (CoST)," *University of Twente*, 2016, `http://dx.doi.org/10.4121/uuid:5ef62345-3b3e-479c-8e1d-c922748c9b29`.

[64] Jung, M. M., Poel, M., Reidsma, D., and Heylen, D. K. J., "A first step towards the automatic understanding of social touch for naturalistic human-robot interaction," *Frontiers in ICT*, vol. 4, no. 3, 2017.

[65] Jung, M. M., Poppe, R., Poel, M., and Heylen, D. K. J., "Touching the void – introducing CoST: Corpus of Social Touch," in *Proceedings of the International Conference on Multimodal Interaction (ICMI), (Istanbul, Turkey)*, 2014, pp. 120–127.

[66] Kächele, M., Schels, M., Meudt, S., Palm, G., and Schwenker, F., "Revisiting the emotiw challenge: How wild is it really?" *Journal on Multimodal User Interfaces*, vol. 10, no. 2, pp. 151–162, 2016.

[67] Kachouie, R., Sedighadeli, S., Khosla, R., and Chu, M.-T., "Socially assistive robots in elderly care: A mixed-method systematic literature review," *International Journal of Human-Computer Interaction*, vol. 30, no. 5, pp. 369–393, 2014.

[68] Kandel, E. R., Schwartz, J. H., Jessell, T. M., Siegelbaum, S. A., and Hudspeth, A. J., *Principles of neural science*. McGraw-Hill, New York, 2013.

[69] Kendon, A., "Gesticulation and speech: Two aspects of the process of utterance," *The relationship of verbal and nonverbal communication*, vol. 25, no. 1980, pp. 207–227, 1980.

[70] Keyes, C. L., "The mental health continuum: From languishing to flourishing in life," *Journal of health and social behavior*, vol. 43, no. 2, pp. 207–222, 2002.

[71] Kim, Y.-M., Koo, S.-Y., Lim, J. G., and Kwon, D.-S., "A robust online touch pattern recognition for dynamic human-robot interaction," *Transactions on Consumer Electronics*, vol. 56, no. 3, pp. 1979–1987, 2010.

[72] Kita, S., Van Gijn, I., and Van Der Hulst, H., "Movement phases in signs and co-speech gestures, and their transcription by human coders," in *International Gesture Workshop*, Springer, 1997, pp. 23–35.

[73] Knight, H., Toscano, R., Stiehl, W. D., Chang, A., Wang, Y., and Breazeal, C., "Real-time social touch gesture recognition for sensate robots," in *Proceedings of the International Conference on Intelligent Robots and Systems (IROS), (St. Louis, MO)*, 2009, pp. 3715–3720.

[74] Kotranza, A. and Lok, B., "Virtual human + tangible interface = mixed reality human an initial exploration with a virtual breast exam patient," in *Proceedings of the Virtual Reality Conference, (Reno, NV)*, 2008, pp. 99–106.

[75] Kotranza, A., Lok, B., Pugh, C. M., and Lind, D. S., "Virtual humans that touch back: Enhancing nonverbal communication with virtual humans through bidirectional touch," in *Proceedings of the Virtual Reality Conference, (Lafayette, LA)*, 2009, pp. 175–178.

[76] Lamers, S., Westerhof, G. J., Bohlmeijer, E. T., Klooster, P. M. ten, and Keyes, C. L., "Evaluating the psychometric properties of the mental health continuum-short form (MHC-SF)," *Journal of clinical psychology*, vol. 67, no. 1, pp. 99–110, 2011.

[77] Löken, L. S., Evert, M., and Wessberg, J., "Pleasantness of touch in human glabrous and hairy skin: Order effects on affective ratings," *Brain research*, vol. 1417, pp. 9–15, 2011.

[78] Löken, L. S., Wessberg, J., McGlone, F., and Olausson, H., "Coding of pleasant touch by unmyelinated afferents in humans," *Nature neuroscience*, vol. 12, no. 5, pp. 547–548, 2009.

[79] McGlone, F, Olausson, H, Boyle, J., Jones-Gotman, M, Dancer, C, Guest, S, and Essick, G, "Touching and feeling: Differences in pleasant touch processing between glabrous and hairy skin in humans," *European Journal of Neuroscience*, vol. 35, no. 11, pp. 1782–1788, 2012.

[80] Morrison, I., "Keep calm and cuddle on: Social touch as a stress buffer," *Adaptive Human Behavior and Physiology*, vol. 2, no. 4, pp. 344–362, 2016.

[81] Morrison, I., Björnsdotter, M., and Olausson, H., "Vicarious responses to social touch in posterior insular cortex are tuned to pleasant caressing speeds," *Journal of Neuroscience*, vol. 31, no. 26, pp. 9554–9562, 2011.

[82] Morrison, I., Löken, L. S., and Olausson, H., "The skin as a social organ," *Experimental brain research*, vol. 204, no. 3, pp. 305–314, 2010.

[83] Nakajima, K., Itoh, Y., Hayashi, Y., Ikeda, K., Fujita, K., and Onoye, T., "Emoballoon: A balloon-shaped interface recognizing social touch interactions," in *Proceedings of Advances in Computer Entertainment (ACE), (Boekelo, The Netherlands)*, 2013, pp. 182–197.

[84] Nass, C., Steuer, J., and Tauber, E. R., "Computers are social actors," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI), (Boston, MA)*, 1994, pp. 72–78.

[85] Naya, F., Yamato, J., and Shinozawa, K., "Recognizing human touching behaviors using a haptic interface for a pet-robot," in *Proceedings of the International Conference on Systems, Man, and Cybernetics (SMC), (Tokyo, Japan)*, vol. 2, 1999, pp. 1030–1034.

[86] Nicolaou, M. A., Gunes, H., and Pantic, M., "Continuous prediction of spontaneous affect from multiple cues and modalities in valence-arousal space," *Transactions on Affective Computing*, vol. 2, no. 2, pp. 92–105, 2011.

[87] Noda, T., Ishiguro, H., Miyashita, T., and Hagita, N., "Map acquisition and classification of haptic interaction using cross correlation between distributed tactile sensors on the whole body surface," in *Proceedings of the International Conference on Intelligent Robots and Systems (IROS), (San Diego, CA)*, 2007, pp. 1099–1105.

[88] Picard, R. W., "Affective computing: Challenges," *International Journal of Human-Computer Studies*, vol. 59, no. 1, pp. 55–64, 2003.

[89] Powell, J. A., "Communication interventions in dementia," *Reviews in Clinical Gerontology*, vol. 10, no. 02, pp. 161–168, 2000.

[90] Rainville, P., Bechara, A., Naqvi, N., and Damasio, A. R., "Basic emotions are associated with distinct patterns of cardiorespiratory activity," *International journal of psychophysiology*, vol. 61, no. 1, pp. 5–18, 2006.

[91] Robinson, H., MacDonald, B. A., Kerse, N., and Broadbent, E., "Suitability of healthcare robots for a dementia unit and suggested improvements," *Journal of the American Medical Directors Association*, vol. 14, no. 1, pp. 34–40, 2013.

[92] Robinson, H., MacDonald, B., and Broadbent, E., "Physiological effects of a companion robot on blood pressure of older people in residential care facility: A pilot study," *Australasian journal on ageing*, vol. 34, no. 1, pp. 27–32, 2015.

[93] Russell, J. A., Weiss, A., and Mendelsohn, G. A., "Affect grid: A single-item scale of pleasure and arousal." *Journal of personality and social psychology*, vol. 57, no. 3, pp. 493–502, 1989.

[94] Russell, S. and Norvig, P., *Artificial intelligence: A modern approach*, Third Edition. Prentice Hall, 2010.

[95]   Saldien, J., Goris, K., Vanderborght, B., Vanderfaeillie, J., and Lefeber, D., "Expressing emotions with the social robot probo," *International Journal of Social Robotics*, vol. 2, no. 4, pp. 377–389, 2010.

[96]   Schneider, O. S. and MacLean, K. E., "Studying design process and example use with macaron, a web-based vibrotactile effect editor," in *Haptics Symposium (HAPTICS), (Philadelphia, PA)*, 2016, pp. 52–58.

[97]   Sefidgar, Y. S., MacLean, K. E., Yohanan, S., Van Der Loos, H. M., Croft, E. A., and Garland, E. J., "Design and evaluation of a touch-centered calming interaction with a social robot," *Transactions on Affective Computing*, vol. 7, no. 2, pp. 108–121, 2016.

[98]   Sharkey, A. and Sharkey, N., "Granny and the robots: Ethical issues in robot care for the elderly," *Ethics and Information Technology*, vol. 14, no. 1, pp. 27–40, 2012.

[99]   Shibata, T., Mitsui, T., Wada, K., Touda, A., Kumasaka, T., Tagami, K., and Tanie, K., "Mental commit robot and its application to therapy of children," in *Proceedings of International Conference on Advanced Intelligent Mechatronics, (Como, Italy)*, vol. 2, 2001, pp. 1053–1058.

[100] Shibata, T. and Wada, K., "Robot therapy: A new approach for mental healthcare of the elderly–a mini-review," *Gerontology*, vol. 57, no. 4, pp. 378–386, 2010.

[101] Silvera-Tawil, D., Rye, D., and Velonaki, M., "Touch modality interpretation for an EIT-based sensitive skin," in *Proceedings of the International Conference on Robotics and Automation (ICRA), (Shanghai, China)*, 2011, pp. 3770–3776.

[102] Silvera-Tawil, D., Rye, D., and Velonaki, M., "Interpretation of the modality of touch on an artificial arm covered with an EIT-based sensitive skin," *The International Journal of Robotics Research*, vol. 31, no. 13, pp. 1627–1641, 2012.

[103] Silvera-Tawil, D., Rye, D., and Velonaki, M., "Interpretation of social touch on an artificial arm covered with an EIT-based sensitive skin," *International Journal of Social Robotics*, vol. 6, no. 4, pp. 489–505, 2014.

[104] Silvera-Tawil, D., Rye, D., and Velonaki, M., "Artificial skin and tactile sensing for socially interactive robots: A review," *Robotics and Autonomous Systems*, vol. 63, pp. 230–243, 2015.

[105] Sparrow, R. and Sparrow, L., "In the hands of machines? the future of aged care," *Minds and Machines*, vol. 16, no. 2, pp. 141–161, 2006.

[106]   Stiehl, W. D., Lieberman, J., Breazeal, C., Basel, L., Lalla, L., and Wolf, M., "Design of a therapeutic robotic companion for relational, affective touch," in *Proceedings of International Workshop on Robot and Human Interactive Communication (ROMAN), (Nashville, TN)*, 2005, pp. 408–415.

[107]   Suvilehto, J. T., Glerean, E., Dunbar, R. I., Hari, R., and Nummenmaa, L., "Topography of social touching depends on emotional bonds between humans," *Proceedings of the National Academy of Sciences*, p. 201 519 231, 2015.

[108]   Ta, V.-C., Johal, W., Portaz, M., Castelli, E., and Vaufreydaz, D., "The Grenoble system for the social touch challenge at ICMI 2015," in *Proceedings of the International Conference on Multimodal Interaction (ICMI), (Seattle, WA)*, 2015, pp. 391–398.

[109]   Tao, J. and Tan, T., "Affective computing: A review," in *Proceedings of the International Conference on Affective Computing and Intelligent Interaction (ACII), (Beijing, China)*, 2005, pp. 981–995.

[110]   Thayer, S., "History and strategies of research on social touch," *Journal of Nonverbal Behavior*, vol. 10, no. 1, pp. 12–28, 1986.

[111]   Van Der Leij, L, "Robotdieren in de ouderenzorg: Een verkennend onderzoek naar de toegevoegde waarde van het verfijnen van tactiele interactie," Master's thesis, University of Twente, Enschede, The Netherlands, 2016.

[112]   Van Erp, J. B. F. and Toet, A., "Social touch in human–computer interaction," *Frontiers in Digital Humanities*, vol. 2, no. 2, pp. 1–14, 2015.

[113]   Van Wingerden, S., Uebbing, T. J., Jung, M. M., and Poel, M., "A neural network based approach to social touch classification," in *Proceedings of Workshop on Emotion Representation and Modelling in Human-Computer-Interaction-Systems (ERM4HCI), (Istanbul, Turkey)*, 2014, pp. 7–12.

[114]   Veevers, J. E., "The social meaning of pets: Alternative roles for companion animals," *Marriage & Family Review*, vol. 8, no. 3-4, pp. 11–30, 1985.

[115]   Vinciarelli, A., Pantic, M., and Bourlard, H., "Social signal processing: Survey of an emerging domain," *Image and Vision Computing*, vol. 27, no. 12, pp. 1743–1759, 2009.

[116] Vinciarelli, A., Pantic, M., Bourlard, H., and Pentland, A., "Social signals, their function, and automatic analysis: A survey," in *Proceedings of the International Conference on Multimodal Interfaces (ICMI), (Chania, Crete, Greece)*, 2008, pp. 61–68.

[117] Wada, K. and Shibata, T., "Living with seal robots – its sociopsychological and physiological influences on the elderly at a care house," *Transactions on Robotics*, vol. 23, no. 5, pp. 972–980, 2007.

[118] Wada, K., Shibata, T., Saito, T., Sakamoto, K., and Tanie, K., "Psychological and social effects of one year robot assisted activity on elderly people at a health service facility for the aged," in *Proceedings of the International Conference on Robotics and Automation (ICRA), (Barcelona, Spain)*, 2005, pp. 2785–2790.

[119] Wada, K., Shibata, T., Saito, T., and Tanie, K., "Effects of robot-assisted activity for elderly people and nurses at a day service center," *Proceedings of the IEEE*, vol. 92, no. 11, pp. 1780–1788, 2004.

[120] Willemse, C. J.A. M., Huisman, G., Jung, M. M., Van Erp, J. B. F., and Heylen, D. K. J., "Observing touch from video: The influence of social cues on pleasantness perceptions," in *International Conference on Human Haptic Sensing and Touch Enabled Computer Applications (EuroHaptics), (London, UK)*, 2016, pp. 196–205.

[121] World Health Organization, *The icd-10 classification of mental and behavioural disorders: Diagnostic criteria for research*, Geneva: World Health Organization, 1993.

[122] World Health Organization, *Dementia: A public health priority*, Geneva: World Health Organization, 2012.

[123] Yohanan, S. and MacLean, K. E., "Design and assessment of the haptic creature's affect display," in *Proceedings of the International conference on Human-Robot Interaction (HRI), (Lausanne, Switzerland)*, 2011, pp. 473–480.

[124] Yohanan, S. and MacLean, K. E., "The role of affective touch in human-robot interaction: Human intent and expectations in touching the haptic creature," *Social Robotics*, vol. 4, no. 2, pp. 163–180, 2012.

[125] Zasloff, R. L., "Measuring attachment to companion animals: A dog is not a cat is not a bird," *Applied Animal Behaviour Science*, vol. 47, no. 1, pp. 43–48, 1996.

[126] Zeng, Z., Pantic, M., Roisman, G., and Huang, T. S., "A survey of affect recognition methods: Audio, visual, and spontaneous expressions," *Transactions on Pattern Analysis and Machine Intelligence*, vol. 31, no. 1, pp. 39–58, 2009.

SIKS DISSERTATION SERIES

2011-01 Botond Cseke (RUN) *Variational Algorithms for Bayesian Inference in Latent Gaussian Models*

2011-02 Nick Tinnemeier (UU) *Organizing Agent Organizations. Syntax and Operational Semantics of an Organization-Oriented Programming Language*

2011-03 Jan Martijn van der Werf (TUE) *Compositional Design and Verification of Component-Based Information Systems*

2011-04 Hado van Hasselt (UU) *Insights in Reinforcement Learning; Formal analysis and empirical evaluation of temporal-difference*

2011-05 Base van der Raadt (VU) *Enterprise Architecture Coming of Age - Increasing the Performance of an Emerging Discipline.*

2011-06 Yiwen Wang (TUE) *Semantically-Enhanced Recommendations in Cultural Heritage*

2011-07 Yujia Cao (UT) *Multimodal Information Presentation for High Load Human Computer Interaction*

2011-08 Nieske Vergunst (UU) *BDI-based Generation of Robust Task-Oriented Dialogues*

2011-09 Tim de Jong (OU) *Contextualised Mobile Media for Learning*

2011-10 Bart Bogaert (UvT) *Cloud Content Contention*

2011-11 Dhaval Vyas (UT) *Designing for Awareness: An Experience-focused HCI Perspective*

2011-12 Carmen Bratosin (TUE) *Grid Architecture for Distributed Process Mining*

2011-13 Xiaoyu Mao (UvT) *Airport under Control. Multiagent Scheduling for Airport Ground Handling*

2011-14 Milan Lovric (EUR) *Behavioral Finance and Agent-Based Artificial Markets*

2011-15 Marijn Koolen (UvA) *The Meaning of Structure: the Value of Link Evidence for Information Retrieval*

2011-16 Maarten Schadd (UM) *Selective Search in Games of Different Complexity*

2011-17 Jiyin He (UVA) *Exploring Topic Structure: Coherence, Diversity and Relatedness*

2011-18 Mark Ponsen (UM) *Strategic Decision-Making in complex games*

2011-19 Ellen Rusman (OU) *The Mind's Eye on Personal Profiles*

2011-20 Qing Gu (VU) *Guiding service-oriented software engineering - A view-based approach*

2011-21 Linda Terlouw (TUD) *Modularization and Specification of Service-Oriented Systems*

2011-22 Junte Zhang (UVA) *System Evaluation of Archival Description and Access*

2011-23 Wouter Weerkamp (UVA) *Finding People and their Utterances in Social Media*

2011-24 Herwin van Welbergen (UT) *Behavior Generation for Interpersonal Coordination with Virtual Humans On Specifying, Scheduling and Realizing Multimodal Virtual Human Behavior*

2011-25 Syed Waqar ul Qounain Jaffry (VU)) *Analysis and Validation of Models for Trust Dynamics*

2011-26 Matthijs Aart Pontier (VU) *Virtual Agents for Human Communication - Emotion Regulation and Involvement-Distance Trade-Offs in Embodied Conversational Agents and Robots*

2011-27 Aniel Bhulai (VU) *Dynamic website optimization through autonomous management of design patterns*

2011-28 Rianne Kaptein (UVA) *Effective Focused Retrieval by Exploiting Query Context and Document Structure*

2011-29 Faisal Kamiran (TUE) *Discrimination-aware Classification*

2011-30 Egon van den Broek (UT) *Affective Signal Processing (ASP): Unraveling the mystery of emotions*

2011-31 Ludo Waltman (EUR) *Computational and Game-Theoretic Approaches for Modeling Bounded Rationality*

2011-32 Nees-Jan van Eck (EUR) *Methodological Advances in Bibliometric Mapping of Science*

2011-33 Tom van der Weide (UU) *Arguing to Motivate Decisions*

2011-34 Paolo Turrini (UU) *Strategic Reasoning in Interdependence: Logical and Game-theoretical Investigations*

2011-35 Maaike Harbers (UU) *Explaining Agent Behavior in Virtual Training*

2011-36 Erik van der Spek (UU) *Experiments in serious game design: a cognitive approach*

2011-37 Adriana Burlutiu (RUN) *Machine Learning for Pairwise Data, Applications for Preference Learning and Supervised Network Inference*

2011-38 Nyree Lemmens (UM) *Bee-inspired Distributed Optimization*

2011-39 Joost Westra (UU) *Organizing Adaptation using Agents in Serious Games*

2011-40 Viktor Clerc (VU) *Architectural Knowledge Management in Global Software Development*

2011-41 Luan Ibraimi (UT) *Cryptographically Enforced Distributed Data Access Control*

2011-42 Michal Sindlar (UU) *Explaining Behavior through Mental State Attribution*

2011-43 Henk van der Schuur (UU) *Process Improvement through Software Operation Knowledge*

2011-44 Boris Reuderink (UT) *Robust Brain-Computer Interfaces*

2011-45 Herman Stehouwer (UvT) *Statistical Language Models for Alternative Sequence Selection*

2011-46 Beibei Hu (TUD) *Towards Contextualized Information Delivery: A Rule-based Architecture for the Domain of Mobile Police Work*

2011-47 Azizi Bin Ab Aziz (VU) *Exploring Computational Models for Intelligent Support of Persons with Depression*

2011-48 Mark Ter Maat (UT) *Response Selection and Turn-taking for a Sensitive Artificial Listening Agent*

2011-49 Andreea Niculescu (UT) *Conversational interfaces for task-oriented spoken dialogues: design aspects influencing interaction quality*

2012-01 Terry Kakeeto (UvT) *Relationship Marketing for SMEs in Uganda*

2012-02 Muhammad Umair (VU) *Adaptivity, emotion, and Rationality in Human and Ambient Agent Models*

2012-03 Adam Vanya (VU) *Supporting Architecture Evolution by Mining Software Repositories*

2012-04 Jurriaan Souer (UU) *Development of Content Management System-based Web Applications*

2012-05 Marijn Plomp (UU) *Maturing Interorganisational Information Systems*

2012-06 Wolfgang Reinhardt (OU) *Awareness Support for Knowledge Workers in Research Networks*

2012-07 Rianne van Lambalgen (VU) *When the Going Gets Tough: Exploring Agent-based Models of Human Performance under Demanding Conditions*

2012-08 Gerben de Vries (UVA) *Kernel Methods for Vessel Trajectories*

2012-09 Ricardo Neisse (UT) *Trust and Privacy Management Support for Context-Aware Service Platforms*

2012-10 David Smits (TUE) *Towards a Generic Distributed Adaptive Hypermedia Environment*

2012-11 J.C.B. Rantham Prabhakara (TUE) *Process Mining in the Large: Preprocessing, Discovery, and Diagnostics*

2012-12 Kees van der Sluijs (TUE) *Model Driven Design and Data Integration in Semantic Web Information Systems*

2012-13 Suleman Shahid (UvT) *Fun and Face: Exploring non-verbal expressions of emotion during playful interactions*

2012-14 Evgeny Knutov (TUE) *Generic Adaptation Framework for Unifying Adaptive Web-based Systems*

2012-15 Natalie van der Wal (VU) *Social Agents. Agent-Based Modelling of Integrated Internal and Social Dynamics of Cognitive and Affective Processes.*

2012-16 Fiemke Both (VU) *Helping people by understanding them - Ambient Agents supporting task execution and depression treatment*

2012-17 Amal Elgammal (UvT) *Towards a Comprehensive Framework for Business Process Compliance*

2012-18 Eltjo Poort (VU) *Improving Solution Architecting Practices*

2012-19 Helen Schonenberg (TUE) *What's Next? Operational Support for Business Process Execution*

2012-20 Ali Bahramisharif (RUN) *Covert Visual Spatial Attention, a Robust Paradigm for Brain-Computer Interfacing*

2012-21 Roberto Cornacchia (TUD) *Querying Sparse Matrices for Information Retrieval*

2012-22 Thijs Vis (UvT) *Intelligence, politie en veiligheidsdienst: verenigbare grootheden?*

2012-23 Christian Muehl (UT) *Toward Affective Brain-Computer Interfaces: Exploring the Neurophysiology of Affect during Human Media Interaction*

2012-24 Laurens van der Werff (UT) *Evaluation of Noisy Transcripts for Spoken Document Retrieval*

2012-25 Silja Eckartz (UT) *Managing the Business Case Development in Inter-Organizational IT Projects: A Methodology and its Application*

2012-26 Emile de Maat (UVA) *Making Sense of Legal Text*

2012-27 Hayrettin Gurkok (UT) *Mind the Sheep! User Experience Evaluation & Brain-Computer Interface Games*

2012-28 Nancy Pascall (UvT) *Engendering Technology Empowering Women*

2012-29 Almer Tigelaar (UT) *Peer-to-Peer Information Retrieval*

2012-30 Alina Pommeranz (TUD) *Designing Human-Centered Systems for Reflective Decision Making*

2012-31 Emily Bagarukayo (RUN) *A Learning by Construction Approach for Higher Order Cognitive Skills Improvement, Building Capacity and Infrastructure*

2012-32 Wietske Visser (TUD) *Qualitative multi-criteria preference representation and reasoning*

2012-33 Rory Sie (OUN) *Coalitions in Cooperation Networks (COCOON)*

2012-34 Pavol Jancura (RUN) *Evolutionary analysis in PPI networks and applications*

2012-35 Evert Haasdijk (VU) *Never Too Old To Learn – On-line Evolution of Controllers in Swarm- and Modular Robotics*

2012-36 Denis Ssebugwawo (RUN) *Analysis and Evaluation of Collaborative Modeling Processes*

2012-37 Agnes Nakakawa (RUN) *A Collaboration Process for Enterprise Architecture Creation*

2012-38 Selmar Smit (VU) *Parameter Tuning and Scientific Testing in Evolutionary Algorithms*

2012-39 Hassan Fatemi (UT) *Risk-aware design of value and coordination networks*

2012-40 Agus Gunawan (UvT) *Information Access for SMEs in Indonesia*

**2012-41** Sebastian Kelle (OU) *Game Design Patterns for Learning*

2012-42 Dominique Verpoorten (OU) *Reflection Amplifiers in self-regulated Learning*

2012-43 Withdrawn

2012-44 Anna Tordai (VU) *On Combining Alignment Techniques*

2012-45 Benedikt Kratz (UvT) *A Model and Language for Business-aware Transactions*

2012-46 Simon Carter (UVA) *Exploration and Exploitation of Multilingual Data for Statistical Machine Translation*

2012-47 Manos Tsagkias (UVA) *Mining Social Media: Tracking Content and Predicting Behavior*

2012-48 Jorn Bakker (TUE) *Handling Abrupt Changes in Evolving Time-series Data*

2012-49 Michael Kaisers (UM) *Learning against Learning - Evolutionary dynamics of reinforcement learning algorithms in strategic interactions*

2012-50 Steven van Kervel (TUD) *Ontology driven Enterprise Information Systems Engineering*

2012-51 Jeroen de Jong (TUD) *Heuristics in Dynamic Sceduling; a practical framework with a case study in elevator dispatching*

2013-01 Viorel Milea (EUR) *News Analytics for Financial Decision Support*

2013-02 Erietta Liarou (CWI) *MonetDB/DataCell: Leveraging the Column-store Database Technology for Efficient and Scalable Stream Processing*

2013-03 Szymon Klarman (VU) *Reasoning with Contexts in Description Logics*

2013-04 Chetan Yadati (TUD) *Coordinating autonomous planning and scheduling*

2013-05 Dulce Pumareja (UT) *Groupware Requirements Evolutions Patterns*

2013-06 Romulo Goncalves (CWI) *The Data Cyclotron: Juggling Data and Queries for a Data Warehouse Audience*

2013-07 Giel van Lankveld (UvT) *Quantifying Individual Player Differences*

2013-08 Robbert-Jan Merk (VU) *Making enemies: cognitive modeling for opponent agents in fighter pilot simulators*

2013-09 Fabio Gori (RUN) *Metagenomic Data Analysis: Computational Methods and Applications*

2013-10 Jeewanie Jayasinghe Arachchige (UvT) *A Unified Modeling Framework for Service Design.*

2013-11 Evangelos Pournaras (TUD) *Multi-level Reconfigurable Self-organization in Overlay Services*

2013-12 Marian Razavian (VU) *Knowledge-driven Migration to Services*

2013-13 Mohammad Safiri (UT) *Service Tailoring: User-centric creation of integrated IT-based homecare services to support independent living of elderly*

2013-14 Jafar Tanha (UVA) *Ensemble Approaches to Semi-Supervised Learning Learning*

2013-15 Daniel Hennes (UM) *Multiagent Learning - Dynamic Games and Applications*

2013-16 Eric Kok (UU) *Exploring the practical benefits of argumentation in multi-agent deliberation*

2013-17 Koen Kok (VU) *The PowerMatcher: Smart Coordination for the Smart Electricity Grid*

2013-18 Jeroen Janssens (UvT) *Outlier Selection and One-Class Classification*

2013-19 Renze Steenhuizen (TUD) *Coordinated Multi-Agent Planning and Scheduling*

2013-20 Katja Hofmann (UvA) *Fast and Reliable Online Learning to Rank for Information Retrieval*

2013-21 Sander Wubben (UvT) *Text-to-text generation by monolingual machine translation*

2013-22 Tom Claassen (RUN) *Causal Discovery and Logic*

2013-23 Patricio de Alencar Silva (UvT) *Value Activity Monitoring*

2013-24 Haitham Bou Ammar (UM) *Automated Transfer in Reinforcement Learning*

2013-25 Agnieszka Anna Latoszek-Berendsen (UM) *Intention-based Decision Support. A new way of representing and implementing clinical guidelines in a Decision Support System*

2013-26 Alireza Zarghami (UT) *Architectural Support for Dynamic Homecare Service Provisioning*

2013-27 Mohammad Huq (UT) *Inference-based Framework Managing Data Provenance*

2013-28 Frans van der Sluis (UT) *When Complexity becomes Interesting: An Inquiry into the Information eXperience*

2013-29 Iwan de Kok (UT) *Listening Heads*

2013-30 Joyce Nakatumba (TUE) *Resource-Aware Business Process Management: Analysis and Support*

2013-31 Dinh Khoa Nguyen (UvT) *Blueprint Model and Language for Engineering Cloud Applications*

2013-32 Kamakshi Rajagopal (OUN) *Networking For Learning; The role of Networking in a Lifelong Learner's Professional Development*

2013-33 Qi Gao (TUD) *User Modeling and Personalization in the Microblogging Sphere*

2013-34 Kien Tjin-Kam-Jet (UT) *Distributed Deep Web Search*

2013-35 Abdallah El Ali (UvA) *Minimal Mobile Human Computer Interaction*

2013-36 Than Lam Hoang (TUe) *Pattern Mining in Data Streams*

2013-37 Dirk Börner (OUN) *Ambient Learning Displays*

2013-38 Eelco den Heijer (VU) *Autonomous Evolutionary Art*

2013-39 Joop de Jong (TUD) *A Method for Enterprise Ontology based Design of Enterprise Information Systems*

2013-40 Pim Nijssen (UM) *Monte-Carlo Tree Search for Multi-Player Games*

2013-41 Jochem Liem (UVA) *Supporting the Conceptual Modelling of Dynamic Systems: A Knowledge Engineering Perspective on Qualitative Reasoning*

2013-42 Léon Planken (TUD) *Algorithms for Simple Temporal Reasoning*

2013-43 Marc Bron (UVA) *Exploration and Contextualization through Interaction and Concepts*

2014-01 Nicola Barile (UU) *Studies in Learning Monotone Models from Data*

2014-02 Fiona Tuliyano (RUN) *Combining System Dynamics with a Domain Modeling Method*

2014-03 Sergio Raul Duarte Torres (UT) *Information Retrieval for Children: Search Behavior and Solutions*

2014-04 Hanna Jochmann-Mannak (UT) *Websites for children: search strategies and interface design - Three studies on children's search performance and evaluation*

2014-05 Jurriaan van Reijsen (UU) *Knowledge Perspectives on Advancing Dynamic Capability*

2014-06 Damian Tamburri (VU) *Supporting Networked Software Development*

2014-07 Arya Adriansyah (TUE) *Aligning Observed and Modeled Behavior*

2014-08 Samur Araujo (TUD) *Data Integration over Distributed and Heterogeneous Data Endpoints*

2014-09 Philip Jackson (UvT) *Toward Human-Level Artificial Intelligence: Representation and Computation of Meaning in Natural Language*

2014-10 Ivan Salvador Razo Zapata (VU) *Service Value Networks*

2014-11 Janneke van der Zwaan (TUD) *An Empathic Virtual Buddy for Social Support*

2014-12 Willem van Willigen (VU) *Look Ma, No Hands: Aspects of Autonomous Vehicle Control*

2014-13 Arlette van Wissen (VU) *Agent-Based Support for Behavior Change: Models and Applications in Health and Safety Domains*

2014-14 Yangyang Shi (TUD) *Language Models With Meta-information*

2014-15 Natalya Mogles (VU) *Agent-Based Analysis and Support of Human Functioning in Complex Socio-Technical Systems: Applications in Safety and Healthcare*

2014-16 Krystyna Milian (VU) *Supporting trial recruitment and design by automatically interpreting eligibility criteria*

2014-17 Kathrin Dentler (VU) *Computing healthcare quality indicators automatically: Secondary Use of Patient Data and Semantic Interoperability*

2014-18 Mattijs Ghijsen (UVA) *Methods and Models for the Design and Study of Dynamic Agent Organizations*

2014-19 Vinicius Ramos (TUE) *Adaptive Hypermedia Courses: Qualitative and Quantitative Evaluation and Tool Support*

2014-20 Mena Habib (UT) *Named Entity Extraction and Disambiguation for Informal Text: The Missing Link*

2014-21 Kassidy Clark (TUD) *Negotiation and Monitoring in Open Environments*

2014-22 Marieke Peeters (UU) *Personalized Educational Games - Developing agent-supported scenario-based training*

2014-23 Eleftherios Sidirourgos (UvA/CWI) *Space Efficient Indexes for the Big Data Era*

2014-24 Davide Ceolin (VU) *Trusting Semi-structured Web Data*

2014-25 Martijn Lappenschaar (RUN) *New network models for the analysis of disease interaction*

2014-26 Tim Baarslag (TUD) *What to Bid and When to Stop*

2014-27 Rui Jorge Almeida (EUR) *Conditional Density Models Integrating Fuzzy and Probabilistic Representations of Uncertainty*

2014-28 Anna Chmielowiec (VU) *Decentralized k-Clique Matching*

2014-29 Jaap Kabbedijk (UU) *Variability in Multi-Tenant Enterprise Software*

2014-30 Peter de Cock (UvT) *Anticipating Criminal Behaviour*

2014-31 Leo van Moergestel (UU) *Agent Technology in Agile Multiparallel Manufacturing and Product Support*

2014-32 Naser Ayat (UvA) *On Entity Resolution in Probabilistic Data*

2014-33 Tesfa Tegegne (RUN) *Service Discovery in eHealth*

2014-34 Christina Manteli (VU) *The Effect of Governance in Global Software Development: Analyzing Transactive Memory Systems.*

2014-35 Joost van Ooijen (UU) *Cognitive Agents in Virtual Worlds: A Middleware Design Approach*

2014-36 Joos Buijs (TUE) *Flexible Evolutionary Algorithms for Mining Structured Process Models*

2014-37 Maral Dadvar (UT) *Experts and Machines United Against Cyberbullying*

2014-38 Danny Plass-Oude Bos (UT) *Making brain-computer interfaces better: improving usability through post-processing.*

2014-39 Jasmina Maric (UvT) *Web Communities, Immigration, and Social Capital*

2014-40 Walter Omona (RUN) *A Framework for Knowledge Management Using ICT in Higher Education*

2014-41 Frederic Hogenboom (EUR) *Automated Detection of Financial Events in News Text*

2014-42 Carsten Eijckhof (CWI/TUD) *Contextual Multidimensional Relevance Models*

2014-43 Kevin Vlaanderen (UU) *Supporting Process Improvement using Method Increments*

2014-44 Paulien Meesters (UvT) *Intelligent Blauw. Met als ondertitel: Intelligence-gestuurde politiezorg in gebiedsgebonden eenheden.*

2014-45 Birgit Schmitz (OUN) *Mobile Games for Learning: A Pattern-Based Approach*

2014-46 Ke Tao (TUD) *Social Web Data Analytics: Relevance, Redundancy, Diversity*

2014-47 Shangsong Liang (UVA) *Fusion and Diversification in Information Retrieval*

2015-01 Niels Netten (UvA) *Machine Learning for Relevance of Information in Crisis Response*

2015-02 Faiza Bukhsh (UvT) *Smart auditing: Innovative Compliance Checking in Customs Controls*

2015-03 Twan van Laarhoven (RUN) *Machine learning for network data*

2015-04 Howard Spoelstra (OUN) *Collaborations in Open Learning Environments*

2015-05 Christoph Bösch (UT) *Cryptographically Enforced Search Pattern Hiding*

2015-06 Farideh Heidari (TUD) *Business Process Quality Computation - Computing Non-Functional Requirements to Improve Business Processes*

2015-07 Maria-Hendrike Peetz (UvA) *Time-Aware Online Reputation Analysis*

2015-08 Jie Jiang (TUD) *Organizational Compliance: An agent-based model for designing and evaluating organizational interactions*

2015-09 Randy Klaassen (UT) *HCI Perspectives on Behavior Change Support Systems*

2015-10 Henry Hermans (OUN) *OpenU: design of an integrated system to support lifelong learning*

2015-11 Yongming Luo (TUE) *Designing algorithms for big graph datasets: A study of computing bisimulation and joins*

2015-12 Julie M. Birkholz (VU) *Modi Operandi of Social Network Dynamics: The Effect of Context on Scientific Collaboration Networks*

2015-13 Giuseppe Procaccianti (VU) *Energy-Efficient Software*

2015-14 Bart van Straalen (UT) *A cognitive approach to modeling bad news conversations*

2015-15 Klaas Andries de Graaf (VU) *Ontology-based Software Architecture Documentation*

2015-16 Changyun Wei (UT) *Cognitive Coordination for Cooperative Multi-Robot Teamwork*

2015-17 André van Cleeff (UT) *Physical and Digital Security Mechanisms: Properties, Combinations and Trade-offs*

2015-18 Holger Pirk (CWI) *Waste Not, Want Not! - Managing Relational Data in Asymmetric Memories*

2015-19 Bernardo Tabuenca (OUN) *Ubiquitous Technology for Lifelong Learners*

2015-20 Loïs Vanhée (UU) *Using Culture and Values to Support Flexible Coordination*

2015-21 Sibren Fetter (OUN) *Using Peer-Support to Expand and Stabilize Online Learning*

2015-23 Luit Gazendam (VU) *Cataloguer Support in Cultural Heritage*

2015-24 Richard Berendsen (UVA) *Finding People, Papers, and Posts: Vertical Search Algorithms and Evaluation*

2015-25 Steven Woudenberg (UU) *Bayesian Tools for Early Disease Detection*

2015-26 Alexander Hogenboom (EUR) *Sentiment Analysis of Text Guided by Semantics and Structure*

2015-27 Sándor Héman (CWI) *Updating compressed colomn stores*

2015-28 Janet Bagorogoza (TiU) *KNOWLEDGE MANAGEMENT AND HIGH PERFORMANCE; The Uganda Financial Institutions Model for HPO*

2015-29 Hendrik Baier (UM) *Monte-Carlo Tree Search Enhancements for One-Player and Two-Player Domains*

2015-30 Kiavash Bahreini (OU) *Real-time Multimodal Emotion Recognition in E-Learning*

2015-31 Yakup Koç (TUD) *On the robustness of Power Grids*

2015-32 Jerome Gard (UL) *Corporate Venture Management in SMEs*

2015-33 Frederik Schadd (TUD) *Ontology Mapping with Auxiliary Resources*

2015-34 Victor de Graaf (UT) *Gesocial Recommender Systems*

2015-35 Jungxao Xu (TUD) *Affective Body Language of Humanoid Robots: Perception and Effects in Human Robot Interaction*

2016-01 Syed Saiden Abbas (RUN) *Recognition of Shapes by Humans and Machines*

2016-02 Michiel Christiaan Meulendijk (UU) *Optimizing medication reviews through decision support: prescribing a better pill to swallow*

2016-03 Maya Sappelli (RUN) *Knowledge Work in Context: User Centered Knowledge Worker Support*

2016-04 Laurens Rietveld (VU) *Publishing and Consuming Linked Data*

2016-05 Evgeny Sherkhonov (UVA) *Expanded Acyclic Queries: Containment and an Application in Explaining Missing Answers*

2016-06 Michel Wilson (TUD) *Robust scheduling in an uncertain environment*

2016-07 Jeroen de Man (VU) *Measuring and modeling negative emotions for virtual training*

2016-08 Matje van de Camp (TiU) *A Link to the Past: Constructing Historical Social Networks from Unstructured Data*

2016-09 Archana Nottamkandath (VU) *Trusting Crowdsourced Information on Cultural Artefacts*

2016-10 George Karafotias (VUA) *Parameter Control for Evolutionary Algorithms*

2016-11 Anne Schuth (UVA) *Search Engines that Learn from Their Users*

2016-12 Max Knobbout (UU) *Logics for Modelling and Verifying Normative Multi-Agent Systems*

2016-13 Nana Baah Gyan (VU) *The Web, Speech Technologies and Rural Development in West Africa - An ICT4D Approach*

2016-14 Ravi Khadka (UU) *Revisiting Legacy Software System Modernization*

2016-15 Steffen Michels (RUN) *Hybrid Probabilistic Logics - Theoretical Aspects, Algorithms and Experiments*

2016-16 Guangliang Li (UVA) *Socially Intelligent Autonomous Agents that Learn from Human Reward*

2016-17 Berend Weel (VU) *Towards Embodied Evolution of Robot Organisms*

2016-18 Albert Meroño Peñuela (VU) *Refining Statistical Data on the Web*

2016-19 Julia Efremova (Tu/e) *Mining Social Structures from Genealogical Data*

2016-20 Daan Odijk (UVA) *Context & Semantics in News & Web Search*

2016-21 Alejandro Moreno Célleri (UT) *From Traditional to Interactive Playspaces: Automatic Analysis of Player Behavior in the Interactive Tag Playground*

2016-22 Grace Lewis (VU) *Software Architecture Strategies for Cyber-Foraging Systems*

2016-23 Fei Cai (UVA) *Query Auto Completion in Information Retrieval*

2016-24 Brend Wanders (UT) *Repurposing and Probabilistic Integration of Data; An Iterative and data model independent approach*

2016-25 Julia Kiseleva (TU/e) *Using Contextual Information to Understand Searching and Browsing Behavior*

2016-26 Dilhan Thilakarathne (VU) *In or Out of Control: Exploring Computational Models to Study the Role of Human Awareness and Control in Behavioural Choices, with Applications in Aviation and Energy Management Domains*

2016-27 Wen Li (TUD) *Understanding Geo-spatial Information on Social Media*

2016-28 Mingxin Zhang (TUD) *Large-scale Agent-based Social Simulation - A study on epidemic prediction and control*

2016-29 Nicolas Höning (TUD) *Peak reduction in decentralised electricity systems -Markets and prices for flexible planning*

2016-30 Ruud Mattheij (UvT) *The Eyes Have It*

2016-31 Mohammad Khelghati (UT) *Deep web content monitoring*

2016-32 Eelco Vriezekolk (UT) *Assessing Telecommunication Service Availability Risks for Crisis Organisations*

2016-33 Peter Bloem (UVA) *Single Sample Statistics, exercises in learning from just one example*

2016-34 Dennis Schunselaar (TUE) *Configurable Process Trees: Elicitation, Analysis, and Enactment*

2016-35 Zhaochun Ren (UVA) *Monitoring Social Media: Summarization, Classification and Recommendation*

2016-36 Daphne Karreman (UT) *Beyond R2D2: The design of nonverbal interaction behavior optimized for robot-specific morphologies*

2016-37 Giovanni Sileno (UvA) *Aligning Law and Action - a conceptual and computational inquiry*

2016-38 Andrea Minuto (UT) *MATERIALS THAT MATTER - Smart Materials meet Art & Interaction Design*

2016-39 Merijn Bruijnes (UT) *Believable Suspect Agents; Response and Interpersonal Style Selection for an Artificial Suspect*

2016-40 Christian Detweiler (TUD) *Accounting for Values in Design*

2016-41 Thomas King (TUD) *Governing Governance: A Formal Framework for Analysing Institutional Design and Enactment Governance*

2016-42 Spyros Martzoukos (UVA) *Combinatorial and Compositional Aspects of Bilingual Aligned Corpora*

2016-43 Saskia Koldijk (RUN) *Context-Aware Support for Stress Self-Management: From Theory to Practice*

2016-44 Thibault Sellam (UVA) *Automatic Assistants for Database Exploration*

2016-45 Bram van de Laar (UT) *Experiencing Brain-Computer Interface Control*

2016-46 Jorge Gallego Perez (UT) *Robots to Make you Happy*

2016-47 Christina Weber (UL) *Real-time foresight - Preparedness for dynamic innovation networks*

2016-48 Tanja Buttler (TUD) *Collecting Lessons Learned*

2016-49 Gleb Polevoy (TUD) *Participation and Interaction in Projects. A Game-Theoretic Analysis*

2016-50 Yan Wang (UVT) *The Bridge of Dreams: Towards a Method for Operational Performance Alignment in IT-enabled Service Supply Chains*

2017-01 Jan-Jaap Oerlemans (UL) *Investigating Cybercrime*

2017-02 Sjoerd Timmer (UU) *Designing and Understanding Forensic Bayesian Networks using Argumentation*

2017-03 Daniël Harold Telgen (UU) *Grid Manufacturing; A Cyber-Physical Approach with Autonomous Products and Reconfigurable Manufacturing Machines*

2017-04 Mrunal Gawade (CWI) *MULTI-CORE PARALLELISM IN A COLUMN-STORE*

2017-05 Mahdieh Shadi (UVA) *Collaboration Behavior*

2017-06 Damir Vandic (EUR) *Intelligent Information Systems for Web Product Search*

2017-07 Roel Bertens (UU) *Insight in Information: from Abstract to Anomaly*

2017-08 Rob Konijn (VU) *Detecting Interesting Differences:Data Mining in Health Insurance Data using Outlier Detection and Subgroup Discovery*

2017-09 Dong Nguyen (UT) *Text as Social and Cultural Data: A Computational Perspective on Variation in Text*

2017-10 Robby van Delden (UT) *(Steering) Interactive Play Behavior*

2017-11 Florian Kunneman (RUN) *Modelling patterns of time and emotion in Twitter #anticipointment*

2017-12 Sander Leemans (UT) *Robust Process Mining with Guarantees*

2017-13 Gijs Huisman (UT) *Social Touch Technology - Extending the reach of social touch through haptic technology*

2017-14 Shoshannah Tekofsky (UvT) *You Are Who You Play You Are: Modelling Player Traits from Video Game Behavior*

2017-15 Peter Berck, Radboud University (RUN) *Memory-Based Text Correction*

2017-16 Aleksandr Chuklin (UVA) *Understanding and Modeling Users of Modern Search Engines*

2017-17 Daniel Dimov (VU) *Crowdsourced Online Dispute Resolution*

2017-18 Ridho Reinanda (UVA) *Entity Associations for Search*

2017-19 Jeroen Vuurens (TUD) *Proximity of Terms, Texts and Semantic Vectors in Information Retrieval*

2017-20 Mohammadbashir Sedighi (TUD) *Fostering Engagement in Knowledge Sharing: The Role of Perceived Benefits, Costs and Visibility*

2017-21 Jeroen Linssen (UT) *Meta Matters in Interactive Storytelling and Serious Gaming (A Play on Worlds)*

2017-22 Sara Magliacane (VU) *Logics for causal inference under uncertainty*

2017-23 David Graus (UVA) *Entities of Interest— Discovery in Digital Traces*

2017-24 Chang Wang (TUD) *Use of Affordances for Efficient Robot Learning*

2017-25 Veruska Zamborlini (VU) *Knowledge Representation for Clinical Guidelines, with applications to Multimorbidity Analysis and Literature Search*

2017-26 Merel Jung (UT) *Socially Intelligent Robots that Understand and Respond to Human Touch*
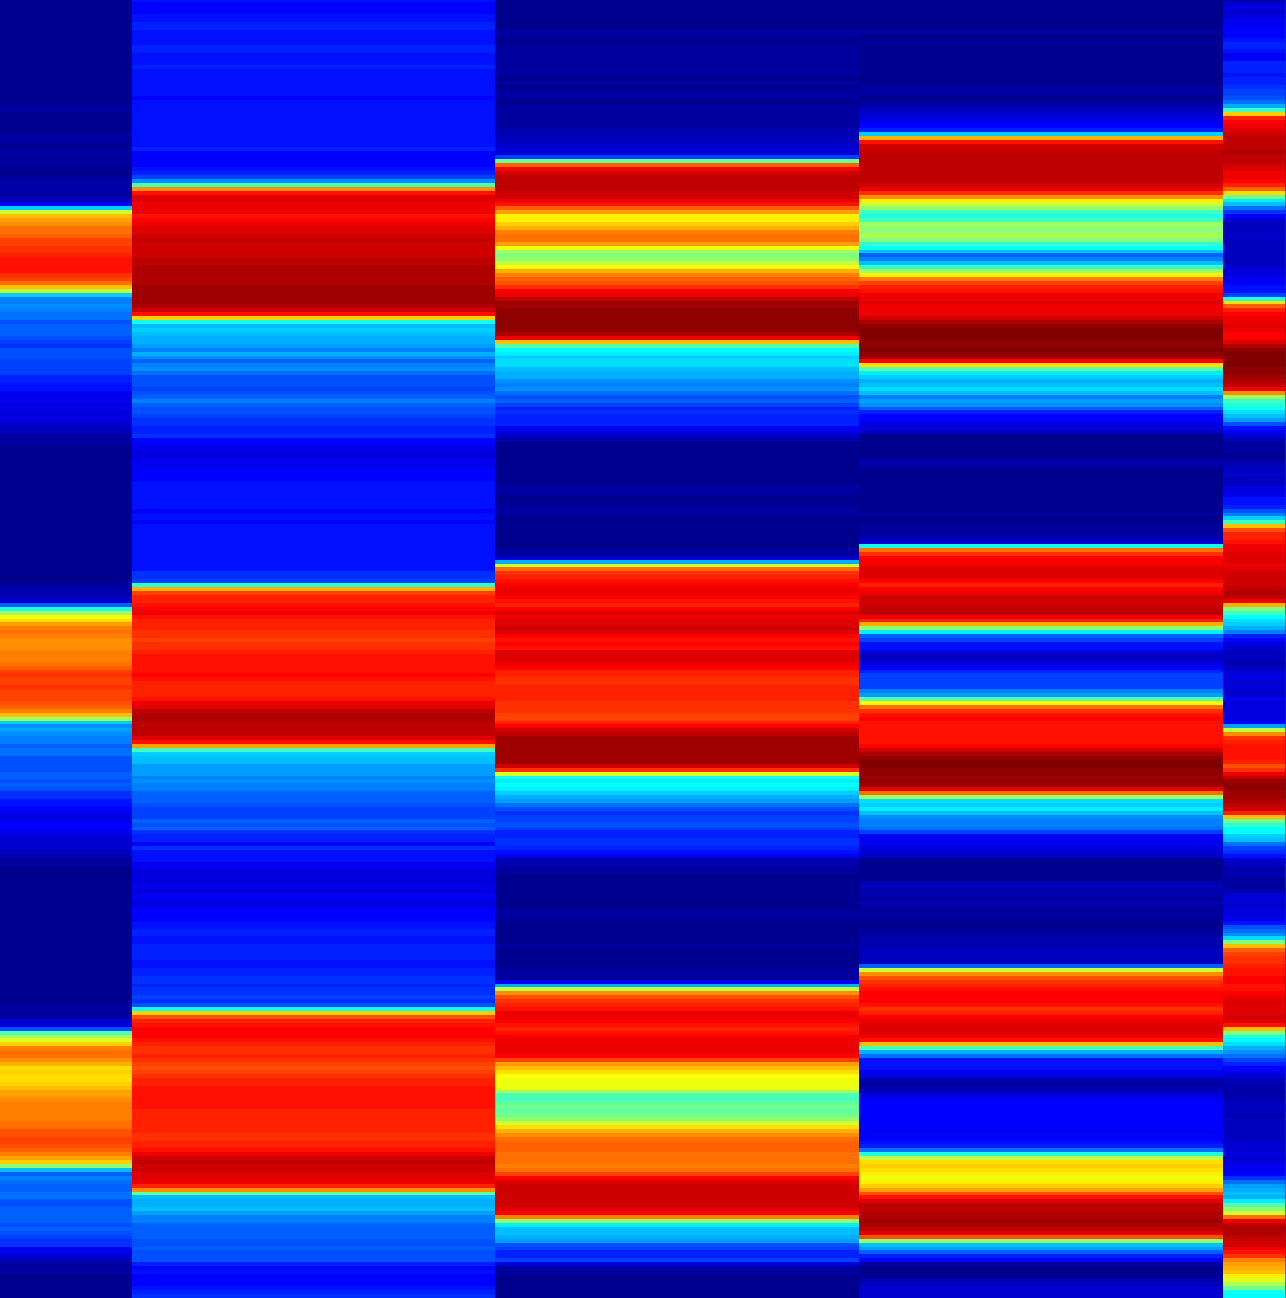
The cover shows a visualization of a rough stroking gesture performed on an 8x8 pressure sensor grid. The mean pressure per sensor grid column over time visualizes the displacement in the row direction. The colors of the heatmap range from blue which indicates low pressure to red which indicates high pressure.

CTIT