Full Papers Session

Tue, March 6

Full Papers Session 1, Track 1

Room W192 (10:50~12:10)

Teleoperation and Shared Manipulation

Session Chair: Henny Admoni (Carnegie Mellon University)

Eye-Hand Behavior in Human-Robot Shared Manipulation (10:50~11:10)

Reuben Aronson (Carnegie Mellon University), **Thiago Santini** (University of Tübingen), **Thomas Kübler** (University of Tübingen), **Enkelejda Kasneci** (University of Tübingen), **Siddhartha Srinivasa** (University of Washington), **Henny Admoni** (Carnegie Mellon University)

Shared autonomy systems enhance people's abilities to perform activities of daily living using robotic manipulators. Recent systems succeed by first identifying their operators' intentions, typically by analyzing the user's joystick input. To enhance this recognition, it is useful to characterize people's behavior while performing such a task. Furthermore, eye gaze is a rich source of information for understanding operator intention. The goal of this paper is to provide novel insights into the dynamics of control behavior and eye gaze in human-robot shared manipulation tasks. To achieve this goal, we conduct a data collection study that uses an eye tracker to record eye gaze during a human-robot shared manipulation activity, both with and without shared autonomy assistance. We process the gaze signals from the study to extract gaze features like saccades, fixations, smooth pursuits, and scan paths. We analyze those features to identify novel patterns of gaze behaviors and highlight where these patterns are similar to and different from previous findings about eye gaze in human-only manipulation tasks. The work described in this paper lays a foundation for a model of natural human eye gaze in human-robot shared manipulation.

"Wait, Can You Move the Robot?": Examining Telepresence Robot Use in Collaborative Teams (11:10~11:30)

Brett Stoll (Cornell University), **Samantha Reig** (Carnegie Mellon University), **Lucy He** (Cornell University), **Ian Kaplan** (Cornell University), **Malte Jung** (Cornell University), **Susan Fussell** (Cornell University)

Telepresence robots provide remote team members with embodied presence, but whether this improves remote teammate participation, remote users' perceptions of team collaboration, or collocated members' perceptions of remote teammates is an open question. We conducted an experiment in which teams of two collocated members and one telepresent (remote) member solved a word puzzle requiring a translation key. We varied who had access to the key to examine effects of resource accessibility in distributed groups: in the Robot Information condition, the remote pilot (RP) possessed the key; in the Shared Information condition, all team members possessed the key; in the Local Information condition, only collocated participants (CPs) possessed the key. Audio transcripts were analyzed for differences in the number of words spoken by each team member. RPs spoke significantly less than CPs, especially when they lacked the translation key. RPs perceived greater task difficulty and less ease of communication than CPs. CPs rated other CPs as more trustworthy than RPs. This suggests an imbalance between collocated and remote collaborators that can negatively affect collaboration. We discuss implications for the design and use of telepresence robots in the workplace.

Shared Dynamic Curves: A Shared-Control Telemanipulation Method for Motor Task Training (11:30~11:50)

Daniel Rakita (University of Wisconsin-

Madison), **Bilge Mutlu** (University of Wisconsin-Madison), **Michael Gleicher** (University of Wisconsin-Madison), **Laura Hiatt** (Naval Research Laboratory)

In this paper, we present a novel shared-control telemanipulation method that is designed to incrementally improve a user's motor ability. Our method initially corrects for the user's suboptimal control trajectories, gradually giving the user more direct control over a series of training trials as he/she naturally gets more accustomed to the task. Our shared-control method, called *Shared Dynamic Curves*, blends suboptimal user translation and rotation control inputs with known translation and rotation paths needed to complete a task. Shared Dynamic Curves provide a translation and rotation path in space along which the user can easily guide the robot, and this curve can bend and flex in real-time as a dynamical system to pull the user's motion gracefully toward a goal. We show through a user study that Shared Dynamic Curves affords effective motor learning on certain tasks compared to alternative training methods. We discuss our findings in the context of shared control and speculate on how this method could be applied in real-world scenarios such as job training or stroke rehabilitation.

It's All in Your Head: Using Priming to Shape an Operator's Perceptions and Behavior during Teleoperation (11:50~12:10)

Daniel Rea (University of Manitoba), **James Young** (University of Manitoba)

Perceptions of a technology can shape the way the technology is used and adopted. Thus, in teleoperation, it is important to understand how a teleoperator's perceptions of a robot can be shaped, and whether those perceptions can impact how people drive robots. Priming, evoking activity in a person by exposing them to learned stimuli, is one way of shaping someone's perception. We investigate priming an operator's impression of a robot's physical capabilities in order to impact their perception of the robot and teleoperation behavior; that is, we examine if we can change operator driving behavior simply by making them believe that a robot is dangerous or safe, fast or slow, etc., without actually changing robot capability. Our results show that priming (with no change to robot behavior or capability) can impact operator perception of the robot, their teleoperation experience, and in some cases may impact teleoperation performance.

Full Papers Session 1, Track 2

Room W194 (10:50~12:10)

Tutoring and Child-Robot Interaction

Session Chair: Kerstin Dautenhahn (University of Hertfordshire)

Do Children Perceive Whether a Robotic Peer is Learning or Not? (10:50~11:10)

Shruti Chandra (INESC-ID, Instituto Superior Tecnico, University of Lisbon & ÉPFL), **Raul Paradeda** (INESC-ID, Instituto Superior Tecnico, University of Lisbon & State University of Rio Grande do Norte), **Hang Yin** (INESC-ID, Instituto Superior Tecnico, University of Lisbon & ÉPFL), **Pierre Dillenbourg** (ÉPFL), **Rui Prada** (INESC-ID, Instituto Superior Técnico, Universidade de Lisbon), **Ana Paiva** (INESC-ID, Instituto Superior Técnico & University of Lisbon)

Social robots are being used to create better educational scenarios, thereby fostering children's learning. In the work presented here, we describe an autonomous social robot that was designed to enhance children's handwriting skills. Exploiting the benefits of the learning-by-teaching method, the system provides a scenario in which a child acts as a teacher and corrects the handwriting difficulties of the robotic agent. To explore the children's perception towards this social robot and the effect on their learning, we have conducted a multi-session study with children that compared two contrasting competencies in the robot: 'learning' vs 'non-learning' and presented as two conditions in the study. The results suggest that the children learned more

in the learning condition compared with the non-learning condition and their learning gains seem to be affected by their perception of the robot. The results did not lead to any significant differences in the children's perception of the robot in the first two weeks of interaction. However, by the end of the 4th week, the results changed. The children in the learning condition gave significantly higher writing ability and overall performance scores to the robot compared with the non-learning condition. In addition, the change in the robot's learning capabilities did not show to affect their perceived intelligence, likability and friendliness towards it.

The Effect of a Robot's Gestures and Adaptive Tutoring on Children's Acquisition of Second Language Vocabularies (11:10~11:30)

Jan de Wit (Tilburg University), **Thorsten Schodde** (Bielefeld University), **Bram Willemsen** (Tilburg University), **Kirsten Bergmann** (Bielefeld University), **Mirjam de Haas** (Tilburg University), **Stefan Kopp** (Bielefeld University), **Emiel Krahmer** (Tilburg University), **Paul Vogt** (Tilburg University)

This paper presents a study in which children, four to six years old, were taught words in a second language by a robot tutor. The goal is to evaluate two ways for a robot to provide scaffolding for students: the use of iconic gestures, combined with adaptively choosing the next learning task based on the child's past performance. The results show a positive effect on long-term memorization of novel words, and an overall higher level of engagement during the learning activities when gestures are used. The adaptive tutoring strategy reduces the extent to which the level of engagement is diminishing during the later part of the interaction.

Thinking Aloud with a Tutoring Robot to Enhance Learning (11:30~11:50)

Aditi Ramachandran (Yale University), **Chien-Ming Huang** (Johns Hopkins University), **Edward Gartland** (Greens Farms Academy), **Brian Scassellati** (Yale University)

Thinking aloud, while requiring extra mental effort, is a metacognitive technique that helps students navigate through complex problem-solving tasks. Social robots, bearing embodied immediacy that fosters engaging and compliant interactions, are a unique platform to deliver problem-solving support such as thinking aloud to young learners. In this work, we explore the effects of a robot platform and the think-aloud strategy on learning outcomes in the context of a one-on-one tutoring interaction. Results from a 2x2 between-subjects study (n=52) indicate that both the robot platform and use of the think-aloud strategy promoted learning gains for children. In particular, the robot platform effectively enhanced immediate learning gains, measured right after the tutoring session, while the think-aloud strategy improved persistent gains as measured approximately one week after the interaction. Moreover, our results show that a social robot strengthened students' engagement and compliance with the think-aloud support while they performed cognitively demanding tasks. Our work indicates that robots can support metacognitive strategy use to effectively enhance learning and contributes to the growing body of research demonstrating the value of social robots in novel educational settings.

"Stop. I see a conflict happening." A Robot Mediator for Young Children's Interpersonal Conflict Resolution (11:50~12:10)

Solace Shen (Cornell University), **Petr Slovak** (University College London), **Malte Jung** (Cornell University)

The ability to constructively resolve interpersonal conflicts is a crucial set of social skills people need to effectively work and live well together. Is it possible to design social robots to support the early development of children's interpersonal conflict resolution skills? To investigate this question, 64 (32 pairs of) children ages 3-6 years engaged in a 50-minute play session consisting of 5 activities facilitated by the robot Keepon. Children were randomly assigned to 1 of 2 conditions. In the mediation condition, Keepon directed the play session flow by indicating when and which activity to switch to, and

whenever possible, signaled the onset of object possession conflicts that occurred between the pair and offered prompts for constructive conflict resolution. In the control condition, Keepon only facilitated and directed the play session and did not intervene during children's conflicts. Results show that children were more likely to resolve conflicts constructively in the mediation condition than in the control condition, and that a key function for a robot mediator within the conflict process is to successfully flag the conflict onset. Drawing from these findings, we discuss design recommendation for a robot mediator.

Full Papers Session 2

Room W192 (16:00~17:40)

Best Paper Nominees

Session Chair: Fumihide Tanaka
(University of Tsukuba)

★★★★ Improving Collocated Robot Teleoperation with Augmented Reality (16:00~16:20)

Hooman Hedayati (University of Colorado Boulder), **Michael Walker** (University of Colorado Boulder), **Daniel Szafir** (University of Colorado Boulder)

Robot teleoperation can be a challenging task, often requiring a great deal of user training and expertise, especially for platforms with high degrees-of-freedom (e.g., industrial manipulators and aerial robots). Users often struggle to synthesize information robots collect (e.g., a camera stream) with contextual knowledge of how the robot is moving in the environment. We explore how advances in augmented reality (AR) technologies are creating a new design space for mediating robot teleoperation by enabling novel forms of intuitive, visual feedback. We prototype several aerial robot teleoperation interfaces using AR, which we evaluate in a 48-participant user study where participants completed an environmental inspection task. Our new interface designs provided several objective and subjective performance benefits over existing systems, which often force users into an undesirable paradigm that divides user attention between monitoring the robot and monitoring the robot's camera feed(s).

★★★★ Expressing Robot Incapability (16:20~16:40)

Minae Kwon (Cornell University), **Sandy Huang** (University of California Berkeley), **Anca Dragan** (University of California Berkeley)

Our goal is to enable robots to express their incapability, and to do so in a way that communicates both what they are trying to accomplish and why they are unable to accomplish it. We frame this as a trajectory optimization problem: maximize the similarity between the motion expressing incapability and what would amount to successful task execution, while obeying the physical limits of the robot. We introduce and evaluate candidate similarity measures, and show that one in particular generalizes to a range of tasks, while producing expressive motions that are tailored to each task. Our user study supports that our approach automatically generates motions expressing incapability that communicate both what and why to end-users, and improve their overall perception of the robot and willingness to collaborate with it in the future.

★★★★ Characterizing the Design Space of Rendered Robot Faces (16:20~16:40)

Alisa Kalegina (University of Washington), **Grace Schroeder** (University of Washington), **Aidan Allchin** (Lakeside School), **Keara Berlin** (Macalester College), **Maya Cakmak** (University of Washington)

Faces are critical in establishing the agency of social robots; however, building expressive mechanical faces is costly and difficult. Instead, many robots built in recent years have faces that are rendered onto a screen. This gives great flexibility in what a robot's face can be and opens up a new design space with which to establish a robot's character and perceived properties. Despite the prevalence of robots with rendered faces, there are no systematic explorations of this design space. Our work aims to fill that gap. We conducted a survey and identified 157 robots with rendered faces and coded them in terms of 76 properties. We present statistics, common patterns, and observations about this data set of faces. Next, we conducted two surveys to understand people's perceptions of rendered robot faces and identify the impact of different face features. Survey results indicate preferences for varying levels of realism and detail in robot faces based on context, and indicate how the presence or absence of specific features affects perception of the face and the types of jobs the face would be appropriate for.

What is Human-like?: Decomposing Robots' Human-like Appearance Using the Anthropomorphic roBOT (ABOT) Database (16:40~17:00)

Elizabeth Phillips (Brown University), **Xuan Zhao** (Brown University), **Daniel Ullman** (Brown University), **Bertram Malle** (Brown University)

Anthropomorphic robots, or robots with human-like appearance features such as eyes, hands, or faces, have drawn considerable attention in recent years. To date, what makes a robot appear human-like has been driven by designers' and researchers' intuitions, because a systematic understanding of the range, variety, and relationships among constituent features of anthropomorphic robots is lacking. To fill this gap, we introduce the ABOT (Anthropomorphic roBOT) Database—a collection of 200 images of real-world robots with one or more human-like appearance features (http://www.abotdatabase.info). Harnessing this database, Study 1 uncovered four distinct appearance dimensions (i.e., bundles of features) that characterize a wide spectrum of anthropomorphic robots and Study 2 identified the dimensions and specific features that were most predictive of robots' perceived human-likeness. With data from both studies, we then created an online estimation tool to help researchers predict how human-like a new robot will be perceived given the presence of various appearance features. The present research sheds new light on what makes a robot look human, and makes publicly accessible a powerful new tool for future research on robots' human-likeness.

Fribo: A Social Networking Robot for Increasing Social Connectedness through Sharing Daily Home Activities from Living Noise Data (17:00~17:20)

Kwangmin Jeong (Yonsei University), **Jihyun Sung** (Yonsei University), **Hae-Sung Lee** (Yonsei University), **Aram Kim** (Yonsei University), **Hyemi Kim** (Yonsei University), **Chan Mi Park** (Yonsei University), **Yuin Jeong** (Yonsei University), **JeeHang Lee** (KAIST), **Jinwoo Kim** (Yonsei University)

The rapid increase in the number of young adults living alone gives rise to a demand for the resolution of social isolation problems. Social robot technologies play a substantial role for this purpose. However, existing technologies try to solve the problem only through one-to-one interaction with robots, which in turn fails to utilize the real-world social relationships. Privacy concern is an additional issue since most social robots rely on the visual information for the interactions. To this end, we propose 'Fribo', auditory information centered social robot that recognizes user's activity by analyzing occupants' living noise and shares the activity information with close friends. A four-week field study with the first prototype of Fribo confirms that activity sharing through the use of anonymized living noise promises a virtual cohabiting experience that triggers more frequent real-world social interactions with less feeling of privacy intrusion. Based on this finding and the further qualitative analysis, we suggest a design principle of sound-based social networking robots and its associated new interactions, then present the second prototype of Fribo inspired by the implications from the field study.

Full Papers Session

Wed, March 7

Full Papers Session 3, Track 1

Room W192 (10:30~12:10)

Machine Learning for HRI

Session Chair: Ross Knepper
(Cornell University)

Active Robot Learning for Temporal Task Models (10:30~10:50)

Mattia Racca (Aalto University), **Ville Kyrki** (Aalto University)

With the goal of having robots learn new skills after deployment, we propose an active learning framework for modelling user preferences about task execution. The proposed approach interactively gathers information by asking questions expressed in natural language. We study the validity and the learning performance of the proposed approach and two of its variants compared to a passive learning strategy. We further investigate the human-robot-interaction nature of the framework conducting a usability study with 18 subjects. The results show that active strategies are applicable for learning preferences in temporal tasks from non-expert users. Furthermore, the results provide insights in the interaction design of active learning robots.

Learning from Richer Human Guidance: Augmenting Comparison-Based Learning with Feature Queries (10:50~11:10)

Chandrayee Basu (University of California Merced), **Mukesh Singhal** (University of California Merced), **Anca Dragan** (University of California Berkeley)

We focus on learning the desired objective function for a robot. Although trajectory demonstrations can be very informative of the desired objective, they can also be difficult for users to provide. Answers to comparison queries, asking which of two trajectories is preferable, are much easier for users, and have emerged as an effective alternative. Unfortunately, comparisons are far less informative.
We propose that there is much richer information that users can easily provide and that robots ought to leverage. We focus on augmenting comparisons with feature queries, and introduce a unified formalism for treating all answers as observations about the true desired reward. We derive an active query selection algorithm, and test these queries in simulation and on real users. We find that richer, feature-augmented queries can extract more information faster, leading to robots that better match user preferences in their behavior.

Learning from Physical Human Corrections, One Feature at a Time (11:10~11:30)

Andrea Bajcsy (University of California Berkeley), **Dylan Losey** (Rice University), **Marcia O'Malley** (Rice University), **Anca Dragan** (University of California Berkeley)

We focus on learning robot objective functions from human guidance: specifically, from physical corrections provided by the person while the robot is acting. Objective functions are typically parametrized in terms of *features*, which capture aspects of the task that might be important. When the person intervenes to correct the robot's behavior, the robot should update its understanding of which features matter, how much, and in what way. Unfortunately, real users do not provide optimal corrections that isolate exactly what the robot was doing wrong. Thus, when receiving a correction, it is difficult for the robot to determine which features the person meant to correct, and which features were changed unintentionally. In this paper, we propose to improve the efficiency of robot learning during physical interactions by reducing *unintended* learning. Our approach allows the human-robot team to focus on learning *one feature at a time*, unlike state-of-the-art techniques that update all

features at once. We derive an online method for identifying the single feature which the human is trying to change during physical interaction, and experimentally compare this one-at-a-time approach to the all-at-once baseline in a user study. Our results suggest that users teaching one-at-a-time perform better, especially in tasks that require changing multiple features.

DNN-HMM based Automatic Speech Recognition for HRI Scenarios (11:30~11:50)

José Novoa (University of Chile), **Jorge Wuth** (University of Chile), **Juan Escudero** (University of Chile), **Josué Fredes** (University of Chile), **Rodrigo Mahu** (University of Chile), **Néstor Yoma** (University of Chile)

In this paper, we propose to replace the classical black box integration of automatic speech recognition technology in HRI applications with the incorporation of the HRI environment representation and modeling, and the robot and user states and contexts. Accordingly, this paper focuses on the environment representation and modeling by training a deep neural network-hidden Markov model based automatic speech recognition engine combining clean utterances with the acoustic-channel responses and noise that were obtained from an HRI testbed built with a PR2 mobile manipulation robot. This method avoids recording a training database in all the possible acoustic environments given an HRI scenario. Moreover, different speech recognition testing conditions were produced by recording two types of acoustics sources, i.e. a loudspeaker and human speakers, using a Microsoft Kinect mounted on top of the PR2 robot, while performing head rotations and movements towards and away from the fixed sources. In this generic HRI scenario, the resulting automatic speech recognition engine provided a word error rate that is at least 26% and 38% lower than publicly available speech recognition APIs with the playback (i.e. loudspeaker) and human testing databases, respectively, with a limited amount of training data.

Deep Reinforcement Learning of Abstract Reasoning from Demonstrations (11:50~12:10)

Madison Clark-Turner (University of New Hampshire), **Momotaz Begum** (University of New Hampshire)

Extracting a set of generalizable rules that govern the dynamics of complex, high-level interactions between humans based only on observations is a high-level cognitive ability. Mastery of this skill marks a significant milestone in the human developmental process. A key challenge in designing such an ability in autonomous robots is discovering the relationships among discriminatory features. Identifying features in natural scenes that are representative of a particular event or interaction (i.e. 'discriminatory features') and then discovering the relationships (e.g., temporal/spatial/spatio-temporal/causal) among those features in the form of generalized rules are non-trivial problems. They often appear as a 'chicken-and-egg' dilemma. This paper proposes an end-to-end learning framework to tackle these two problems in the context of learning generalized, high-level rules of human interactions from structured demonstrations. We employed our proposed deep reinforcement learning framework to learn a set of rules that govern a behavioral intervention session between two agents based on observations of several instances of the session. We also tested the accuracy of our framework with human subjects in diverse situations.

Full Papers Session 3, Track 2

Room W194 (10:30~12:10)

Societal Issues: Abuse, Trust, Racism

Session Chair: Cindy Bethel
(Mississippi University)

Inducing Bystander Interventions During Robot Abuse with Social Mechanisms (10:30~10:50)

Xiang Zhi Tan (Carnegie Mellon University), **Marynel Vázquez** (Stanford University),

Elizabeth Carter (Carnegie Mellon University), **Cecilia Morales** (Carnegie Mellon University), **Aaron Steinfeld** (Carnegie Mellon University)

We explored whether a robot can leverage social influences to motivate nearby bystanders to intervene and defend them from human abuse. We designed a between-subjects study where 48 participants took part in a memorization task and observed a confederate mistreating a robot both verbally and physically. The robot was either empathetic towards the participant's performance in the task or indifferent. When the robot was mistreated, it ignored the abuse, shut down in response to it, or reacted emotionally. We found that the majority of the participants intervened to help the robot after it was abused. Interventions happened for a wide range of reasons. Interestingly, the empathetic robot increased the proportion of participants that self-reported intervening in comparison to the indifferent robot, but more participants moved the robot as a response to abuse in the latter case. The participants also perceived the robot being verbally mistreated more and reported higher levels of personal distress when the robot briefly shut down after abuse in comparison to when it reacted emotionally or did not react at all.

The Ripple Effects of Vulnerability: The Effects of a Robot's Vulnerable Behavior on Trust in Human-Robot Teams (10:50~11:10)

Sarah Strohkorb Sebo (Yale University), **Margaret Traeger** (Yale University), **Malte Jung** (Cornell University), **Brian Scassellati** (Yale University)

Successful teams are characterized by high levels of trust between team members, allowing the team to learn from mistakes, take risks, and entertain diverse ideas. We investigated a robot's potential to shape trust within a team through the robot's expressions of vulnerability. We conducted a between-subjects experiment (N=35 teams, 105 participants) comparing the behavior of three human teammates collaborating with either a social robot making vulnerable statements or with a social robot making neutral statements. We found that, in a group with a robot making vulnerable statements, participants responded more to the robot's comments and directed more of their gaze to the robot, displaying a higher level of engagement with the robot. Additionally, we discovered that during times of tension, human teammates in a group with a robot making vulnerable statements were more likely to explain their failure to the group, console team members who had made mistakes, and laugh together, all actions that reduce the amount of tension experienced by the team. These results suggest that a robot's vulnerable behavior can have 'ripple effects' on their human team members' expressions of trust-related behavior.

Humans Conform to Robots: Disambiguating Trust, Truth, and Conformity (11:10~11:30)

Nicole Salomons (Yale University), **Michael van der Linden** (Yale University), **Sarah Strohkorb Sebo** (Yale University), **Brian Scassellati** (Yale University)

Asch's conformity experiment has shown that people are prone to adjusting their view to match those of group members even when they believe the answer of the group to be wrong. Previous studies have attempted to replicate Asch's experiment with a group of robots but have failed to observe conformity. One explanation can be made using Hodges and Geyers work, in which they propose that people consider distinct criteria (truth, trust, and social solidarity) when deciding whether to conform to others. In order to study how trust and truth affect conformity, we propose an experiment in which participants play a game with three robots, in which there are no objective answers. We measured how many times participants changed their preliminary answers to match the group of robots' in their final answer. We conducted a between-subjects study (N = 30) in which there were two conditions: one in which participants saw the group of robots' preliminary answer before deciding their final answer, and a control condition in which they did not know the robots' preliminary answer. Participants in the experimental condition conformed significantly more (29%) than participants in the control condition (6%). Therefore we have shown that groups of robots can cause people to conform to them. Additionally trust plays a role in conformity:

initially, participants conformed to robots at a similar rate to Asch's participants, however, many participants stop conforming later in the game when trust is lost due to the robots choosing an incorrect answer.

Robots and Racism (11:30~11:50)

Christoph Bartneck (University of Canterbury), **Kumar Yogeeswaran** (University of Canterbury), **Qi Min Ser** (University of Canterbury), **Graeme Woodward** (University of Canterbury), **Robert Sparrow** (Monash University), **Siheng Wang** (Guizhou University of Engineering Science), **Friederike Eyssel** (University of Bielefeld)

Most robots currently being sold or developed are either stylized with white material or have a metallic appearance. In this research we used the shooter bias paradigm and several questionnaires to investigate if people automatically identify robots as being racialized, such that we might say that some robots are 'White' while others are 'Asian', or 'Black'. To do so, we conducted an extended replication of the classic social psychological shooter bias paradigm using robot stimuli to explore whether effects known from human-human intergroup experiments would generalize to robots that were racialized as Black and White. Reaction-time based measures revealed that participants demonstrated 'shooter-bias' toward both Black people and robot racialized as Black. Participants were also willing to attribute a race to the robots depending on their racialization and demonstrated a high degree of inter-subject agreement when it came to these attributions.

Mindless Robots get Bullied (11:50~12:10)

Merel Keijsers (University of Canterbury), **Christoph Bartneck** (University of Canterbury)

Humans recognise and respond to robots as social agents, to such extent that they occasionally attempt to bully a robot. The current paper investigates whether aggressive behaviour directed towards robots is influenced by the same social processes that guide human bullying behaviour. More specifically, it measured the effects of dehumanisation primes and anthropomorphic qualities of the robot on participants' verbal abuse of a virtual robotic agents.
Contrary to previous findings in human-human interaction, priming participants with power did not result in less mind attribution. However, evidence for dehumanisation was still found, as the less mind participants attributed to the robot, the more aggressive responses they gave. In the main study this effect was moderated by the manipulations of power and robot anthropomorphism; the low anthropomorphic robot in the power prime condition endured significantly less abuse, and mind attribution remained a significant predictor for verbal aggression in all conditions save the low anthropomorphic robot with no prime.
It is concluded that dehumanisation occurs in human-robot interaction and that like in human-human interaction, it is linked to aggressive behaviour. Moreover, it is argued that this dehumanisation is different from anthropomorphism as well as human-human dehumanisation, since anthropomorphism itself did not predict aggressive behaviour and dehumanisation of robots was not influenced by primes that have been established in human-human dehumanisation research.

Full Papers Session 4, Track 1
Room W192 (13:40~15:00)

Designing Robots and Interactions

Session Chair: Maya Cakmak
(University of Washington)

User-Centered Robot Head Design: a Sensing Computing Interaction Platform for Research Robotics (SCIPRR) (13:40~14:00)

Anthony Harrison (Naval Research Laboratory), **Wendy Xu** (Oregon State University), **J. Gregory Trafton** (Naval Research Laboratory)

We developed and evaluated a novel humanoid

head, SCIPRR (Sensing, Computing, Interacting Platform for Robotics Research). SCIPRR is a head shell that was iteratively created with additive manufacturing. SCIPRR contains internal scaffolding that allows sensors, small form computers, and a back-projection system to display an animated face on a front-facing screen. SCIPRR was developed using User Centered Design principles and evaluated using three different methods. First, we created multiple, small-scale prototypes through additive manufacturing and performed polling and refinement of the overall head shape. Second, we performed usability evaluations of expert HRI mechanics as they swapped sensors and computers within the the SCIPRR head. Finally, we ran and analyzed an experiment to evaluate how much novices would like a robot with our head design to perform different social and traditional robot tasks. We made both major and minor changes after each evaluation and iteration. Overall, expert users liked the SCIPRR head and novices wanted a robot with the SCIPRR head to perform more tasks (including social tasks) than a more traditional robot.

Bioluminescence-Inspired Human-Robot Interaction: Designing Expressive Lights that Affect Human's Willingness to Interact with a Robot (14:00~14:20)

Sichao Song (The Graduate University for Advanced Studies, SOKENDAI), **Seiji Yamada** (National Institute of Informatics & The Graduate University for Advanced Studies, SOKENDAI)

Bioluminescence is the production and emission of light by a living organism. It, as a means of communication, is of importance for the survival of various creatures. Inspired by bioluminescent light behaviors, we explore the design of expressive lights and evaluate the effect of such expressions on a human's perception of and attitude toward an appearance-constrained robot. Such robots are in urgent need of finding effective ways to present themselves and communicate their intentions due to a lack of social expressivity. We particularly focus on the expression of attractiveness and hostility because a robot would need to be able to attract or keep away human users in practical human-robot interaction (HRI) scenarios. In this work, we installed an LED lighting system on a Roomba robot and conducted a series of two experiments. We first worked through a structured approach to determine the best light expression designs for the robot to show attractiveness and hostility. This resulted in four recommended light expressions. Further, we performed a verification study to examine the effectiveness of such light expressions in a typical HRI context. On the basis of the findings, we offer design guidelines for expressive lights that HRI researchers and practitioners could readily employ.

"Haru": Hardware Design of an Experimental Table top Robot Assistant (14:20~14:40)

Randy Gomez (Honda Research Institute Japan), **Deborah Szapiro** (University of Technology Sydney), **Kerl Galindo** (University of Technology Sydney), **Keisuke Nakamura** (Honda Research Institute Japan)

This paper discusses the design and development of an experimental tabletop robot called 'Haru' based on design thinking methodology. Right from the very beginning of the design process, we have brought an interdisciplinary team that includes animators, performers and sketch artists to help create the first iteration of a distinctive anthropomorphic robot design based on a concept that leverages form factor with functionality. Its unassuming physical affordance is intended to keep human expectation grounded while its actual interactive potential stokes human interest. The meticulous combination of both subtle and pronounced mechanical movements together with its stunning visual displays, highlight its affective affordance. As a result, we have developed the first iteration of our tabletop robot rich in affective potential for use in different research fields involving long-term human-robot interaction.

Iterative Design of an Upper Limb Rehabilitation Game with Tangible Robots (14:40~15:00)

Arzu Guneysu Ozgur (ÉPFL), **Maximilian Wessel** (ÉPFL), **Wafa Johal** (ÉPFL), **Kshitij Sharma** (NTNU), **Ayberk Özgür** (ÉPFL), **Philippe Vuadens** (Clinique Romande de Réadaptation), **Francesco Mondada** (ÉPFL), **Friedhelm Hummel** (ÉPFL & Clinique Romande de Réadaptation), **Pierre Dillenbourg** (ÉPFL)

Rehabilitation aims to ameliorate deficits in motor control via intensive practice with the affected limb. Current strategies, such as one-on-one therapy done in rehabilitation centers, have limitations such as treatment frequency and intensity, cost and requirement of mobility. Thus, a promising strategy is home-based therapy that includes task specific exercises. However, traditional rehabilitation tasks may frustrate the patient due to their repetitive nature and may result in lack of motivation and poor rehabilitation. In this article, we propose the design and verification of an effective upper extremity rehabilitation game with a tangible robotic platform named Cellulo as a novel solution to these issues. We first describe the process of determining the design rationales to tune speed, accuracy and challenge. Then we detail our iterative participatory design process and test sessions conducted with the help of stroke, brachial plexus and cerebral palsy patients (18 in total) and 7 therapists in 4 different therapy centers. We present the initial quantitative results, which support several aspects of our design rationales and conclude with our future study plans.

Full Papers Session 4, Track 2

Room W194 (13:40~15:00)

Groups and Teams

Session Chair: Adriana Tapus
(ENSTA, ParisTech)

Human-Robot Similarity and Willingness to Work with Robotic Co-Worker (13:40~14:00)

Sangseok You (Syracuse University), **Lionel Robert Jr.** (University of Michigan)

Organizations now face a new challenge of encouraging their employees to work alongside robots. In this paper, we address this problem by investigating the impacts of human-robot similarity, trust in a robot, and the risk of physical danger on individuals' willingness to work with a robot and their willingness to work with a robot over a human co-worker. We report the results from an online experimental study involving 200 participants. Results showed that human-robot similarity promoted trust in a robot, which led to willingness to work with robots and ultimately willingness to work with a robot over a human co-worker. However, the risk of danger moderated not only the positive link between the surface-level similarity and trust in a robot, but also the link between intention to work with the robot and willingness to work with a robot over a human co-worker. We discuss several implications for the theory of human-robot interaction and design of robots.

Group-based Emotions in Teams of Humans and Robots (14:00~14:20)

Filipa Correia (INESC-ID & Instituto Superior Técnico, Universidade de Lisboa), **Samuel Mascarenhas** (INESC-ID & Instituto Superior Técnico, Universidade de Lisboa), **Rui Prada** (INESC-ID & Instituto Superior Técnico, Universidade de Lisboa), **Francisco Melo** (INESC-ID & Instituto Superior Técnico, Universidade de Lisboa), **Ana Paiva** (INESC-ID & Instituto Superior Técnico, Universidade de Lisboa)

Providing social robots an internal model of emotions can help them guide their behaviour in a more humane manner by simulating the ability to feel empathy towards others. Furthermore, the growing interest in creating robots that are capable of collaborating with other humans in team settings provides an opportunity to explore another side of human emotion, namely, group-based emotions.

This paper contributes with the first model on group-based emotions in social robotic partners. We defined a model of group-based emotions for social robots that allowed us to create two distinct robotic characters that express either individual or group-based emotions. This paper also contributes with a user study where two autonomous robots embedded the previous characters, and formed two human-robot teams to play a competitive game. Our results showed that participants perceived the robot that expresses group-based emotions as more likeable and attributed higher levels of group identification and group trust towards their teams, when compared to the robotic partner that expresses individual-based emotions.

Where Should Robots Talk?: Spatial Arrangement Study from a Participant Workload Perspective (14:20~14:40)

Takahiro Matsumoto (Keio University), **Mitsuhiro Goto** (NTT Service Evolution Laboratories), **Ryo Ishii** (NTT Media Intelligence Laboratories), **Tomoki Watanabe** (NTT Service Evolution Laboratories), **Tomohiro Yamada** (NTT Service Evolution Laboratories), **Michita Imai** (Keio University)

Several benefits obtained using multiple robots in conversation have been reported in the human-robot interaction field.This paper first presents pre-trial results by which elderly people assigned a lower rating to a conversation with two robots than to one with a single robot. Observations of the trial suggest the hypothesis that an inappropriate spatial arrangement between robots and humans increases the workload in a conversation. Reducing the workload is important, especially when robots are used by elderly people. Therefore, we specifically examine the workload that is influenced by the spatial arrangement in group conversation. To verify the hypothesis, we use a NASA-TLX and a dual-task method to evaluate the workload and to conduct a comparative experiment in which the participant talks with two robots in two spatial arrangements. We also conduct a case study for elderly people in the same conversational conditions. From these experiments, we demonstrate that the spatial arrangement in which people cannot see both robots simultaneously increases their conversational workload and decreases their evaluation of the dialogue compared to a spatial arrangement by which people can see both robots simultaneously. We also show that the primary cause of the workload by positioning is not physical but mental.

Friends or Foes? Socioemotional Support and Gaze Behaviors in Mixed Groups of Humans and Robots (14:40~15:00)

Raquel Oliveira (Instituto Universitário de Lisboa ISCTE-IUL, CIS-IUL), **Patrícia Arriaga** (Instituto Universitário de Lisboa ISCTE-IUL, CIS-IUL), **Patrícia Alves-Oliveira** (Instituto Universitário de Lisboa ISCTE-IUL, CIS-IUL & INESC-ID), **Filipa Correia** (Instituto Superior Técnico, Universidade de Lisboa & INESC-ID), **Sofia Petisca** (Instituto Universitário de Lisboa ISCTE-IUL, CIS-IUL & INESC-ID), **Ana Paiva** (Instituto Superior Técnico, Universidade de Lisboa & INESC-ID)

This study investigated non-verbal behavior and socioemotional interactions in small-groups of humans and robots. Sixty-participants were involved in a group setting in which they were required to play a card game with another human and two robots (playing as partners or as opponents). The two robots displayed different goal orientations: a competitive robot (named Emys-) and a relationship-driven cooperative robot (named Glin+). Video recordings of the interactions were analyzed in three game play sessions. Eye gaze and socioemotional support behaviors were coded based on Bales' Interaction Process Analysis. Results indicated that gaze behavior towards partners was more frequently displayed to the relationship-driven robot than to the competitive robot and the human partners. In contrast, gaze towards opponents occurred more often towards the competitive robot than to the relationship-driven robot and the human opponents. Socioemotional support occurred more frequently towards partners than opponents, and was also displayed more often towards humans than towards robots. Moreover, in the sessions where the robots were opponents, participants provided more support to the competitive robot. This investigation in small groups of humans and

robots provided evidence of different interaction patterns towards robots displaying distinct orientation goals, which can be useful in guiding the successful design of social robots.

Full Papers Session 5

Room W192 (15:30~17:10)

Best Paper Nominees

Five 20-Minute Presentations
(5 full papers)

Social Robots for Engagement in Rehabilitative Therapies: Design Implications from a Study with Therapists (15:30~15:50)

Katie Winkle (Bristol Robotics Laboratory), **Praminda Caleb-Solly** (Bristol Robotics Laboratory), **Ailie Turton** (Bristol Robotics Laboratory), **Paul Bremner** (Bristol Robotics Laboratory)

In this paper we present the results of a qualitative study with therapists to inform social robotics and human robot interaction (HRI) for engagement in rehabilitative therapies. Our results add to growing evidence that socially assistive robots (SARs) could play a role in addressing patients' low engagement with self-directed exercise programmes. Specifically, we propose how SARs might augment or offer more pro-active assistance over existing technologies such as smartphone applications, computer software and fitness trackers also designed to tackle this issue. In addition, we present a series of design implications for such SARs based on therapists' expert knowledge and best practices extracted from our results. This includes an initial set of SAR requirements and key considerations concerning personalised and adaptive interaction strategies.

"Thank You for Sharing that Interesting Fact!": Effects of Capability and Context on Indirect Speech Act Use in Task-Based Human-Robot Dialogue (15:50~16:10)

Tom Williams (Colorado School of Mines), **Daria Thames** (Tufts University), **Julia Novakoff** (Tufts University), **Matthias Scheutz** (Tufts University)

Naturally interacting robots must be able to understand natural human speech. As such, recent work has sought to allow robots to infer the intentions behind commonly used non-literal utterances such as indirect speech acts (ISAs). However, it is still unclear to what extent ISAs will actually be used in task-based human-robot dialogue, and to what extent robots could function without the ability to understand ISAs. In this paper, we present the results of a Wizard-of-Oz experiment that examined human ISA use in scenarios that did or did not have conventionalized social norms, and analyzed both ISA use and perceptions of robots when robots were or were not capable of understanding ISAs. Our results suggest that (1) ISAs are commonly used in task-based human-robot dialogues, even when robots show themselves unable to understand ISAs; (2) ISA use is more common in contexts with conventionalized social norms; and (3) a robot's inability to understand ISAs harms both the robot's task performance and human perception of the robot.tegies.

Planning with Trust for Human-Robot Collaboration (16:10~16:30)

Min Chen (National University of Singapore), **Stefanos Nikolaidis** (Carnegie Mellon University), **Harold Soh** (National University of Singapore), **David Hsu** (National University of Singapore), **Siddhartha Srinivasa** (University of Washington)

Trust is essential for human-robot collaboration and user adoption of autonomous systems, such as robot assistants. This paper introduces a computational model which integrates trust into robot decision-making. Specifically, we learn from data a partially observable Markov decision process (POMDP) with human trust as a latent

variable. The trust-POMDP model provides a principled approach for the robot to (i) infer the trust of a human teammate through interaction, (ii) reason about the effect of its own actions on human behaviors, and (iii) choose actions that maximize team performance over the long term. We validated the model through human subject experiments on a table-clearing task in simulation (201 participants) and with a real robot (20 participants). The results show that the trust-POMDP improves human-robot team performance in this task. They further suggest that maximizing trust in itself may not improve team performance.

Communicating Robot Motion Intent with Augmented Reality (16:30~16:50)

Michael Walker (University of Colorado Boulder), **Hooman Hedayati** (University of Colorado Boulder), **Jennifer Lee** (University of Colorado Boulder), **Daniel Szafir** (University of Colorado Boulder)

Humans coordinate teamwork by conveying intent through social cues, such as gestures and gaze behaviors. However, these methods may not be possible for appearance-constrained robots that lack anthropomorphic or zoomorphic features, such as aerial robots. We explore a new design space for communicating robot motion intent by investigating how augmented reality (AR) might mediate human-robot interactions. We develop a series of explicit and implicit designs for visually signaling robot motion intent using AR, which we evaluate in a user study. We found that several of our AR designs significantly improved objective task efficiency over a baseline in which users only received physically-embodied orientation cues. In addition, our designs offer several trade-offs in terms of intent clarity and user perceptions of the robot as a teammate.

An Autonomous Dynamic Camera Method for Effective Remote Teleoperation (16:50~17:10)

Daniel Rakita (University of Wisconsin-Madison), **Bilge Mutlu** (University of Wisconsin-Madison), **Michael Gleicher** (University of Wisconsin-Madison)

In this paper, we present a method that improves the ability of remote users to teleoperate a \textit{manipulation} robot arm by continuously providing them with an effective viewpoint using a second *camera-in-hand* robot arm. The user controls the manipulation robot using *any* teleoperation interface, and the camera-in-hand robot automatically servos to provide a view of the remote environment that is estimated to best support effective manipulations. Our method avoids occlusions with the manipulation arm to improve visibility, provides context and detailed views of the environment by varying the camera-target distance, utilizes motion prediction to cover the space of the user's next manipulation actions, and actively corrects views to avoid disorienting the user as the camera moves. Through two user studies, we show that our method improves teleoperation performance over alternative methods of providing visual support for teleoperation. We discuss the implications of our findings for real-world teleoperation and for future research.

Full Papers Session

Thu, March 8

Full Papers Session 6, Track 1

Room W192 (10:30~12:10)

Communicating with and Without Speech

Session Chair: Kerstin Fischer (University of Southern Denmark)

Multimodal Expression of Artificial Emotion in Social Robots Using Color, Motion and Sound (10:30~10:50)

Diana Löffler (University of Würzburg), **Nina Schmidt** (University of Würzburg), **Robert Tscharn** (University of Würzburg)

Artificial emotion display is a key feature of social robots to communicate internal states and behaviors in familiar human terms. While humanoid robots can draw on signals such as facial expressions or voice, emotions in appearance-constrained robots can only be conveyed through less-anthropomorphic output channels. While previous work focused on identifying specific expressional designs to convey a particular emotion, little work has been done to quantify the information content of different modalities and how they become effective in combination. Based on emotion metaphors that capture mental models of emotions, we systematically designed and validated a set of 28 different uni- and multimodal expressions for the basic emotions joy, sadness, fear and anger using the most common output modalities color, motion and sound. Classification accuracy and users' confidence of emotion assignment were evaluated in an empirical study with 33 participants and a robot probe. The findings are distilled into a set of recommendations about which modalities are most effective in communicating basic artificial emotion. Combining color with planar motion offered the overall best cost/benefit ratio by making use of redundant multimodal coding. Furthermore, modalities differed in their degree of effectiveness to communicate single emotions. Joy was best conveyed via color and motion, sadness via sound, fear via motion and anger via color. robotics.

Getting to Know Each Other: The Role of Social Dialogue in Recovery from Errors in Social Robots (10:50~11:10)

Gale Lucas (University of Southern California), **Jill Boberg** (University of Southern California), **David Traum** (University of Southern California), **Ron Artstein** (University of Southern California), **Jonathan Gratch** (University of Southern California), **Alesia Gainer** (University of Southern California), **Emmanuel Johnson** (University of Southern California), **Anton Leuski** (University of Southern California), **Mikio Nakano** (Honda Research Institute Japan)

This work explores the extent to which social dialogue can mitigate (or exacerbate) the loss of trust caused when robots make conversational errors. Our study uses a NAO robot programmed to persuade users to agree with its rankings on two tasks. We perform two manipulations: (1) The timing of conversational errors—the robot exhibited errors either in the first task, the second task, or neither; (2) The presence of social dialogue—between the two tasks, users either engaged in a social dialogue with the robot or completed a control task. We found that the timing of the errors matters: replicating previous research, conversational errors reduce the robot's influence in the second task, but not on the first task. Social dialogue interacts with the timing of errors, acting as an intensifier: social dialogue helps the robot recover from prior errors, and actually boosts subsequent influence; but social dialogue backfires if it is followed by errors, because it extends the period of good performance, creating a stronger contrast effect with the subsequent errors. The design of social robots should therefore be more careful to avoid errors after periods of good performance than early on in a dialogue.

alt.HRI - Agreeing to Interact: Understanding Interaction as Human-Robot Goal Conflicts (11:10~11:30)

Kazuhiro Sasabuchi (University of Tokyo), **Katsushi Ikeuchi** (Microsoft), **Masayuki Inaba** (University of Tokyo)

In human-robot interaction (HRI) people have studied user preferable robot actions in various social situations. The role of the robot is often designed, and situations are assumed that the robot will interact with the human. However, there are also situations where either the human or robot may not be willing to interact. In such situations, the human and robot are under a goal conflict and must first agree to begin an interaction. In this paper, we re-explore interaction beginnings and endings as a confliction and agreement between human and robot goals—the willingness of whether to interact or not. Through our discussion, we categorize conflict/agreement interactions into nine situations. Using a probabilistic analysis approach and 93 HRI recordings, we evaluate the different human behaviors in different interaction situations. We further question whether learning from typical existing HRI would benefit other scenarios when a robot has physical task capabilities.We conclude that the benefits of understanding different agreement situations would largely depend on a robot's task capability as well as the human's expectation toward these capabilities; however, conflict and agreement should not be neglected when applying interaction capability to physical-task-capable robots. Our research also suggests the probabilistic drawbacks of robot speech in situations where both the human and robot are not willing to interact.

Observing Robot Touch in Context: How Does Touch and Attitude Affect Perceptions of a Robot's Social Qualities? (11:30~11:50)

Thomas Arnold (Tufts University), **Matthias Scheutz** (Tufts University)

The complex role of touch is an increasingly appreciated horizon for HRI research. The explicit and implicit registers of touch, both human-to-robot and robot-to-human, have opened up pressing questions in design and HRI ethics about embodiment, communication, care, and human affection. In this paper we present results of an MTurk survey about robot-initiated touch in a social context. We examine how a positive or negative attitude from the robot, a well as whether the robot touches an interactant, affects how a robot is judged as a worker and teammate. Our findings confirm previous empirical support for the idea of touch as enhancing social appraisals of a robot, though the extent of that positive tactile role was complicated and tempered by the survey responses' gender effects.

Social Momentum: A Framework for Legible Navigation in Dynamic Multi-Agent Environments (11:50~12:10)

Christoforos Mavrogiannis (Cornell University), **Wil Thomason** (Cornell University), **Ross Knepper** (Cornell University)

Intent-expressive robot motion has been shown to result in increased efficiency and reduced planning efforts for copresent humans. Existing frameworks for generating intent-expressive robot behaviors have typically focused on applications in static or structured environments. Under such settings, emphasis is placed towards communicating the robot's intended final configuration to other agents. However, in dynamic, unstructured and multi-agent domains, such as pedestrian environments, knowledge of the robot's final configuration is not sufficiently informative as it completely ignores the complex dynamics of interaction among agents. To address this problem, we design a planning framework that aims at generating motion that clearly communicates an agent's intended collision avoidance strategy rather than its destination. Our framework estimates the most likely intended avoidance protocols of others based on their past behaviors, superimposes them, and generates an expressive and socially compliant robot action that reinforces the expectations of others regarding these avoidance protocols. This action facilitates inference and decision making for everyone, as illustrated in the simplified topological pattern of agents' trajectories. Extensive simulations demonstrate that our framework consistently achieves significantly lower topological complexity, compared against common benchmark approaches in multi-agent collision avoidance. The significance of this result for real world applications is demonstrated by a user study that reveals statistical evidence suggesting that multi-agent trajectories of lower topological complexity tend to facilitate inference for observers.

Full Papers Session 6, Track 2

Room W194 (10:30~12:10)

Psychology and HRI

Session Chair: Christoph Bartneck
(University of Canterbury)

alt.HRI - Social Psychology and Human-Robot Interaction: An Uneasy Marriage (10:30~10:50)

Bahar Irfan (University of Plymouth), **James Kennedy** (University of Plymouth), **Séverin Lemaignan** (University of Plymouth), **Fotios Papadopoulos** (University of Plymouth), **Emmanuel Senft** (University of Plymouth), **Tony Belpaeme** (University of Plymouth)

The field of Human-Robot Interaction (HRI) lies at the intersection of several disciplines, and is rightfully perceived as a prime interface between engineering and the social sciences. In particular, our field entertains close ties with social and cognitive psychology, and there are many HRI studies which build upon commonly accepted results from psychology to explore the novel relation between humans and machines. Key to this endeavour is the trust we, as a field, put in the methodologies and results from psychology, and it is exactly this trust that is now being questioned across psychology and, by extension, should be questioned in HRI.
The starting point of this paper are a number of failed attempts by the authors to replicate old and established results on social facilitation, which leads us to discuss our arguable over-reliance and over-acceptance of methods and results from psychology. We highlight the recent 'replication crisis' in psychology, which directly impacts the HRI community and argue that our field should not shy away from developing its own reference tasks.

The Peculiarities of Robot Embodiment (EmCorp-Scale): Development, Validation and Initial Test of the Embodiment and Corporeality of Artificial Agents Scale (10:50~11:10)

Laura Hoffmann (Bielefeld University), **Nikolai Bock** (University of Duisburg-Essen), **Astrid Rosenthal v.d. Pütten** (RWTH Aachen University)

We propose a new theoretical framework assuming that embodiment effects in HAI and HRI are mediated by users' perceptions of an artificial entity's body-related capabilities. To enable the application of our framework to foster more theoretical-driven research, we developed a new self-report measurement that assesses bodilyrelated perceptions of the embodiment and corporeality—which we reveal as not being a binary characteristic of artificial entities. For the development and validation of the new scale we conducted two surveys and one video-based experiment. Exploratory factor analysis reveal a four-factorial solution with good reliability (Study 2, n=442), which was confirmed via confirmatory factor analysis (Study 3, n=260). In addition, we present first insights into the explanatory power of the scale: We reveal that humans' perceptions of an artificial entity's capabilities vary between virtual and physical embodiments, and that the evaluation of the artificial counterpart can be explained through the perceived capabilities. Practical applications and future research lines are discussed.

Be More Transparent and Users Will Like You: A Robot Privacy and User Experience Design Experiment (11:10~11:30)

Jonathan Vitale (University of Technology Sydney), **Meg Tonkin** (University of Technology Sydney), **Sarita Herse** (University of Technology Sydney), **Suman Ojha** (University of Technology Sydney), **Jesse Clark** (University of Technology Sydney), **Mary-Anne Williams** (University of Technology

Sydney), **Xun Wang** (Commonwealth Bank of Australia), **William Judge** (Commonwealth Bank of Australia)

Robots interacting with humans in public spaces often need to collect users' private information in order to provide the required services. Current privacy legislation in major jurisdictions requires organisations to disclose information about their data collection process and obtain user's consent prior to collecting privacy sensitive information. In this study, we consider a privacy-sensitive design of a data collection system for face identification. We deployed a face enrolment system on a humanoid robot with human-like gesturing and speech. We compared it with an equivalent system, in terms of capability and interactive process, on a screen-based interactive kiosk.
In our previous contribution, we investigated the effects that embodiment has on users' privacy considerations. We found that an embodied humanoid robot is capable of collecting more private information from users in comparison to a disembodied interactive kiosk. However, this effect was statistically significant only when the two compared systems were using a transparent interface, i.e. an interface communicating to users the privacy policies for data processing and storage. Thus, in this work, we aim to further investigate the effects of transparency on users' privacy considerations and their experience with robot applications.
We found that when comparing a non-transparent vs. transparent interface within the same system (i.e. on an embodied robot or on a disembodied kiosk) transparency does not lead to significant effects on users' privacy considerations. However, we found that transparency leads to a significantly better user experience for both systems.
Therefore, our overall analyses suggest that both the interactive robot and the interactive kiosk are capable of enhancing the user experience by providing transparent information to users, which is required by privacy legislation. However, an interactive kiosk providing transparent information elicits significantly more privacy concerns in users as compared to the robot supplying the very same transparent information.
This exploratory study provides conclusions that provide valuable insights for designing robot applications dealing with users privacy and it discusses the related legal implications, concluding with recommendations for privacy policymakers.

alt.HRI - Design Strategies for Representing the Divine in Robots (11:30~11:50)

Gabriele Trovato (Waseda University), **Cesar Lucho** (Pontificia Universidad Católica del Perú), **Alexander Huerta-Mercado** (Pontificia Universidad Católica del Perú), **Francisco Cuellar** (Pontificia Universidad Católica del Perú)

Robot appearance morphology can be divided in anthropomorphic, zoomorphic and functional. In previous recent work, a new category was introduced, called 'theomorphic robots', in which robots carry the shape and the identity of a supernatural creature or object within a religion. This approach can bring some advantages for certain categories of users such as children and elders. This paper is an exploratory discussion over practical design strategies for representing the divine in robots, based on theoretical insights on the historical intertwinements between sacred art and robotics. The illustrated concepts will be followed in the realisation of the prototypes of the first theomorphic robots.

Futuristic Autobiographies: Weaving Participant Narratives to Elicit Values around Robots (11:50~12:10)

EunJeong Cheon (Indiana University), **Norman Su** (Indiana University)

In this paper, we motivate and introduce Futuristic Autobiographies, a method inspired by design fiction for eliciting values and perspectives on the future of technologies from participants such as users, designers, and researchers. Futuristic autobiographies are the creative work of the researchers and participants. Grounded in empirical and background work, researchers pose several stories involving the participant as a character about a future state with robots. Participants are then asked to weave fictional autobiographies to explain what led to this future state. Via a case study in which futuristic autobiographies were used with 23 roboticists, we detail the process involved in developing and implementing this method. When futuristic autobiographies are employed and carefully crafted from background research, they allow informants to speak for themselves on how their practices and values are intertwined now and in the future. We highlight both the benefits and challenges of futuristic autobiographies as a way to elicit rich stories about values. We argue that futuristic autobiographies are a promising addition to the current qualitative methods toolkit used in HRI.

Full Papers Session 7, Track 1
Room W192 (13:40~15:00)

Rethinking Human-Robot Relationships

Session Chair: Alan Wagner
(Penn State University)

alt.HRI - Interacting with Anatomically Complete Robots - A Discussion about Human-robot Relationships (13:40~14:00)

Christoph Bartneck (University of Canterbury), **Matthew McMullen** (Abyss Creations LLC)

Abyss Creations LLC is selling highly realistic and anatomically complete dolls for decades. The interaction that customers have with the dolls has been subject to many controversial discussions. David Levy already showed that humanity has been rather inventive when it comes to machines that useful for sexual purposes. Abyss Creations recently revealed its Harmony platform that consists of a robotic head that is attached to their doll bodies. This paper present an interview with Matthew McMullen, CEO and Creative director of Abyss Creations about the Harmony platform and its implications on human-robot relationships.

Social Cobots: Anticipatory Decision-Making for Collaborative Robots Incorporating Unexpected Human Behaviors (14:00~14:20)

Orhan Can Görür (Technische Universität Berlin), **Benjamin Rosman** (CSIR & University of Witwatersrand), **Fikret Sivrikaya** (GT-ARC gemeinnützige GmbH), **Sahin Albayrak** (Technische Universität Berlin)

We propose an architecture as a robot's decision-making mechanism to anticipate a human's state of mind, and so plan accordingly during a human-robot collaboration task. At the core of the architecture lies a novel stochastic decision-making mechanism that implements a partially observable Markov decision process anticipating a human's state of mind in two-stages. In the first stage it anticipates the human's task related availability, intent (motivation), and capability during the collaboration. In the second, it further reasons about these states to anticipate the human's true need for help. Our contribution lies in the ability of our model to handle these unexpected conditions: 1) when the human's intention is estimated to be irrelevant to the assigned task and may be unknown to the robot, e.g., motivation is lost, another assignment is received, onset of tiredness, and 2) when the human's intention is relevant but the human doesn't want the robot's assistance in the given context, e.g., because of the human's changing emotional states or the human's task-relevant distrust for the robot. Our results show that integrating this model into a robot's decision-making process increases the efficiency and naturalness of the collaboration.

alt.HRI - Crucial Answers about Humanoid Capital (14:20~14:40)

Brian Beaton (California Polytechnic State University)

Inside AI research and engineering communities, explainable artificial intelligence (XAI) is one of the most provocative and promising lines of AI research and development today. XAI has the potential to make expressible the context and domain-specific benefits of particular AI applications to a diverse and inclusive array of stakeholders and audiences. In addition, XAI has the potential to make AI benefit claims more deeply evidenced. Outside AI research and engineering communities, one of the most provocative and promising lines of research happening today is the work on 'humanoid capital' at the edges of the social, behavioral, and economic sciences. Humanoid capital theorists renovate older discussions of 'human capital' as part of trying to make calculable and provable the domain-specific capital value, value-adding potential, or relative worth (i.e., advantages and benefits) of different humanoid models over time. Bringing these two exciting streams of research into direct conversation for the first time is the larger goal of this landmark paper. The primary research contribution of the paper is to detail some of the key requirements for making humanoid robots explainable in capital terms using XAI approaches. In this regard, the paper not only brings two streams of provocative research into much-needed conversation but also advances both streams.

A Design Methodology for the UX of HRI: Example Case Study of a Social Robot at an Airport (14:40~15:00)

Meg Tonkin (University of Technology Sydney), **Jonathan Vitale** (University of Technology Sydney), **Sarita Herse** (University of Technology Sydney), **Mary-Anne Williams** (University of Technology Sydney), **William Judge** (Commonwealth Bank of Australia), **Xun Wang** (Commonwealth Bank of Australia)

Research in robotics and human-robot interaction is becoming more and more mature. Additionally, more affordable social robots are being released commercially. Thus, industry is currently demanding ideas for viable commercial applications to situate social robots in public spaces and enhance customers experience. However, present literature in human-robot interaction does not provide a clear set of guidelines and a methodology to (i) identify commercial applications for robotic platforms able to position the users' needs at the centre of the discussion and (ii) ensure the creation of a positive user experience. With this paper we propose to fill this gap by providing a methodology for the design of robotic applications including these desired features, suitable for integration by researchers, industry, business and government organisations. As we will show in this paper, we successfully employed this methodology for an exploratory field study involving the trial implementation of a commercially available, social humanoid robot at an airport.

Full Papers Session 7, Track 2

Room W194 (13:40~15:00)

Coordination in Time and Space

Session Chair: Greg Trafton
(Naval Research Lab)

Compact Real-time Avoidance on a Humanoid Robot for Human-robot Interaction (13:40~14:00)

Dong Hai Phuong Nguyen (Istituto Italiano di Tecnologia), **Matej Hoffmann** (Czech Technical University Prague & Istituto Italiano di Tecnologia), **Alessandro Roncone** (Yale University), **Ugo Pattacini** (Istituto Italiano di Tecnologia), **Giorgio Metta** (Istituto Italiano di Tecnologia)

With robots leaving factories and entering less controlled domains, possibly sharing the space with humans, safety is paramount and multimodal awareness of the body surface and the surrounding environment is fundamental. Taking inspiration from peripersonal space representations in humans, we present a framework on a humanoid robot that dynamically maintains such a protective safety zone, composed of the following main components: (i) a human 2D keypoints estimation pipeline employing a deep learning based algorithm, extended here into 3D using disparity; (ii) a distributed peripersonal space representation around the robot's body parts; (iii) a reaching controller that incorporates all obstacles entering the robot's safety zone on the fly into the task. Pilot experiments demonstrate that an effective safety margin between the robot's and the human's body parts is kept. The proposed

solution is flexible and versatile since the safety zone around individual robot and human body parts can be selectively modulated—here we demonstrate stronger avoidance of the human head compared to rest of the body. Our system works in real time and is self-contained, with no external sensory equipment and use of onboard cameras only.

Detecting Contingency for HRI in Open-World Environments (14:00~14:20)

Elaine Short (University of Texas Austin), **Mai Lee Chang** (University of Texas Austin), **Andrea Thomaz** (University of Texas Austin)

This paper presents a novel algorithm for detecting contingent reactions to robot behavior in noisy real-world environments with naive users. Prior work has established that one way to detect contingency is by calculating a difference metric between sensor data before and after a robot probe of the environment. Our algorithm, CIRCLE (Contingency for Interactive Real-time CLassification of Engagement) provides a new approach to calculating this difference and detecting contingency, improving the running time for the difference calculation from 2.5 seconds to approximately 0.001 seconds on an 1100-sample vector, and effectively enabling real-time detection of contingent events. We show accuracy comparable to the best offline results for detecting contingency in this way (89.5% vs 91% in prior work), and demonstrate the utility of the real-time contingency detection in a field study of a survey-administering robot in a noisy open-world environment with naive users, showing that the robot can decrease the number of requests it makes (from 38 to 13) while more efficiently collecting survey responses (30% response rate rather than 26.3%).

Effects of Robot Sound on Auditory Localization in Human-Robot Collaboration (14:20~14:40)

Elizabeth Cha (University of Southern California), **Naomi Fitter** (University of Southern California), **Yunkyung Kim** (NASA Ames Research Center), **Terry Fong** (NASA Ames Research Center), **Maja Mataric** (University of Southern California)

Auditory cues facilitate situational awareness by enabling humans to infer what is happening in the nearby environment. Unlike humans, many robots do not continuously produce perceivable state-expressive sounds. In this work, we propose the use of iconic auditory signals that mimic the sounds produced by a robot's operations. In contrast to artificial sounds (e.g., beeps and whistles), these signals are primarily functional, providing information about the robot's actions and state. We analyze the effects of two variations of robot sound, tonal and broadband, on auditory localization during a human-robot collaboration task. Results from 24 participants show that both signals significantly improve auditory localization, but the broadband variation is preferred by participants. We then present a computational formulation for auditory signaling and apply it to the problem of auditory localization using a human-subjects data collection with 18 participants to learn optimal signaling policies.

Evaluating Social Perception of Human-to-Robot Handovers using the Robot Social Attributes Scale (RoSAS) (14:40~15:00)

Matthew Pan (University of British Columbia), **Elizabeth Croft** (Monash University), **Günter Niemeyer** (Disney Research)

This work explores social perceptions of robots within the domain of human-to-robot handovers. Using the Robotic Social Attributes Scale (RoSAS), we explore how users socially judge robot receivers as three factors are varied: initial position of the robot arm prior to handover, grasp method employed by the robot when receiving a handover object trading off perceived object safety for time efficiency, and retraction speed of the arm following handover. Our results show that over multiple handover interactions with the robot, users gradually perceive the robot receiver as being less discomforting and having more emotional warmth. Additionally, we have found that by varying grasp method and retraction speed, users may hold significantly different judgments of robot competence and discomfort. With these results, we recognize empirically that users are able to develop social perceptions of robots which can change through modification of robot receiving behaviour and through repeated interaction with the robot. More widely, this work suggests that measurement of user social perceptions should play a larger role in the design and evaluation of human-robot interactions and that the RoSAS can serve as a standardized tool in this regard.

Full Papers Session 8

Room W192 (15:30~16:30)

Transactions on HRI

Session Chair: Guy Hoffman
(Cornell University)

Closed-loop Global Motion Planning for Reactive Collision-free Execution of Learned Tasks (15:30~15:50)

Chris Bowen (University of North Carolina Chapel Hill), **Ron Alterovitz** (University of North Carolina Chapel Hill)

We present a motion planning approach for performing a learned task while reacting to the movement of obstacles and task-relevant objects. We employ a closed-loop sampling-based motion planner operating multiple times a second that senses obstacles and task-relevant objects and generates collision-free motion plans based on a learned task model. The task model is learned from expert demonstrations prior to task execution and is represented as a hidden Markov model. During task execution, our motion planner quickly searches in the Cartesian product of the task model and a probabilistic roadmap for a collision-free plan with features most similar to the demonstrations given the current locations of the task-relevant objects. We accelerate replanning using a fast bidirectional search and by biasing the sampling distribution using information from the learned task model. We show the efficacy of our approach with the Baxter robot performing two tasks.

Adapting a General Purpose Social Robot for Paediatric Rehabilitation through In-situ Design (15:50~16:10)

Felip Marti Carillo (Swinburne University of Technology), **Joanna Butchart** (Royal Children's Hospital Melbourne, Murdoch Children's Research Institute), **Sarah Knight** (Murdoch Children's Research Institute, Royal Children's Hospital Melbourne), **Adam Scheinberg** (Royal Children's Hospital Melbourne, Murdoch Children's Research Institute), **Lisa Wise** (Swinburne University of Technology), **Leon Sterling** (Swinburne University of Technology), **Chris McCarthy** (Swinburne University of Technology)

Socially Assistive Robots (SARs) offer great promise for improving outcomes in paediatric rehabilitation. However, the design of software and interactive capabilities for SARs must be carefully considered in the context of their intended clinical use. While previous work has explored specific roles and functionalities to support paediatric rehabilitation, few have considered the design of such capabilities in the context of ongoing clinical deployment. In this paper we present a two-phase In-situ design process for SARs in health care, emphasising stakeholder engagement and on-site development. We explore this in the context of developing the humanoid social robot NAO as a socially assistive rehabilitation aid for children with cerebral palsy. We present and evaluate our design process, outcomes achieved, and preliminary results from ongoing clinical testing with 9 patients and 5 therapists over 14 sessions. We argue that our in-situ Design methodology has been central to the rapid and successful deployment of our system.

Reframing Assistive Robots to Promote Successful Aging (16:10~16:30)

Hee Rin Lee, Laurel Riek (UC San Diego)

We are living in an exciting time, as people are living longer, active lives. This is reshaping how we think about aging. Rather than viewing aging as a problem to be fixed (i.e., a deficit model of aging), many aging researchers are viewing aging as a developmental stage of life to be celebrated and supported, e.g., 'Successful Aging'. In this paper, we embrace this new approach and consider it in the context of assistive robot design, in an aim to steer the conversation away from deficit models that have limited robot design possibilities. To explore an alternative design approach to the study of aging in HRI, we invited five aging researchers including geriatricians and nine older adults to our research process. In the study, participants illustrated their interpretations of aging and suggested potential assistive robots. We found that while all participants perceived disabilities due to aging as important to understand it, they considered disabilities as only one aspect of the experience of aging. Based on the findings, we suggest a social model of disability as an alternative approach to reframe assistive robots and we call for 'robots for successful aging.'