Report from Dagstuhl Seminar 19411

Social Agents for Teamwork and Group Interactions

Edited by
Elisabeth André[1], Ana Paiva[2], Julie Shah[3], and Selma Šabanovic[4]

1 **Universität Augsburg, DE,** andre@informatik.uni-augsburg.de
2 **INESC-ID – Porto Salvo, PT,** ana.paiva@inesc-id.pt
3 **MIT – Cambridge, US,** julie_a_shah@csail.mit.edu
4 **Indiana University, US,** selmas@indiana.edu

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 19411 "Social Agents for Teamwork and Group Interactions". It summarises the three talks that were held during the seminar on three different perspectives: the impact of robots in human teamwork, mechanisms to support group interactions in virtual settings, and affect analysis in human-robot group settings. It also details the considerations of six working groups covering the following topics: datasets, design, team dynamics, social cognition, scenarios, and social behaviours.

Seminar October 6–11, 2019 – http://www.dagstuhl.de/19411
2012 ACM Subject Classification Human-centered computing → Collaborative interaction, Human-centered computing → Interaction design, Human-centered computing → User studies, Computing methodologies → Cognitive robotics
Keywords and phrases Social Agents, Social Robotics, Multi-Agent Systems, Human-Agent Interaction, Groups and Teams
Digital Object Identifier 10.4230/DagRep.9.10.1
Edited in cooperation with Filipa Correia and Sarah Strohkorb Sebo

1 Executive Summary

Elisabeth André (Universität Augsburg, DE, andre@informatik.uni-augsburg.de)
Ana Paiva (INESC-ID & University of Lisbon, PT, ana.paiva@inesc-id.pt)
Julie Shah (Massachusetts Institute of Technology, US, julie_a_shah@csail.mit.edu)
Selma Šabanovic (Indiana University, US, selmas@indiana.edu)

As artificial agents and social robots become more prominent in our lives, they will also increasingly become parts of the groups and teams in which people spend much of their time. The objective of this Dagstuhl Seminar was to explore and discuss theories, methods, and techniques for building embodied social agents (including robots) that can operate in groups as members of a mixed team consisting of humans and agents. Recent advances in AI, and particularly in conversational agents, are likely to lead to an increased placement of agents in groups, covering a variety of application scenarios including healthcare, education, the workplace, and the home. Platforms such as Amazon Echo, Google Home, and new social robots such as Nao, Pepper, and Aibo facilitate such placement. Studies with robots in open-ended environments, including homes and public spaces, also suggest that people often engage with robots in such contexts in groups, rather than just individually. Yet, existing research on human-agent interaction and human-robot interaction so far focuses mostly on one-on-one interactions between a human and a social agent. To stimulate growing research in settings where one or more humans interact with multiple agents or robots, this seminar focused on human-agent communication, interaction, and teamwork in groups. As such, we discussed how agents shape the dynamics of groups, how agents and robots are able to perceive other members of a group and how they relate to each other, and how to move from one-to-one interactions to multi-party interactions of agents and humans in groups and teams. By bringing together researchers from different communities, such as human-robot interaction, multi-agent systems, social psychology, and organizational studies, we aim to generate common ground and new approaches in this interdisciplinary area. While this new domain of inquiry relies on existing research at the intersection between AI, robotics, and the social sciences, our aim is to highlight open questions that current work has not sufficiently addressed.

2 Table of Contents

3 Overview of Talks

3.1 Teamwork with Robots

Malte Jung (Cornell University, US, mfj28@cornell.edu)

Research on human-robot interaction (HRI) to date has largely focused on examining a single human interacting with a single robot. This work has led to advances in fundamental understanding about the psychology of HRI (e.g. how specific design choices affect interactions with and attitudes towards robots) and about the effective design of HRI (e,g. how novel mechanisms or computational tools can be used to improve HRI). However, the single-robot-single-human focus of this growing body of work stands in stark contrast to the complex social contexts in which robots are increasingly placed. While robots increasingly support teamwork across a wide range of settings covering search and rescue missions, minimally invasive surgeries, space exploration missions, or manufacturing, we have limited understanding of how groups people will interact with robots and how robots will affect how people interact with each other in groups and teams. In this talk I present empirical findings from several studies that show how robots can shape in direct, but also subtle ways how people interact and collaborate with each other in teams.

References

1 Malte F Jung. Coupling interactions and performance: Predicting team performance from thin slices of conflict. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 23(3):18, 2016.
2 Malte F Jung. Affective grounding in human-robot interaction. In *2017 12th ACM/IEEE International Conference on Human-Robot Interaction (HRI*, pages 263–273. IEEE, 2017.
3 Malte F Jung, Dominic DiFranzo, Brett Stoll, Solace Shen, Austin Lawrence, and Houston Claure. Robot assisted tower construction-a resource distribution task to study human-robot collaboration and interaction with groups of people. *arXiv preprint arXiv:1812.09548*, 2018.
4 Nikolas Martelaro, Malte Jung, and Pamela Hinds. Using robots to moderate team conflict: The case of repairing violations. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction Extended Abstracts*, pages 271–271. ACM, 2015.
5 Solace Shen, Petr Slovak, and Malte F Jung. Stop. i see a conflict happening.: A robot mediator for young children's interpersonal conflict resolution. In *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*, pages 69–77. ACM, 2018.
6 Hamish Tennent, Solace Shen, and Malte Jung. Micbot: A peripheral robotic object to shape conversational dynamics and team performance. In *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 133–142. IEEE, 2019.

3.2 Supporting Interactions in Online Groups

Kobi Gal (Ben-Gurion University, IL, kobig@bgu.ac.il)

Advances in network technologies and interface design are enabling group activities of varying complexities to be carried out, in whole or in part, over the internet (e.g., citizen science, Massive Online Open Courses (MOOC) and questions-and-answers sites). The need to support these highly diverse interactions brings new and significant challenges to AI; how to design efficient representations for describing online group interactions; how to provide incentives that keep participants motivated and productive; and how to provide useful, non-intrusive information to system designers to help them decide whether and how to intervene with the group's work. I describe two ongoing projects that address these challenges in the wild, and discuss the potential impact of this work to environment design.

References

1 Avi Segal, Kobi Gal, Ece Kamar, Eric Horvitz, and Grant Miller. Optimizing interventions via offline policy evaluation: Studies in citizen science. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
2 Avi Segal, Kobi Gal, Ece Kamar, Eric Horvitz, Alex Bower, and Grant Miller. Intervention Strategies for Increasing Engagement in Volunteer-Based Crowdsourcing. In *International Joint Conference on Artificial Intelligence (IJCAI)*, 2016.
3 Avi Segal, Ya'akov Kobi Gal, Robert J Simpson, Victoria Victoria Homsy, Mark Hartswood, Kevin R Page, and Marina Jirotka. Improving productivity in citizen science through controlled intervention. In *Proceedings of the 24th International Conference on World Wide Web*, pages 331–337. ACM, 2015.

3.3 Affect and Personality Analysis in Human-Human-Robot Interaction Settings

Hatice Gunes (University of Cambridge, GB, hatice.gunes@cl.cam.ac.uk)

Designing intelligent systems and interfaces with socio-emotional skills is a challenging task. Past works have mainly focussed on automatically analysing expressions, affect and personality of people in individual settings. However, when we move from single user settings to multi-user and group ones, the process of affect analysis calls for new definitions, new datasets with meaningful annotations, and appropriate feature extraction and classification mechanisms in space and time. This talk questions some of the initial assumptions made in this area, and presents an overview of the works we have undertaken in recent years in human-human and human-human-robot interaction settings.

Firstly, the talk presents a set of experiments for affect analysis of subjects in group settings. These individuals were recorded watching videos alone and watching videos as part of a group [3, 4]. Our results show that: 1) facial appearance representation (i.e., the proposed Volume Quantized Local Zernike Moments Fisher Vector) outperforms other unimodal features in affect analysis in both settings; 2) temporal learning models perform better than the static learning models; 3) it is possible to predict the context, i.e., whether a person is alone or in-a-group, using their non-verbal behavioural cues; 4) people in the same group share similarities in facial behaviours which contributes to automatic affect prediction; and 5) when the expressive behaviour of one subject in a group setting is not available, behaviours expressed by the other subject(s) can be used for affect prediction [4].

Secondly, the talk introduces a novel dataset we have collected, the Multimodal Human-Human-Robot-Interactions (MHHRI) dataset [2], acquired with the aim of studying personality simultaneously in human-human interactions (HHI) and human-human-robot interactions (HRI) and its relationship with engagement. Multimodal data was collected during a controlled interaction study where dyadic interactions between two human participants and triadic interactions between two human participants and a robot took place with interactants asking/answering a set of personal questions. Interactions were recorded using two static and two dynamic cameras as well as two biosensors, and meta-data was collected by having participants to fill in two types of questionnaires, for assessing their own personality traits and their perceived engagement with their partners (self-annotated labels) and for assessing personality traits of the other participants partaking in the study (acquaintance labels).

Thirdly, using the MHHRI dataset, the talk introduces a number of experiments we have conducted for automatic prediction of personality and engagement. We analyse interactions with the robot from the viewpoint of human participants through an ego-centric camera placed on their forehead [1]. We focus on human participants' and robot's personalities and their impact on the human-robot interactions. We automatically extract nonverbal cues (e.g., head movement) from first-person perspective and explore the relationship of nonverbal cues with participants' self-reported personality and their interaction experience. We generate two types of behaviours for the robot (i.e., extroverted vs. introverted) and examine how robot's personality and behaviour affect the findings. Significant correlations are obtained between the extroversion and agreeable-ness traits of the participants and the perceived enjoyment with the extroverted robot. Plausible relationships are also found between the measures of interaction experience and personality and the first-person vision features [1].

Finally, using the MHHRI dataset, the talk introduces work that focuses on the automatic analysis and classification of engagement based on humans' and robot's personality profiles in the triadic human-human-robot interaction setting [5]. More explicitly, the study investigates how participants' personalities can be used together with the robot's personality to predict the engagement state of each participant as well as the engagement of the overall group. The fully automatic system is first trained to predict the Big Five personality traits of each participant by extracting individual and interpersonal features from their nonverbal behavioural cues. Then the output of the personality prediction system is used as an input to the engagement classification system. Third, we focus on the concept of "group engagement", which we define as the collective engagement of the participants with the robot, and analyse the impact of similar and dissimilar personalities on the engagement classification. Our experimental results show that: 1) using the automatically predicted personality labels for engagement classification yields an F-measure on par with using the manually annotated personality labels, demonstrating the effectiveness of the automatic personality prediction module proposed; 2) using the individual and interpersonal features without utilizing personality information is not sufficient for engagement classification, instead incorporating the participants and robots personalities with individual/interpersonal features increases engagement classification performance; and 3) the best classification performance is achieved when the participants and robot are extroverted, while the worst results are obtained when all are introverted [5].

References

1 Oya Celiktutan and Hatice Gunes. Computational analysis of human-robot interactions through first-person vision: Personality and interaction experience. In *2015 24th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, pages 815–820. IEEE, 2015.

2 Oya Celiktutan, Efstratios Skordos, and Hatice Gunes. Multimodal human-human-robot interactions (mhhri) dataset for studying personality and engagement. *IEEE Transactions on Affective Computing*, 2017.

3 Wenxuan Mou, Hatice Gunes, and Ioannis Patras. Alone versus in-a-group: A multi-modal framework for automatic affect recognition. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)*, 15(2):47, 2019.

4 Wenxuan Mou, Hatice Gunes, and Ioannis Patras. Your fellows matter: Affect analysis across subjects in group videos. In *2019 14th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2019)*, pages 1–5. IEEE, 2019.

5 Hanan Salam, Oya Celiktutan, Isabelle Hupont, Hatice Gunes, and Mohamed Chetouani. Fully automatic analysis of engagement and its relationship to personality in human-robot interactions. *IEEE Access*, 5:705–721, 2016.

4 Working Groups

4.1 Working Group on Datasets with Humans and Agents in Groups

Patrícia Alves-Oliveira (ISCTE-IUL & INESC-ID, PT & Cornell University, US, patricia_alves_oliveira@iscte-iul.pt)
Elisabeth André (Universität Augsburg, DE, andre@informatik.uni-augsburg.de)
Tony Belpaeme (Ghent University, BE & University of Plymouth, UK, tony.belpaeme@ugent.be)
Filipa Correia (INESC-ID & University of Lisbon, PT, filipacorreia@tecnico.ulisboa.pt)
Kobi Gal (Ben Gurion University, IL, kobig@bgu.ac.il)
James Kennedy (Disney Research, US, james.r.kennedy@disney.com)
Iolanda Leite (KTH Royal Institute of Technology, SE, iolanda@kth.se)
Catharine Oertel (TU Delft, NL, C.R.M.M.Oertel@tudelft.nl)
André Pereira (KTH Royal Institute of Technology, SE, atap@kth.se)

Data is, to put it simply, another word for information. Recent years have seen a growth in datasets capturing data relevant for group interactions. While data was predominantly collected to form a quantitative understanding of social interactions, the recent attention for machine learning has resulted in datasets that contain features and annotations which are aimed at classifying or predicting social behavior. Due to the data-hungry nature of machine learning and the low cost of storage, the size of typical datasets is now several orders of magnitude larger than a decade ago.

We describe what constitutes a good dataset, and provide a number of resources related to datasets with a specific focus on group interactions, either between human-human, or human-agent groups. We provide dimensions upon which datasets for group interactions involving social agents could be characterized and classify existing datasets accordingly. Such a resource would be helpful for the research community to quickly identify data to work with, or identify gaps in available data. To conclude, we discuss several challenges in producing and annotating datasets, with a focus specifically on aspects pertinent for groups.

4.1.1 Where to find and deposit datasets

While datasets used to be stored on personal repositories, there is now an opportunity to host data on more persistent and curated sites. There are many online data repositories and as a researcher, care needs to be taken in choosing which repository offers the best conditions for your data. The most popular repositories are listed below, in order of relevance to the study of group interactions.

- **Zenodo** (https://zenodo.org/) is a general-purpose open-access repository developed under the European OpenAIRE program and operated by CERN, using the same systems to store data from the Large Hadron Collider. Datasets are given a digital object identifier (DOI), making it easily citable, and works together with GitHub allowing code and data to coexist.
- **Open Science Framework** (https://osf.io/) is run by the US non-profit organisation Center for Open Science, which facilitates open collaboration in science research. It was created in response to the replication crisis in psychology. Data and code can be stored in OSF Storage, the OSF public repository.
- **FigShare** (https://figshare.com/) started out as an online statistical figure repository, allowing authors to share high-resolution versions of printed material which can be interactively explored. Today it is an open access repository for a range of research outputs, including figures, reports, datasets, images, and videos. All deposited material is given a DOI. It is owned by Digital Science, a technology company based in the UK.
- **Github** (https://github.com/) started out as a versioning control system for computer code, but has increasingly been used to store datasets. Often datasets are stored in combination with code to analyse or report on the data.
- **Linguistic Data Consortium** (https://www.kaggle.com/) is an open consortium of universities, companies and government research laboratories. It stores data related to language, such as audio recordings of speech and text corpora, and is often used for Natural Language Processing and machine learning. Dryad Digital Repository (datadryad.org) is a repository with a focus on medical sciences. It charges a fee for submitting data, but guarantees free access for academic purposes.
- **Kaggle** (https://www.kaggle.com/) stores a wide range of datasets aimed predominantly at data science and machine learning. Many datasets serve machine learning competitions or challenges, and have been released by commercial organisation looking for new ways to capitalise on data they hold. The website allows online data science exploration, offers cloud computing, serves as a code repository and learning centre for data scence and machine learning. Kaggle is owned by Google.
- **IDIAP Data Distribution Portal** (https://www.idiap.ch/dataset) is hosted by the IDIAP Research Institute in Switzerland. It holds a small collection of data, mainly aimed at machine learning. It is unclear if and how new data can be added.

A record of all data repositories for scientific research is maintained at Re3Data (https://www.re3data.org/), which lists and tracks anything from single datasets to large data repositories. Data On The Mind (http://www.dataonthemind.org/) holds a modest list of datasets spread over the web. A promising development, with the potential to grow into a powerful tool for researchers, is Google's Dataset Search (https://toolbox.google.com/datasetsearch). It aggregates data collections from across the web, but at the moment does not have a function to return only results useful for particular purposes, such as machine learning or scientific analysis.

4.1.2 Available datasets for group interactions

Human-human datasets are important for the development of social agents in groups and teams. They can provide insights for understanding human behavior, be used to derive agent behavior (e.g., through machine learning techniques), or be used in evaluating algorithms. Many datasets that contain human-human interactions are publicly available. There is also good variability in terms of the activity and settings in which they have been collected.

The availability of human group interaction datasets partially arises from multiple communities that have a focused problem that they wish to solve. For example, the Emotion Recognition in the Wild Challenge 2019, which has a focus on group cohesion prediction (https://sites.google.com/view/emotiw2019). Additionally, datasets receive more attention in some communities, to the extent of having dedicated conferences (e.g., the International Conference on Language Resources and Evaluation – LREC).

In contrast, there are few datasets currently available that contain group interactions with artificial agents. This is likely in part due to the greater quantity of work that has been conducted in dyadic scenarios with artificial agents, but could also be due to the greater diversity of goals for research involving agents in social groups. Creating challenges for the community could provide impetus for collecting and sharing datasets for specific problems. There would subsequently be the possibility of reusing these datasets for solving other issues.

Furthermore, while many research institutes run empirical studies involving agents and humans in groups, few of these studies are captured and shared as datasets. This additional step poses many significant challenges (see 4.1.6), but could also provide a great source of data for the community.

Group datasets with one agent or robot

This section provides all of the publicly available group datasets including at least one agent that we are aware of:

- The Vernissage Dataset (http://vernissage.humavips.eu/). This dataset includes human-robot interactions with multiple participants and the commonly used robot platform NAO. It was collected using a Wizard-of-Oz protocol and has multiple camera views, audio streams, robot behavior logs, and some annotations.
- UE-HRI dataset (https://www.tsi.telecom-paristech.fr/aao/en/2017/05/18/ue-hri-dataset). The UE-HRI dataset consists of recordings of humans interacting with the social robot Pepper. It has spontaneous interactions in a naturalistic setting. There are a mixture of dyadic and triadic interactions. A variety of sensor data is available including video, audio, depth, sonar, laser and user touch inputs.

Group datasets with only humans

This section provides an overview of some human group datasets. The list here is by no means comprehensive, but is used to give some insight into the types of datasets that are publicly available and how they have been used in research:

- Elea Dataset (https://www.idiap.ch/dataset/elea).
  The aim of the Elea corpus was to create a resource to study group interaction. It is a multi modal corpus featuring both audio and video data. Annotations of both gaze and

voice activity among others are available. The corpus is particularly suited for studying group dynamics in terms of group performance measures.

- AMI (http://groups.inf.ed.ac.uk/ami/download/). The AMI corpus is one of the largest openly available corpora. It is a multi modal meeting corpus. Both audio and video data is available as well as some annotations in gaze and voice activity. The corpus has been used amongst others to study dominance.
- WOLF (https://www.idiap.ch/dataset/wolf). The WOLF corpus is based on a role-playing game where some people take on the roles of werewolves or villagers. The game is designed around deception, which has also been the research question most commonly addressed with this corpus. It is a multi modal corpus including both audio and video data.

4.1.3 Dataset dimensions

Datasets should clearly state the characteristics and procedures that were utilized during the data collection. The following list tries to specify several dimensions to describe the content of the dataset. Examples of existing datasets using these dimensions are present in Table 1.

1. **Resource name**: State the full name for newly created resources, followed by the acronym, if any.
2. **Size of the dataset**: Put the size of your resource on the basis of the group and data amount (interactions, time length, group sizes, and make-up).
3. **Demographics**: Describe the participants present in your dataset (e.g. "Newborns", "Children", "Teenagers", "Adults", "Elderly", "Mixed").
4. **Activity**: Explain what people are doing (e.g. playing a game, watching a movie together, ...) and the type of agent that is being used.
5. **Languages**: State the language spoken and/or read in the dataset (e.g., English, German, Japanese, none).
6. **Modalities**: Choose an appropriate label or combination of labels for describing the recorded data: "Visual", "Audio", "Physiological", ...
7. **Annotation**: Describe how and if data was labeled by using: "Human Labelled", "Automatic Labelled", "Not Labelled".
8. **Availability**: State the availability of the resource for the community. If the dataset is available on the web, at least for research ("Open"); if data is associated to an institution ("From Data Centers"); dataset distributed directly by the owner, usually associated with informed consent restrictions ("From Owner"); or other ("Other").
9. **Setting**: Describe where the dataset was collected (e.g. in the lab, public open space, school, etc.).
10. **Resource production status**: State is the resource already existed or if it was newly created. For newly created resources, describe if the production is completed ("Complete") or if work is still in progress ("Work-in-Progress"). In the case of an existing resource, describe whether it has been simply used ("Existing-used") or you have updated or modified it ("Existing-updated").
11. **URL/DOI/Publication** (if available): Indicate the URL of the resource/tool/guidelines described, if it exists, including the URL of the resource documentation if available.

4.1.4 Annotation tools

The following list contains examples of tools that can be used to perform the annotation and the statistical analysis, in some cases, of the dataset.

Table 1 Dataset dimensions for Humans and Agents in Groups.

Resource Name	Size	Demographics	Activity	Languages	Modalities	Annotations	Availability	Setting	Status
Vernissage	13 sessions (about 11 minutes each) of NAO interacting with two persons	Adults	Quiz activity about art and culture	English	Audio Video Mocap Robot logs	Speech transcriptions and several nonverbal cues such as 2D head-location, nodding, visual focus of attention (VFOA) and addressees.	Open	Lab	Complete
UE-HRI	54 interactions Aprox. 9 hours 1 agent, 1 or more humans	Adults	Social chit-chat	French?	Audio Video Depth Sonar Laser UI input Robot logs	Engagement	Open	Public space – university hallway	Partially annotated

- ELAN https://tla.mpi.nl/tools/tla-tools/elan/
- Noldus ObserverXT https://www.noldus.com
- NOVA https://github.com/hcmlab/nova – Nonverbal Analyzer is a tool for annotating and analyzing behaviours in social interactions. It supports Annotators using Machine Learning already during the coding process. Further it features both, discrete labels and continuous scores and a visualization of streams recorded with the SSI Framework

4.1.5 Feature extraction tools / feature set

The following list contains examples of tools that automatically detect some of the features commonly used in the annotation of human behaviour.

- OpenSmile (audio features) https://www.audeering.com/opensmile/
- EmoVoice (audio features) https://github.com/hcmlab/emovoice
- GeMAPS: Geneva Minimalist Acoustic Parameter Set [5]
- OpenFace https://github.com/TadasBaltrusaitis/OpenFace
- OpenPose https://github.com/CMU-Perceptual-Computing-Lab/openpose

4.1.6 Open Challenges

Datasets gain value from being shared, but data on social interactions will always contain data collected by recording people. There are formidable challenges facing dataset collection and dissemination:

- **Logistics:** Collecting data with groups is challenging for practical reasons. If our aim is to have a fixed group size, recruiting participants and ensuring that the same group size number in all sessions can be difficult. This is especially challenging when considering repeated interactions between the agent an the same group of users.
- **Annotation:** Group interactions are inherently more complex than dyadic interactions which means that the collected data and consequent annotations can become more noisy. For example, many tools for automatic feature extraction/annotation do not necessarily support multiple users (or only consider a limited set of users), and even human-annotated labels can become more subjective (e.g. coder agreement might decrease as the number of participants increase).
- **Gold-Standard:** The interpretation of nonverbal data is highly subjective. Thus the question arises of how to get a golden standard. Should it rely on the assessment of

(multiple) observers/annotators or on the assessment of the human group members involved in the interaction with the robot(s)? In general, obtaining gold standard data requires significant human effort.
- **Generalizability:** Datasets are collected in heterogeneous settings (see above). Can we transfer findings from one setting to the other? For example, can we transfer findings from a setting with two robots and one human to a setting with three robots and two humans? Can we generalize data from agents with different embodiments (e.g., robots vs virtual agents)?
- **Privacy:** There is an increasing understanding of the importance to share data to facilitate a comparison of algorithms and approaches. On the other hand, there are also increasing awareness to privacy, which constraints the use of data. Given recent initiatives to protect citizens, such as the European General Data Protection Regulation (GDPR)[1] or the Children's Online Privacy Protection Act (COPPA)[2] in the USA, the storage and use of data are now heavily regulated.
- **Persistence/Access:** Terms and conditions of the data collection require a legal support and ethical revision. Regarding the data collection of human behavior in group settings, one of the main challenges is to grant access to any participant without exposing the data of another participant. A possible solution is to apply anonymization tools (e.g. [1]) that would allow: (1) to grant participants their right to access their own data without revealing the identity of other participants; (2) the extraction and labeling of behaviours without revealing the identity of participants. Another challenge is the actual persistence and storage of the data and the fact that the withdrawal of one participation may compromise the data of other participant(s).
  In addition to existing previously mentioned challenges, terms and conditions and written consent are often worded to support the immediate goal of the study or technology for which the data was collected, but are not necessarily drawn up to support future use by others. For example, a consent form can stipulate that "data can be used by the research team for future analyses", but this wording limits future use only to the research team affiliated with the institution who collected the data and does not allow non-affiliates to use the data. Also, the insistence that data can only be used for academic purposes often sits in the way of current developments in the industry, where commercial software, often hosted on cloud services, is used to analyze data, thereby blurring the distinction between academic and commercial access to the data. Another issue is that it may be technically difficult to collect consent from every individual that generates data.

4.1.7 What makes a good dataset?

There is a huge proliferation of available datasets for designing and testing machine learning algorithms. Most of these datasets are collected and designed to solve a small defined machine learning problem and too specific to extract principles for humans and agents in groups. However, some recommendations for designing good datasets in the machine learning community can be useful as a basis for datasets that have been collected for studying human/agent group interactions.

Datasets should contain examples of research questions that can be answered or studied by using the data. Ideally, it should enable the community to work on their own research

[1] https://en.wikipedia.org/wiki/General_Data_Protection_Regulation
[2] https://en.wikipedia.org/wiki/Children%27s_Online_Privacy_Protection_Act

questions or models by using the same data. Most datasets extracted from real-life applications contain noise, missing values or irrelevant data. A good dataset should be cleaned from these problems before it is shared. This makes the dataset easy to work with and new contributors or users do not have to repeat this process again. Data pre-processing should also be performed to optimize the data for visualization (e.g., a txt file for Elan) or machine learning (e.g., data normalization) purposes.

The annotation schemes and features contained in the dataset should be clearly labelled and motivated by research in human-agent group interactions. All behaviors that contribute to understanding group behavior in human-agent group interactions should be annotated. This includes both extracting important features from the humans in the interaction but also logging all of the important timings and behaviors from the intelligent agents. Finally, to increase the usefulness and reusability of these datasets, they should be collected in settings that promote authentic and natural group interactions.

4.1.8 Conclusions and future challenges

Datasets are one of the most cited research outcomes, demonstrating the potential of sharing data with the community (e.g., [4, 8, 6]). Several fields of research have already acknowledged this impact and are actively contributing to data sharing. For example, the International Conference on Language Resources and Evaluation encourages and publishes datasets as a central part of conference contributions; MediaEval[3], a benchmarking initiative dedicated to evaluating new algorithms for multimedia access and retrieval, makes their datasets publicly available after work is concluded; the AVEC conference (Audio-Visual Emotion Challenge) and the Interspeech Computational Paralinguistics Challenges (ComPaRe) has challenges that release the associated datasets in open-access.

Despite the obvious benefits of dataset sharing, this culture have not yet been fully adopted by the field of Humans and Agents in Groups. Many factors are associated with the lack of dataset sharing culture among this field of research. Firstly, there is no venue that specifically values dataset sharing and publication. Secondly, researchers face many challenges when facing the option to share data, mainly because this field collects data from human participants which are associated with with legal, ethical, and privacy policies and restrictions. However, there are several benefits associated with sharing data. Specifically, sharing data advances the pace of research. Usually, in the Humans and Agents in Groups field, researchers collect new data for every new study performed. This requires extra research time due to the recruitment of participants and data collection. By adopting a dataset sharing culture, studies can be performed by re-using already existing materials and resources and analysis can be performed with already collected data. Additionally, research quality and transparency increases since analysis can be performed with already existing data, supporting reproducibility in this domain of research [7].

We present in this document the first attempt to share datasets within the field of Humans and Agents in Groups. We have defined best policies for datasets by defining what a good dataset is, what are the challenges associated with dataset sharing, and proposed interesting solutions to attend those. A major contribution of this work was to state the characteristics of the datasets since researchers can use the defined dimensions to share their datasets and also will ease information retrieval. Future directions to stimulate dataset sharing among this research field includes the organization of a workshop on datasharing (possible venues: HRI, CHI, CSCW), and the creation of challenges or competitions (similarly to DARPA or a simple competition in Kaggle).

[3] MediEval dataset: http://www.multimediaeval.org/datasets/

References

1 Pascal Birnstill, Daoyuan Ren, and Jürgen Beyerer. A user study on anonymization techniques for smart video surveillance. In *2015 12th IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS)*, pages 1–6. IEEE, 2015.

2 Rupert Brown and Samuel Pehrson. *Group processes: Dynamics within and between groups.* John Wiley & Sons, 2019.

3 Ewart J. de Visser, Marieke M. M. Peeters, Malte F. Jung, Spencer Kohn, Tyler H. Shaw, Richard Pak, and Mark A. Neerincx. Towards a theory of longitudinal trust calibration in human-robot teams. *Int. Journal of Social Robotics*, in press.

4 Li Deng. The mnist database of handwritten digit images for machine learning research [best of the web]. *IEEE Signal Processing Magazine*, 29(6):141–142, 2012.

5 Florian Eyben, Klaus R Scherer, Björn W Schuller, Johan Sundberg, Elisabeth André, Carlos Busso, Laurence Y Devillers, Julien Epps, Petri Laukka, Shrikanth S Narayanan, et al. The geneva minimalistic acoustic parameter set (gemaps) for voice research and affective computing. *IEEE Transactions on Affective Computing*, 7(2):190–202, 2015.

6 Nils Petter Gleditsch, Peter Wallensteen, Mikael Eriksson, Margareta Sollenberg, and Håvard Strand. Armed conflict 1946-2001: A new dataset. *Journal of peace research*, 39(5):615–637, 2002.

7 Steven N Goodman, Daniele Fanelli, and John PA Ioannidis. What does research reproducibility mean? *Science translational medicine*, 8(341):341ps12–341ps12, 2016.

8 Brian McFee, Thierry Bertin-Mahieux, Daniel PW Ellis, and Gert RG Lanckriet. The million song dataset challenge. In *Proceedings of the 21st International Conference on World Wide Web*, pages 909–916. ACM, 2012.

9 Jurriaan Van Diggelen, Mark Neerincx, Marieke Peeters, and Jan Maarten Schraagen. Developing effective and resilient human-agent teamwork using team design patterns. *IEEE intelligent systems*, 34(2):15–24, 2018.

4.2 Working Group on Design

Cindy L. Bethel (Mississippi State University, US, cbethel@cse.msstate.edu)
Kobi Gal (Ben-Gurion University, IL, kobig@bgu.ac.il)
Michio Okada (Toyohashi University of Technology, JP, okada@tut.jp)
André Pereira (KTH Royal Institute of Technology, SE, atap@kth.se)
Samuel Mascarenhas (INESC-ID & University of Lisbon, PT, samuel.mascarenhas@gaips.inesc-id.pt)

4.2.1 Introduction

Design is a cycle based on the target: redesign and evolution. It can be applied in several stages of the development of human-agent teamwork. We address different perspectives where design considerations are relevant to create of multi-party settings for humans and agents.

4.2.2 Interactive Design Ideas

We identified three different categories where design principles can guide decisions to successfully implement human-agent teamwork: the target, the levers, and the application domain.

Target

One category is the target which reflects the focus on the goals and the metrics to evaluate and optimize the system. Examples of targets include:

- Ethical Interactions (e.g. fairness, interpretability), which constrains the design space for each of the other areas considered
- Partnerships/Teammates (e.g. best representations, composition)
- Goals (e.g. education, help, companionship, entertainment, entertainment, work collaboration, persuasion)

Levers

Another category is the levers, which are the factors that can be controlled. Examples of levers include:

- Form Factor/Embodiment
- Actions/Behaviors
- Environment
  - Incentives
  - Visability
  - Communication protocols
  - Social Setting (e.g., collaborative, competitive)

Application Domain

Another category is the application domain, which includes the aspects related to the task and physical space where the agent will operate. The application domain both informs the physical environment design (e.g., Amazon Warehouse) and also informs the social environment of who they are interacting with. Examples of application domains include:

- Manufacturing
- Entertainment
- Education
- Assistance (e.g., physical or social)
- Healthcare
- First Responder (e.g., law enforcement, fire, search & rescue)

4.2.3 Design principles for robots and agents to interact in groups and teams

From the website 'The Undercover Recruiter'[4], design criteria for successful team interactions include:

- They communicate well with each other.
- They focus on goals and results.
- Everyone contributes their fair share.
- They offer each other support.
- Team members are diverse.
- Good leadership.
- They are organised.
- They have fun.

[4] https://theundercoverrecruiter.com/qualities-successful-work-team/

Inspired by other design principles (e.g., Dieter Rams[5] and Yves Béhar[6]), we propose the following design principles for robots and agents to interact in groups and teams:

1. Good design for agents in teams should support fairness and ethical interactions.
2. Good design for agents in teams should consider the emotions and reactions of its members.
3. Good design for agents in teams should optimize:
   - Performance
   - Communication (e.g. support, feedback)
   - Affective signals (e.g.,team motivation, enjoyment)
   - Interpretability and simplicity of interactions
   - Organization and structure
   - Role allocation and dynamic changes to roles
   - Diversity (e.g., cultures, gender)
   - Support and feedback
   - Fairness and responsibilities
   - Aesthetics and affordances appropriate to accomplish the goal
   - Simplicity of interactions
   - Usability
   - User experience

4.2.4 Open Problems

In our opinion, the largest open problem is that the scientific study of designing a heterogeneous systems is difficult because these three above design principles all inherently interact.

Other open problems include:

- How to consider adaptation in team dynamics? How do you adapt each team member to have positive results for the overall team goals and subgoals?
- How to design incentives for team participants that maximize team performance and also caring about the motivation of team participants? How do you incentivize the right people to do the right tasks?
- How to construct a team to address the joint goal in the best way to maximize results?
- When to intervene in a team interaction to help out maximizing for optimal team performance?
- How to delegate roles within a team for maximum benefit?
- How do you decide in mixed teams when to transfer control for decision problems? How do you decide when to hand off a problem to someone else?
- How to infer participants' plans and goals?
- How to design behaviors for long-term interactions with agents that are involved in team dynamics?
- How do you evaluate the policy online?

4.2.5 Design cycle for agents to interact in groups and teams

The design cycle for creating an agent to interact in groups and teams involves:

[5] https://hackernoon.com/dieter-rams-10-principles-of-good-design-e7790cc983e9
[6] https://www.fastcompany.com/3067632/10-principles-for-design-in-the-age-of-ai

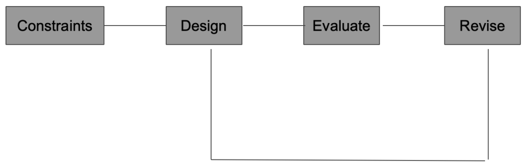

Figure 1 Design cycle.

1. **Elicitation** – Ask stakeholders about their needs, capabilities, and their issues. Consider the environment. How many agents? For how long? What exactly is the communication? How are you representing the information gathered (aggregate or personal data)?
2. **Representation** – Find the best representation that you can use to satisfy the targets/criteria that you have. (e.g. how are you giving badges in stack overflow?)
3. **Optimization** – Optimize either the environment or the policies of the agents.
4. **Evaluation** – Run randomized controlled studies to validate or determine what needs to be modified, then return to step 1 and repeat until optimal.

4.2.6 Promising Ideas

- Develop participatory design techniques to develop the best possible solutions over time in iterative design in partnership with the stakeholders. This allows for deployable systems in real-world applications and domains.
- Run randomized control studies on collaborative systems (e.g. stack overflow, reddit) that have large amounts of available data.
- Evaluate with actual team members in real-world scenarios and take it out of the lab.

4.2.7 Conclusions

In order to successfully create groups of humans and agents, the development of such agents and their tasks must follow certain criteria and guidelines. The discussion of this breakout group made a step forward towards the definition of those guidelines by creating a set of design principles for robots and agents to interact in group settings. Moreover, we propose a design cycle to address the stages and process of designing those interactions and we identified open problems and promising ideas within this topic.

References

1 Nihan Karatas, Soshi Yoshikawa, and Michio Okada. Namida: Sociable driving agents with multiparty conversation. In *Proceedings of the Fourth International Conference on Human Agent Interaction*, pages 35–42. ACM, 2016.

2 Ryosuke Mayumi, Naoki Ohshima, and Michio Okada. Pocketable-bones: A portable robot sharing interests with user in the breast pocket. In *Proceedings of the 7th International Conference on Human-Agent Interaction*, pages 211–213. ACM, 2019.

4.3 Working Group on Human-Agent Team Dynamics

Filipa Correia (INESC-ID & University of Lisbon, PT, filipacorreia@tecnico.ulisboa.pt)
Hatice Gunes (University of Cambridge, GB, hatice.gunes@cl.cam.ac.uk)
Stefan Kopp (Universität Bielefeld, DE, skopp@techfak.uni-bielefeld.de)
Iolanda Leite (KTH Royal Institute of Technology, SE, iolanda@kth.se)
Mark Neerincx (TU Delft & TNO, NL, mark.neerincx@tno.nl)
Catharine Oertel (TU Delft, NL, C.R.M.M.Oertel@tudelft.nl)
Rui Prada (INESC-ID & University of Lisbon, PT, rui.prada@tecnico.ulisboa.pt)
Marynel Vázquez (Yale University, US, marynel.vazquez@yale.edu)

4.3.1 Introduction

The Human-Agent Team (HAT) working group focused on the identification of key HAT-concepts and challenges. Working *definitions* of the key concepts were formulated, the corresponding *dimensions of change* and characteristics of *team processes and patterns* were worked out, example *scenarios* were proposed to exemplify the HAT dynamics in research and development, and *challenges* were derived from this exploration.

4.3.2 Definitions

The first discussions centred on the core definitions and the particular properties that make a collection of people a group or a team. The three following concepts have an inheritance-type of relationship: "collection of agents", "group" and "team".

Collection of agents
- Multiple agents with individual goals, abilities, skill, expertise;
- Having a minimal degree of autonomy (ability to decide on their actions);
- Not necessarily co-located;
- But with ways to interact with each other, closely or loosely coupled.

Group
- Agents are individually aware of having a shared identity (group, commonality);
- Awareness of in-group/out-group agents.

Team
- Agents have a joint goal or task they are working on;
- Agents are aware of their working on it together;
- Agents are committed to it and mutually support each other;
- (not necessarily interdependent).

4.3.3 Dimensions of change

Given the definition of teams in section 4.3.2, we identified the following dimension of change and some examples of each:
- **Team organisation**: structure, roles, and norms
- **Team members**: number, capabilities, and autonomy
- **Relationships**: trust, liking, intimacy level, and power

- **Group attitude**: commitment, (group) identity, and trust
- **Shared cognition**: experiences, knowledge, skills, and awareness
- **Group properties**: cohesion, interdependencies, performance (change), resilience, and distribution of participation
- **Context**: environment, task, resources, and tools

4.3.4 Team Processes & Patterns

The members of a human-agent team adapt their behavior to each other and the dynamic environment in which the team operates. Constructive and destructive behaviour patterns can be implicitly or explicitly brought forward [3], e.g. based on experience or just "emerging".

The following team processes can be distinguished [1]:

Forming the group. This first stage includes the tentative communication, uncertainty and exchange of personal information of the group members. It leads to the sense of belonging and the shared identity of the group.

Establishing norms and common ground. At this stage, the group develops standards and/or agrees on the procedures to operate.

Assigning roles with responsibilities. The roles attribute a certain structure to the group and usually are intended to improve the communication among members.

Planning, executing, monitoring and repairing the (group) tasks. This stage hold the actual performance of the group on the defined task.

Managing relationships and conflicts. Trust is an important construct, characterizing relationships in a team. Trust develops over time and adequate trust calibration is crucial for collaboration. When the benefits and costs of the relationship outcomes are harmonized for the concerning team members (i.e., there is a balanced *relationship equity*), these team members will collaborate well [2].

4.3.5 Envisioned Scenarios

Goal: autonomous agent (virtual/robot) that can perform in a team such that the team performs better. Possible scenarios (5 to 10 year vision):

- Health-care in a hospital setting by human-robot teams (healthcare professionals, nurses, social robots, ...)
- Collaborative assembly
- Search and rescue, disaster response, ...
- Entertainment, education, citizen-science...

4.3.6 Open Challenges

Open challenges from an individual level to a group level at the same time as from a low-level sensing to a cognitive decision-making level (see Figure 2):

- What sensing/perception and acting/expressive capabilities do robots need to have in teams?
- How can it maintain the necessary level of social awareness and task-related awareness?
- How to accommodate for social norms, culture, individuality of team-members, etc.?
- How can robots engage in constructive team coordination patterns and processes?
- How can robots/agents learn what role is best suited for that particular team?

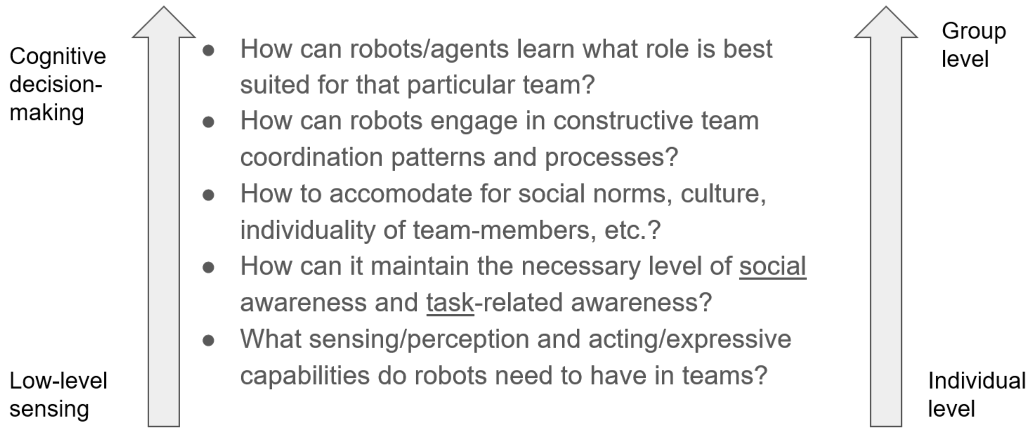

Figure 2 Open challenges for research & development of team agents, from low-level sensing to joint decision-making and from individual to team level.

4.3.7 Conclusions

The presented definitions, dimensions and team processes were discussed in the light of interpersonal group interactions and the some of the well-established dynamics of those groups [1, 4]. Therefore, this discussion highlighted the importance of redefining this concepts for groups (and teams) in which both humans and agents are part of. Nevertheless, some of the proposed challenges and the envisioned scenarios identify and guide future approaches and avenues to explore this topic.

References

1 Rupert Brown and Samuel Pehrson. *Group processes: Dynamics within and between groups.* John Wiley & Sons, 2019.
2 Ewart J. de Visser, Marieke M. M. Peeters, Malte F. Jung, Spencer Kohn, Tyler H. Shaw, Richard Pak, and Mark A. Neerincx. Towards a theory of longitudinal trust calibration in human-robot teams. *Int. Journal of Social Robotics*, in press.
3 Jurriaan Van Diggelen, Mark Neerincx, Marieke Peeters, and Jan Maarten Schraagen. Developing effective and resilient human-agent teamwork using team design patterns. *IEEE intelligent systems*, 34(2):15–24, 2018.
4 Donelson R Forsyth. *Group dynamics.* Cengage Learning, 2018.

4.4 Working Group on Social Cognition for Robots and Virtual Agents

Cindy L. Bethel (Mississippi State University, US, cbethel@cse.msstate.edu)
Merijn Bruijnes (University of Twente, NL, Merijn Bruijnes)
Malte Jung (Cornell University, US, mfj28@cornell.edu)
Samuel Mascarenhas (INESC-ID & University of Lisbon, PT, samuel.mascarenhas@gaips.inesc-id.pt)
Christoforos Mavrogiannis (University of Washington, US, cmavro@cs.uw.edu)
Simon Parsons (University of Lincoln, GB, simon.parsons@kcl.ac.uk)
Catherine Pelachaud (Sorbonne University, FR, catherine.pelachaud@upmc.fr)
Rui Prada (INESC-ID & University of Lisbon, PT, rui.prada@tecnico.ulisboa.pt)
Laurel Riek (University of California, US, lriek@eng.ucsd.edu)
Sarah Strohkorb Sebo (Yale University, US, sarah.sebo@yale.edu)
Julie Shah (Massachusetts Institute of Technology, US, julie_a_shah@csail.mit.edu)
Elaine Short (Tufts University, US, elaine.short@tufts.edu)
Marynel Vázquez (Yale University, US, marynel.vazquez@yale.edu)

4.4.1 Introduction

This section discusses the cognition required by physically embodied social agents that interact in a group context. We consider a *group* to be a set of agents (e.g., humans, robots, and virtual agents) that share certain characteristics or relationships. For example, group members might be connected by identity, location, or the beliefs that they hold. Consistent with existing work (e.g. [7, 30]) we consider teams as a specific type of group whose members "are interdependent in their tasks" [7], interact socially, and share a common goal.

While the term "social agents" covers a broad spectrum, ranging from chatbots to virtual characters to physical robots, in this article we focus specifically on *physically embodied agents*. This includes both physical robots able to enact physical change in the world and embodied virtual agents, who may appear within physical objects capable of engaging in groups with humans. We will thus refer to these as embodied social agents (ESA) throughout this work.

Social Cognition

Social cognition refers to how people process, store, and apply information about other people and social situations [45]. While social cognition is integral to how people perform in both work and non-work related interactions, research on computational models of social cognition largely focus on purely non-work related social interactions. There is evidence that this delineation undermines the task effectiveness and acceptance of robots and virtual agents. For example, medical professionals were found to actively sabotage a robot that was navigating within a hospital environment without consideration of the social norms that protect medical professionals from interruption during high workload situations [37]. Considering social cognition for groups and teams brings with it novel challenges that are yet to be addressed.

Social Cognition in Groups

To distinguish social cognition in groups from that in individuals, we draw from Brown and Pehrson's [5] work that highlights three concepts that are important when characterizing how people behave as group members versus individuals. The first concept, *social identity*, refers to how people define themselves in terms of a group and how they attach emotion and value to these self-categorizations. The second concept is *social context*, where people's social identities are dependent on the context in which they find themselves and the groups with which they are part of. This can have both macro-level societal structure as well as the micro-context of specific social environments. The third concept, *social actions*, refers to the affordances that groups offer in helping people enact change in the world. It also reflects the notion of collective action to achieve social change.

These distinctions are useful in framing how ESAs team in groups, where an ESA will have its own social identity of capabilities and functionalities that it can employ in group settings, as well a set of roles [53, 18], which may be known *a priori* or may be learned during interactions. For example, a robot tour guide may take the role of a mentor, teaching a group of students about an exhibit. After the museum closes it may adopt the role of a mechanic, helping workers take down an exhibit. In these different roles, the robot may exhibit very different capabilities and functionalities.

The social context can also be informative for an ESA to support its cognition, particularly with regards to supporting constraints on its behavior, establishing common ground, and informing its decision making [38, 41, 60]. Additionally, using context can help ESAs to reduce the problem space and bootstrap problem satisfaction [48]. Finally, social action can provide ESAs affordances for actions within groups. For example, a human group may be engaged in a shared workspace activity, and a robot can observe human-human interaction as a model for how it should behave within it [20, 22, 65].

4.4.2 Characterizing ESA-Group Interactions

We distinguish between different levels of cognition depending on how an ESA's actions take groups into consideration (Figure 3). First and foremost, ESAs need to be able to act with *groups in mind.* Next, ESAs can *act on groups* as outsiders, influencing or shaping aspects of group behavior. Finally, they can *act in groups* as another group member. At this level, they not only influence the group but also need to take part directly in the task and inherent dynamics of the interaction.

Figure 3 captures the idea that ESAs who *act in groups* are a subset of those who *act on groups* or can influence the outcomes of a group. Both of these types of ESA would be subsets of those who can *act with groups in mind.* This structure emerges from the fact that acting on groups requires all the abilities of acting with groups in mind plus some additional abilities. Similarly, acting in groups requires all the abilities of acting on groups but also requires some additional abilities.

Acting with Groups in Mind

This is the most basic level in which group cognition can be considered. At this level, ESAs account for different aspects of the group with its cognition irrespective of the number of people/ESAs who are part of the current interaction.

We consider any social interaction with two or more ESAs as group-situated to some degree. This position has implications for human-robot interaction (HRI) research as groups can be taken into account independent of whether the interaction that takes place is dyadic

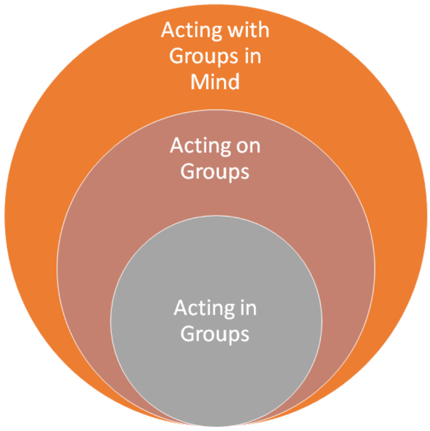

Figure 3 A framework for considering group cognition.

or more complex. For example, consider a robot that enters a room and hands over an object to a single person who is alone in the room. Here the robot is in a dyadic interaction, but can (and arguably should) take group membership (e.g. culture) into account when reasoning about how it completes its task. Of course, if the person is not alone, then there will be other aspects of group membership that will need to be taken into account, such as not interrupting a speaker.

Acting on Groups

Acting on a group refers to the level of group cognition that is involved when the purpose of an ESA's action is to influence or shape a group's structure or behavior without the ESA being a member of the group. Such influence is analogous to a dog herding a group of sheep. The dog shapes the behavior of the group of sheep without being a member of that group. An example of a robot designed to purely act on a group without being part of it is the Micbot [62]. The robot influences a group's participation dynamics by signaling engagement and by nudging group members to participate, though it is not directly a part of the group it influences. A key, implicit, idea in acting on a group is that the robot may work in a pro-social way to improve the functioning of the group. Doing this in a way that does not hijack the task that the group is working on will require the robot to provide subtle interventions, such as the Micbot probing members who have spent less time speaking. Being able to do this in a general, flexible, way, needs the robot to be equipped with a deep knowledge of group behaviour, and the ability to reason about this (in particular the ability to predict the consequences of specific interventions).

For ESAs to be able to effect groups, it is helpful to consider the kinds of specific affordances that groups offer for interaction and influence. Affordances are "properties of the world that are compatible with and relevant for people's interactions" [13]. Gaver's notion of affordances is different than Gibson's [15] original concept of affordances or Norman's [40] concept of perceived affordances. For Gibson, affordances naturally exist as properties of the world. Norman's concept of perceived affordances highlights that interactions are driven less by the actual properties of an object but rather by how it is perceived. For example a door might have the affordance to be pulled (based on its mechanical design) but the particular door handle design leads to a perceived affordance as "pushable." Since a robot is not dependent on human-like perception of the world, we adopt Gaver's definition of affordances here. Drawing from Gaver [13], we define group affordances as properties of *groups that are compatible with and relevant for interactions.*

Groups offer unique affordances dependent on characteristics such as structural composition (e.g. a leaderless group affords different interactions than a team with a pronounced leadership structure) or size (e.g. a small three-person group affords different interactions than a stadium full of people). ESAs can leverage such group affordances to influence their behavior. For example, Kwon and colleagues [31] showed that a robot can leverage a group's unique hierarchical structure to shape its behavior in systematic ways. While to our understanding there are no current approaches to reason about group affordances specifically, existing work has developed approaches for reasoning about affordances more generally. For example, Sarathy and Scheutz proposed a computational framework for inferring affordances [51]. Similarly, Shu, Ryu, and Zhoo [58] introduced an approach for learning "social affordances."

Acting in Groups

Acting in groups refers to a level of cognition that is required for robots to serve as members of a group and not just as tools used by the group. This involves abilities to influence groups and a general consideration of cognition as group-situated. The robot can still influence the actions of other group members as in the previous level but it also needs to have cognitive abilities that specifically pertain to the task or activity that the group is performing. As such, this level offers more opportunities for a robot to have more impact on the group than what would be possible for an outsider. It may also be more difficult to generalize its impact across different group tasks or activities. Groom and Nass argued that it is impossible for robots to become members of work groups or teams as they lack the basic abilities to build and maintain trust [19].

For an ESA to become a member of a group, it is not necessarily enough that it is able to identify whether a group exists or not. Additionally, the robot needs to be able to make sense of how the other group members perceive its membership status as well as what roles are expected to be allocated to the robot within the group. While not so important in highly controlled settings, these capabilities become quite relevant in scenarios where group formation and role allocation occurs in a highly dynamic manner.

4.4.3 Situation Awareness – Decision-making – Action: Group Interaction Throughout the Sense-Think-Act Cycle

Parasuraman, et al. [44] proposed a taxonomy that guides functional automation design. They identified four stages of human information processing that may be supported by automation: information acquisition (Stage 1), information assessment (Stage 2), decision and action selection (Stage 3), and action implementation (Stage 4). We explore how the same framework can guide the design of social interactions with agents. Many aspects of social interaction (depicted in Fig. 4) involve these four stages or processes of human cognition.

In our analogy, situational assessment is composed of processes for: *perception*, to perceive the state of other agents, the group or team, environment; *comprehension*, to assess the relationships among entities in the environment, including interactions between humans and robots; and *projection*, to infer the intent of other agents and predict plans, behaviors, trajectories. Under the Endsley model [11] for human situational awareness, this subsystem represents the comprehension and projection elements of situation assessment. Decision-making may be a skill-, rule-, or knowledge-based [46]. Action can include taking physical action, or implicit/explicit communication and signaling.

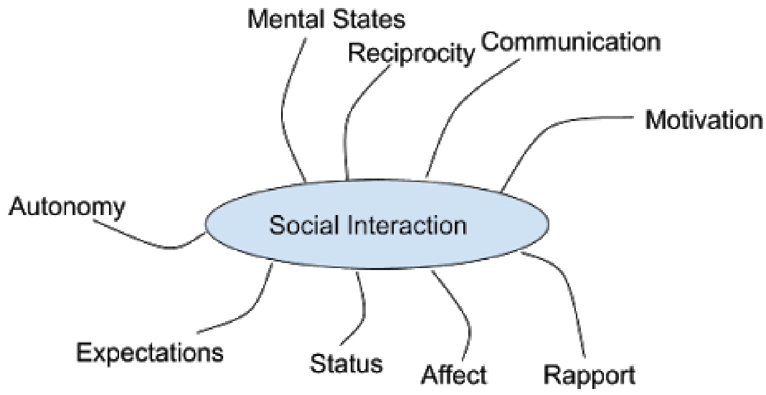

Figure 4 Dimensions of social interaction.

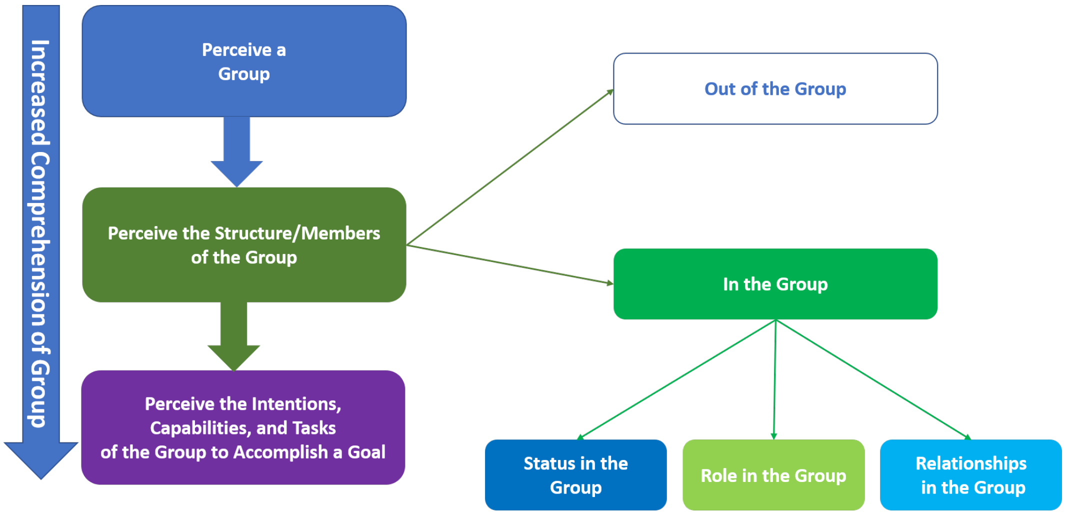

Figure 5 Overview of increased comprehension of group understanding.

As depicted in Figure 5, when agents (including humans and/or robots), encounter one another they first perceive the group. Once the group is perceived or identified, the next phase is to perceive the structure or members of the group. As part of that process, a determination needs to be made of whether individual agents are members in the group or if they are outside of the perceived group. For those agents perceived as being in the group, members will have roles, statuses, and there are relationships between those members that must be recognized and understood. The next phase is to then perceive the actual intentions, capabilities, and tasks that the group and its members need to consider to accomplish the goals or intentions of the group. From the moment of perception of a group there are increasing levels of comprehension or understanding of the dynamics of the group and its members.

In the case of new members joining a group, the members first tend to imitate the actions and behaviors of existing members of the group (refer to Figure 6). They follow the examples of those established members until they obtain more knowledge. As learning occurs over time, a certain level of knowledge is achieved and the member more fully integrates into the group. Imitation becomes knowledge through the process of learning about the intentions, capabilities, and tasks of the group, which allows for members to better understand why they are performing certain behaviors to be more fully integrated as part of the group.

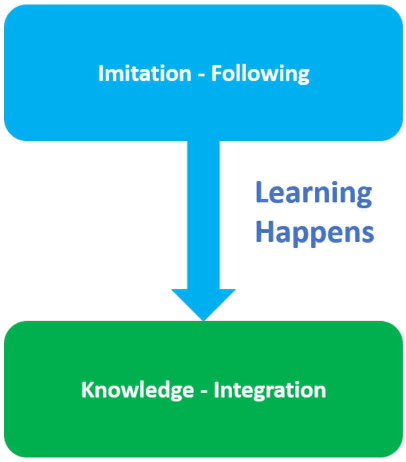

Figure 6 Members of groups transform from imitation to integration in their behaviors.

To further expand on the transformation from imitation to integration, group dynamics will often begin with a system of reactive behaviors. This can be especially true when new members are involved in the group. During these types of group interactions, the behaviors tend to be more reactive in nature. A situation in sensed or perceived and the group responds or reacts. Over time, as learning occurs, the dynamics begin to transform and the behaviors become a hybrid of deliberative decisions and reactive responses [1]. Members of the group have a better understanding based on experiences and knowledge of how to decide to respond to the situation encountered. During these types of interactions, a situation or encounter is perceived, followed by comprehension of the potential impacts of the situation, and then the agents are better able to decide the best response to perform. Additionally, there may be an element of planning or deliberative thoughts that occur once comprehension of the situation happens, planning can then occur for how to react, and then the appropriate response is performed to accomplish the goals or tasks for the group.

We are interested in the unique challenges posed by social cognition in groups, where assessing relationships among agents and projecting intent and behavior must be performed across varying time-scales and levels of information abstraction.

For example, an agent addressing specific people in the group can affect not only the behavior of the target user, but also the cohesion of the group as a whole, in both the short- and long-term [57]. Furthermore, the effect of agents on groups includes short-term phenomena, such as rapport or engagement, that lead to long-term effects such as trust or friendship. This means that agents need to detect, reason about, and take actions that will affect the group dynamics on multiple time scales, from seconds to years. Another challenge is the difficulty of computationally representing these social phenomena. Ideas such as engagement, trust, intent, and discomfort need to be translated into reward functions, actions, preconditions, and so on. In the current state-of-the-art that accounts for social considerations, it is often heuristic, because it is hard to measure social features and outcomes in a computational way (and to enumerate all of the possible failure modes). Assigning credit for social interactions and reasoning about the consequences of actions needs to be done over time and among group members. Finally, acting in groups magnifies the effects of actions: for example, humans can ally with each other with or against the agent, forming subgroups. Effects of agent behavior are also magnified in multi-agent scenarios because it increases the influence of the agent(s) as a group.

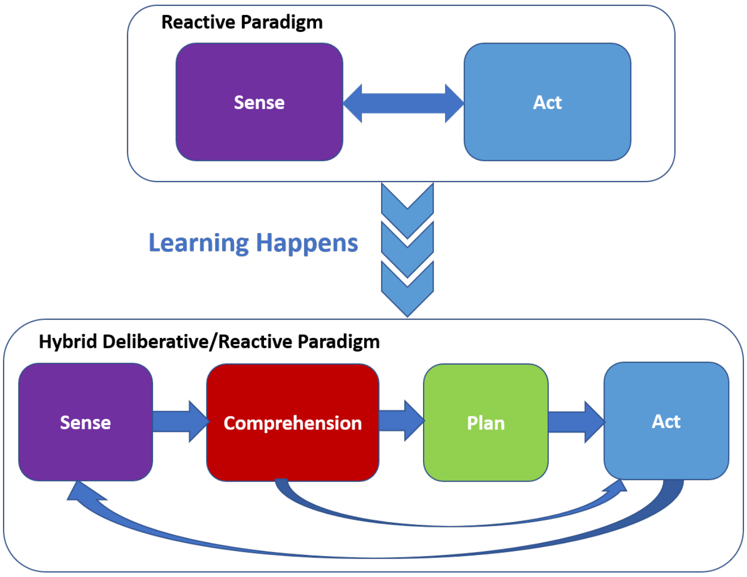

Figure 7 Reactive and Hybrid Deliberative/Reactive Paradigms for interactions in groups.

4.4.4 Group-Specific Considerations

Clearly it is harder for an ESA to deal with a group rather than an individual. We distinguish two aspects to this. First, it is often the case that group settings lead to phenomena not typically observed in one-to-one interactions. Second, groups give rise to elements that are not present in one-to-one interactions. Among the group-specific phenomena, we can identify the following:

Diffusion of responsibility. This has two effects: first, humans might be less willing to just take care of things, because there are more agents, and second, each ESA has less responsibility, even if the group as a whole has more effect.

Conforming. ESAs may conform to people within a group and people also will likely conform to agents within a group [50]. It is important to know when to conform and how to wield their influence when trying to have people conform to ESAs.

F-Formations. Another type of group phenomena that can emerge are *F-Formations*: spatial formations typical of situated group conversations which result from the need of interactants to share information and perceive one another during the interaction. While dyadic formations can be observed in one-to-one interactions, groups may lead to other types of configurations worth considering in HRI. For instance, detecting F-Formations can aid in identifying social interactions in human environments and, ultimately, help enable appropriate ESA behavior in human environments [64]. The disposition in space of the group should dynamically evolve as a new member join or leave the group. F-Formation should be updated on the fly to reflect these changes.

Modeling social influence. In groups, the ability to reason about social influence is critical to being able to determine the effects of one's actions. For example, the ESA may directly influence (nodes connect within a graph) some people and indirectly influence other people (as a result of the direct connections). Feedback will likely also be helpful in refining these models after actions are taken.

Turning to elements that are not present in one-to-one interactions, we have:

Group size. The size of the group in which interaction is taking place affects that interaction [66]. Large groups behave differently than small groups [63]: there may be local patterns, and an ESA may not need to know everything else. There is also a need to consider sub-groups, which may be dynamic, so relations may differ between members of the overall group, the members of a subgroup, and relations between subgroups.

Context. Context helps to constrain what is important in the interaction. This can happen through constraints on behavior (norms), common ground, and what actions other members of the group are likely to take. For example, in a negotiation (or other adversarial interaction), an individual would want to track how open the other group members seem to their proposal, and the individual might want to be less open about emotional states, and would expect other members of the group to engage in negotiation behavior.

Social status. Successful interactions must involve an awareness of the social status of group members. While issues of status are most salient when working in groups, even when interacting with a single person, status is still relevant because a person or ESA's status is a result of relationship with groups of people.

Finally, group interactions involve unique **conversational decision making skills** that ESAs must develop in order to verbally communicate. Examples of these conversational decision making skills include selecting whom to address and choosing whether or not to interrupt someone.

4.4.5 Social goals of ESAs within groups and teams

There are both long and short term social goals that aid an ESA's interactions on a group, in a group, and keeping groups in mind. These social goals may apply to dyadic interactions, however, become exponentially harder in groups due to the increased complexity of relationships and interactions.

Long-term group dynamics

ESAs must build and maintain relationships with people in their group or team, as well as the group itself. Stable relationships can be important for resilience to agent failure. Maintaining relationships involves, for example, taking into account the affective state of the group and its members. But, specifically in groups, the concept of in-group / out-group should be taken into account [3, 67] as the inter-group relations may influence the interpersonal ones. Properly building and maintaining relationships is important to manage group cohesion, stability, sense of shared identity and common purpose. Complying to politeness theory [4], for example, could be a means to achieve such goals. Some dimensions to be considered are:

Managing affect. At a minimum, in order to have an effective role within a group or team, relationships with the people in the group must remain positive and the ESA must avoid hindering the goals and work of the people with whom it is interacting [25, 37].

Building and maintaining interpersonal trust. Interpersonal trust is distinct from task-related trust. Interpersonal trust is related to long-term rapport and more of a function of the relationship between ESAs, independent of context. McKnight and Chervany [36] identify trust as conceptual categories competence, benevolence, integrity, predictability, and other characteristics such as open, careful, safe, shared understanding, and being personally attractive.

Managing social influence. ESAs should identify and manage different sources of influence among group members, for example, social power, friendship, expertise and status. They should find a balance between the leadership dynamics supported by these sources.

Building and maintaining group commitment. ESAs can be committed to be part of a group and/or to actively achieve the group's shared goal. Ensuring an ESA's commitment, as well as other members' commitment, over time is the result of a process that involves coordinating the ESA's and other members' participation in the group, working toward and maintaining the group stability, task progress and social cohesion.

Building and maintaining shared cognition. Through the group interactions the members work towards establishing common ground, shared knowledge and representations, shared experiences and mutual understanding. Moreover, each member of the group builds a cognitive model of self and all other members of the group and updates this as the interaction unfolds. This model may include individual mental and affective state, level of commitment, task responsibility and expertise, and social norms.

Short-term interaction in groups

There are shorter-term goals that will build into the long-term ESA-group relationship. Many of these phenomena have been studied in one-on-one HRI, but relatively fewer in group interactions. For example:

Avoiding discomfort. ESA behaviors may often cause human discomfort. The mechanisms underlying the emergence of discomfort as a result of agent behaviors are not sufficiently understood. This results in ESA behavior design resulting in undesired phenomena, such as the reciprocal dance in navigation [12].

Adhering to social norms. In order to maintain their membership in groups as well as relationship with group members, an ESA must adhere to social norms [33]. This may include being polite, being respectful toward others, and knowing when it is appropriate to be more direct [16].

Expressing intent. In order to ensure effective social interactions with groups of humans, it is crucial for an ESA to be able to succinctly express its own intentions. Ongoing work looks at encoding intentions into ESA actions (e.g., collaborative manipulation [10], navigation [35]). However, depending on the task and the context, different mechanisms, modalities, and strategies of conveying intent may be more applicable or effective. Research on identifying and understanding these mechanisms would enable ESAs to adapt to the social context more naturally.

Inferring human intent. Social interaction with people requires an understanding of human intentions. By leveraging signals encoded in human behavior, ESAs may infer latent human desires, preferences and goals. Past work has focused on interpreting task-specific human intentions to achieve desired performance in shared autonomy applications (e.g., [2, 24]). Future work could expand this work into the group domain to enable seamless interaction and coordination.

Repairing trust and expectations after errors. The ESA must assume that it will make errors and will need to detect and repair those issues [54]. This is important so that when a large mistake comes along, the consequences are not as extreme.

Setting expectations. ESA expectation setting can help manage trust and prevent large loss of trust if expectations are too high. Members of a group share responsibility for the social affairs in the group and compromising the initial expectations can have a detrimental effect. In human-human interactions such violations often trigger repairs. Compare the door in the face (DITF) and the foot in the door (FITD) techniques [9], where asking for an initial favor influences the acceptance of a request for a further favor.

Monitoring group state and status. A probe could be used to affirm or check one's status within a group, or to get the "vibe" of a group. Repairs can also act as probes to ensure that relationships are maintained.

Adhering to social roles expectations. When taking a specific role in the group, the agent must comply with the expected responsibilities and affordable behaviours.

Choosing the right member to address. In the case of groups, carefully considering the choice of the interaction partner in a particular moment is very important as it signals all participating members. Addressing the group as a whole is an option as well. This choice highly affects the group dynamics and the development of interpersonal relations, in particular, the attitudes towards the ESA.

Recognizing and expressing group-based emotions. Group-based emotions are particular emotions that come into play when considering group interactions. These are emotions that result from an appraisal process in which group identity is higher in salience than self identity [17]. ESAs that are able to express appropriate group-based emotions can promote a stronger sense of group identification from all the members [8].

Expressing affect. Expressing affect (both positive and negative) plays a crucial role in social control between group members [52, 26]. For example, expressing negative affect may be important for some interventions in the group dynamics, such as, to signal members to change their current behaviours. On the other hand, expressing positive can encourage behaviors of others. An ESA can apply such expressions strategically related to the task (e.g. social skills training [6, 61]).

4.4.6 First Steps to Advance the State-of-the-Art

In this section, we propose some promising areas for development of new research, building on the state-of-the-art. We divide the areas into three categories: designing and studying ESA behaviors in groups; creating theories and models to inform ESA decision making within groups; and developing control and learning algorithms for ESAs to use in group interactions.

Designing and studying agent behaviors in groups

There are many human behaviors that ESAs can emulate to further both long and short term group interaction goals (see Section 4.4.5). These behaviors are likely best investigated through design research and user studies. We believe that the following examples will lead to fruitful advances in our understanding of ESA cognition:

- **Using repair behaviors.** Despite growing body of work within HRI on repair (e.g. dyadic human-robot trust repair [32, 54] and robot conflict repair between human group members [27, 56]), there are still many aspects of repair worth exploring. For example, how does the repair of a dyadic relationship may influence other members in a group? Does it improve overall group cohesiveness? Or, when is a repair necessary to recover the reputation of the team member who the ESA maligned?

- **Setting expectations.** How can an ESA appropriately set expectations of its capabilities in group contexts that may vary in group membership as well as the knowledge and background of individual members? Expectation setting is critical to perceptions of robots and the ensuing interaction [43]. Besides the initial introduction, this can also be a recurring process during the interaction. For example, when the goal of the group changes different abilities might be needed and a new negotiation on group roles can ensue (and thus the expectations of each group member). Another example, in an education setting a teachable agent (an agent that is introduced as a peer that understands less than the student) is shown to make students put more effort in learning [59].
- **Rapport.** What factors influence the rapport within a human-ESA group? How does rapport evolve over time in a group interaction? Are there differences in rapport between human-human and human-ESA interactions? How can we expand prior work on one-on-one rapport [47] to group interactions?
- **Expressing intent.** How can an ESA express its intent in a group where the members have a diverse set of backgrounds, viewpoints, and knowledge bases? Additionally, how can an ESA express its intent exclusively to a subset of a group and not the other members?
- **Inferring and respecting social norms.** Human interactions are governed by norms [14]. ESAs that participate in human group activities will need to be able to conform to those norms [33] and, ultimately, to be able to infer those norms from observing the interactions between group members. How can they effectively correct for wrong or incomplete models of social norms? Furthermore, how should ESAs deal with situations in which their group members conform to different norms?
- **Expressing and detecting discomfort.** What features of group human-ESA interaction are indicative of discomfort with the ESA? How can an ESA express discomfort with human behavior in order to better meet its needs without damaging the group interaction?
- **Tracking long-term relationships.** How do group human-ESA relationships change over long-term interactions? What features of the interaction are important to understanding these relationships?

Creating theories and models to inform ESA decision making within groups

In order to understand a group, ESAs will need to use appropriate representations and abstractions of the group and its functions. We highlight several challenges that require the development of new representations, abstractions, and models.

- **Representing state.** How can we effectively represent states? There are cases in which the information that needs to be represented in different ways as a function of varying groups or contexts. There is also the problem of the curse of dimensionality.
- **Modeling Influence.** How can we represent influence within a group? And which representations (e.g. graphs with nodes and weighted ties) would be helpful to show the dispersion of influence? How can we dynamically alter this representation?
- **Abstracting behavior.** How can we abstract/represent behaviors over space and time in a computationally efficient fashion? What is the right type and level of abstraction to capture salient yet computationally practical aspects of group behaviors? And how could an ESA decide on where to focus its attention within a long-term memory of observations? Existing work has looked at abstracting collective group behavior into symbolic structures (e.g., in navigation [34]). This allows for the use of classical planning and inference methods. However, it is still unclear how to formalize the abstraction across different tasks and contexts.

- **Modifying representations over time.** If an ESA receives feedback signals from people around it that are not congruent with its current representations, how will we allow the ESA to change its representations? For example, if a pizza delivery robot makes a delivery within a library and communicates at a normal volume with the librarian who ordered the pizza, those in the library may provide the robot with negative feedback signals. How can the robot modify its representations of delivery locations and appropriate behavior within them (e.g. speaking volume), so that it behaves appropriately in all of its delivery locations?
- **Modeling Adaptation/Consensus.** What is the right direction towards modeling adaptation and consensus within groups of humans and ESAs? How can we leverage multi-modal signals inherent in ESAs' behaviors to model the state of consensus within a group? And what is an effective way of incorporating such a model into the ESA's decision making? Existing literature has proposed decision-theoretic [39] and information-theoretic approaches [29]. Can we unify such models towards formalizing a general theory of adaptation for human-ESA interaction within groups? Some work on modeling human-multi-ESA interactions has been done [28].

Developing control and learning algorithms for ESAs to use in group interactions

In this section, we propose new areas for technical development of intelligent behavior for ESAs, including promising new areas for developing computational models and methods that will enable ESAs to successfully accomplish the goals discussed in the previous sections.

- **Probing for interaction and group state.** Active information gathering can enable ESAs to collect specific information faster than with passive observation [55]. It can help robots understand the mental states of other ESAs [49]. In the context of group interactions, it would be interesting to explore what types of behavior allow an ESA to most effectively obtain information about a group? How do groups respond to information-seeking and group-evaluating ESAs? How can active information gathering actions help repair interactions in groups? Additionally, what new methods are needed to take advantage of the agent's embodiment and role in the group?
- **Sensing and assessing groups.** How can we measure key components of group interactions using a virtual agent or robot group member? How can we build on current work, for example on expertise assessment [20, 21, 22, 23]? How can we infer intentions, dominance relationships, and group affect? Another challenge is to deal with groups of multiple independent ESAs and one or more humans [42]. Especially when groups of multiple agents form 'in the wild', our models might not be able to recognise social behavior or other ESAs.
- **Simultaneous planning for task and social goals.** How do current planning methods need to be modified to allow planning for at least one social goal at the same time as task goals?
- **Using transfer learning.** How can we develop knowledge over time about a single group? How can we transfer modeled social dynamics from one task to another?
- **Building rapport** How can we use control theory or other continuous representations to build rapport between an ESA and a group? How can we make social agents in groups more responsive to human behavior? How can we make the ESA's responses more intelligent and relevant?
- **Detecting anomalies.** What anomalies can occur in social signals? What new methods might allow us to detect these anomalies and reason about their implications?

4.4.7 Conclusion

This report outlines the social cognitive abilities required for ESAs to perform effectively in a group context. The report characterizes types of ESA-group interactions, introduces a cognitive behavioral architecture for robots to act in groups, and highlights group dimensions relevant for the successful integration of ESAs into contexts. Based on the assertion that any social interaction is a group social interaction to some degree, we claim that our recommendations have relevance not only for the design of virtual agents and robots that are deployed in immediate group and team contexts but for HRI in general.

References

1 Ronald C Arkin. *Behavior-based robotics*. MIT Press, 1998.
2 Reuben M. Aronson and Henny Admoni. Semantic gaze labeling for human-robot shared manipulation. In *Proceedings of the 11th ACM Symposium on Eye Tracking Research & Applications*, ETRA '19, pages 2:1–2:9. ACM, 2019.
3 Marilynn B Brewer. The psychology of prejudice: Ingroup love and outgroup hate? *Journal of social issues*, 55(3):429–444, 1999.
4 P. Brown and S. Levinson. *Politeness. Some Universals in Language Usage*. Cambridge University Press, Cambridge, 1987.
5 Rupert Brown and Samuel Pehrson. *Group processes: Dynamics within and between groups*. John Wiley & Sons, 2019.
6 Merijn Bruijnes, Rieks op den Akker, Arno Hartholt, and Dirk Heylen. Virtual suspect william. In *International Conference on Intelligent Virtual Agents*, pages 67–76. Springer, 2015.
7 Susan G Cohen and Diane E Bailey. What makes teams work: Group effectiveness research from the shop floor to the executive suite. *Journal of Management*, 23(3):239–290, 1997.
8 Filipa Correia, Samuel Mascarenhas, Rui Prada, Francisco S Melo, and Ana Paiva. Group-based emotions in teams of humans and robots. In *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*, pages 261–269. ACM, 2018.
9 James P Dillard, John E Hunter, and Michael Burgoon. Sequential-request persuasive strategies: Meta-analysis of foot-in-the-door and door-in-the-face. *Human Communication Research*, 10(4):461–488, 1984.
10 Anca Dragan and Siddhartha Srinivasa. Integrating human observer inferences into robot motion planning. *Autonomous Robots*, 37(4):351–368, 2014.
11 Mica R Endsley and Daniel J Garland. *Situation awareness analysis and measurement*. CRC Press, 2000.
12 Franck Feurtey. Simulating the collision avoidance behavior of pedestrians. Master's thesis, University of Tokyo, Tokyo, Japan, 2000.
13 William W Gaver. Technology affordances. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 79–84. ACM, 1991.
14 Michele J Gelfand, Jesse R Harrington, and Joshua Conrad Jackson. The strength of social norms across human groups. *Perspectives on Psychological Science*, 12(5):800–809, 2017.
15 James J Gibson. The theory of affordances. *Hilldale, USA*, 1(2), 1977.
16 Erving Goffman. *Interaction ritual: Essays on face-to-face interaction*. Aldine, 1967.
17 Amit Goldenberg, Eran Halperin, Martijn van Zomeren, and James J Gross. The process model of group-based emotion: Integrating intergroup emotion and emotion regulation perspectives. *Personality and Social Psychology Review*, 20(2):118–141, 2016.
18 Michael A Goodrich and Alan C Schultz. Human-robot interaction: a survey. *Foundations and Trends in Human-Computer Interaction*, 1(3):203–275, 2007.
19 Victoria Groom and Clifford Nass. Can robots be teammates?: Benchmarks in human-robot teams. *Interaction Studies*, 8(3):483–500, 2007.

20 Tariq Iqbal, Samantha Rack, and Laurel D Riek. Movement coordination in human-robot teams: a dynamical systems approach. *IEEE Transactions on Robotics*, 32(4):909–919, 2016.

21 Tariq Iqbal and Laurel D Riek. A method for automatic detection of psychomotor entrainment. *IEEE Transactions on Affective Computing*, 7(1):3–16, 2015.

22 Tariq Iqbal and Laurel D Riek. Coordination dynamics in multihuman multirobot teams. *IEEE Robotics and Automation Letters*, 2(3):1712–1717, 2017.

23 Tariq Iqbal and Laurel D Riek. Human-robot teaming: Approaches from joint action and dynamical systems. *Humanoid Robotics: A Reference*, pages 2293–2312, 2019.

24 Shervin Javdani, Henny Admoni, Stefania Pellegrinelli, Siddhartha S. Srinivasa, and J. Andrew Bagnell. Shared autonomy via hindsight optimization for teleoperation and teaming. *The International Journal of Robotics Research*, 37(7):717–742, 2018.

25 Malte F Jung. Coupling interactions and performance: Predicting team performance from thin slices of conflict. *ACM Transactions on Computer-Human Interaction*, 23(3):18, 2016.

26 Malte F Jung. Affective grounding in human-robot interaction. In *2017 12th ACM/IEEE International Conference on Human-Robot Interaction*, pages 263–273. IEEE, 2017.

27 Malte F Jung, Nikolas Martelaro, and Pamela J Hinds. Using robots to moderate team conflict: the case of repairing violations. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction*, pages 229–236. ACM, 2015.

28 Reshmashree B Kantharaju, Alison Pease, Dennis Reidsma, Catherine Pelachaud, Mark Snaith, Merijn Bruijnes, Randy Klaassen, Tessa Beinema, Gerwin Huizing, Donatella Simonetti, et al. Integrating argumentation with social conversation between multiple virtual coaches. In *Proceedings of the 19th ACM International Conference on Intelligent Virtual Agents*, pages 203–205. ACM, 2019.

29 Ross A. Knepper, Christoforos Mavrogiannis, Julia Proft, and Claire Liang. Implicit communication in a joint action. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction*, pages 283–292, 2017.

30 Steve WJ Kozlowski and Bradford S Bell. Work groups and teams in organizations. *Handbook of Psychology*, pages 333–375, 2003.

31 Minae Kwon, Mengxi Li, Alexandre Bucquet, and Dorsa Sadigh. Influencing leading and following in human-robot teams. *Proceedings of Robotics: Science and Systems, FreiburgimBreisgau, Germany*, 2019.

32 Min Kyung Lee, Sara Kiesler, Jodi Forlizzi, Siddhartha Srinivasa, and Paul Rybski. Gracefully mitigating breakdowns in robotic services. In *2010 5th ACM/IEEE International Conference on Human-Robot Interaction*, pages 203–210. IEEE, 2010.

33 Bertram F Malle, Matthias Scheutz, and Joseph L Austerweil. Networks of social and moral norms in human and robot agents. In *A World with Robots*, pages 3–17. Springer, 2017.

34 Christoforos I Mavrogiannis and Ross A Knepper. Multi-agent path topology in support of socially competent navigation planning. *The International Journal of Robotics Research*, 38(2-3):338–356, 2019.

35 Christoforos I. Mavrogiannis, Wil B. Thomason, and Ross A. Knepper. Social momentum: A framework for legible navigation in dynamic multi-agent environments. In *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*, pages 361–369. ACM, 2018.

36 D. Harrison McKnight and Norman L. Chervany. What is trust? A conceptual analysis and an interdisciplinary model. *AMCIS 2000 Proceedings*, page 382, 2000.

37 Bilge Mutlu and Jodi Forlizzi. Robots in organizations: the role of workflow, social, and environmental factors in human-robot interaction. In *Proceedings of the 3rd ACM/IEEE International Conference on Human Robot Interaction*, pages 287–294. ACM, 2008.

38 Aastha Nigam and Laurel D Riek. Social context perception for mobile robots. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 3621–3627. IEEE, 2015.

39 Stefanos Nikolaidis, David Hsu, and Siddhartha S. Srinivasa. Human-robot mutual adaptation in collaborative tasks: Models and experiments. *The International Journal of Robotics Research*, 36(5-7):618–634, 2017.

40 Donald A Norman. Affordance, conventions, and design. *interactions*, 6(3):38–43, 1999.

41 Maria Francesca O'Connor and Laurel D Riek. Detecting social context: A method for social event classification using naturalistic multimodal data. In *2015 11th IEEE International Conference and Workshops on Automatic Face and Gesture Recognition*, volume 3, pages 1–7. IEEE, 2015.

42 Harm op den Akker, Rieks op den Akker, Tessa Beinema, Oresti Banos, Dirk Heylen, Björn Bedsted, Alison Pease, Catherine Pelachaud, Vicente Traver Salcedo, Sofoklis Kyriazakos, et al. Council of coaches a novel holistic behavior change coaching approach. In *4th International Conference on Information and Communication Technologies for Ageing Well and e-Health, ICT4AWE 2018*, pages 219–226. SciTePress, 2018.

43 Steffi Paepcke and Leila Takayama. Judging a bot by its cover: an experiment on expectation setting for personal robots. In *2010 5th ACM/IEEE International Conference on Human-Robot Interaction*, pages 45–52. IEEE, 2010.

44 Raja Parasuraman, Thomas B Sheridan, and Christopher D Wickens. A model for types and levels of human interaction with automation. *IEEE Transactions on Systems, Man, and Cybernetics, Part A: Systems and Humans*, 30(3):286–297, 2000.

45 Mina Park, Jae-Jin Song, Seo Jin Oh, Min-Sup Shin, Jun Ho Lee, and Seung Ha Oh. The relation between nonverbal IQ and postoperative CI outcomes in cochlear implant users: Preliminary result. *BioMed Research International*, 2015, 2015.

46 Jens Rasmussen. Skills, rules, and knowledge; signals, signs, and symbols, and other distinctions in human performance models. *IEEE Transactions on Systems, Man, and Cybernetics*, (3):257–266, 1983.

47 Laurel D Riek, Philip C Paul, and Peter Robinson. When my robot smiles at me: Enabling human-robot rapport via real-time head gesture mimicry. *Journal on Multimodal User Interfaces*, 3(1-2):99–108, 2010.

48 L.D. Riek. The Social Co-Robotics Problem Space: Six Key Challenges. *In Proceedings of Robotics: Science, and Systems, Robotics Challenges and Visions*, 2013.

49 Dorsa Sadigh, S Shankar Sastry, Sanjit A Seshia, and Anca Dragan. Information gathering actions over human internal state. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 66–73. IEEE, 2016.

50 Nicole Salomons, Michael van der Linden, Sarah Strohkorb Sebo, and Brian Scassellati. Humans conform to robots: Disambiguating trust, truth, and conformity. In *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*, pages 187–195. ACM, 2018.

51 Vasanth Sarathy and Matthias Scheutz. A logic-based computational framework for inferring cognitive affordances. *IEEE Transactions on Cognitive and Developmental Systems*, 10(1):26–43, 2016.

52 Matthias Scheutz, Paul Schermerhorn, James Kramer, and Christopher Middendorff. The utility of affect expression in natural language interactions in joint human-robot tasks. In *Proceedings of the 1st ACM SIGCHI/SIGART Conference on Human-Robot Interaction*, volume 2, pages 226–233, 2006.

53 J Scholtz. Theory and evaluation of human robot interactions. In *Proceedings of the 36th Annual Hawaii International Conference on System Sciences*, page 10 pp. IEEE, 2003.

54 Sarah Strohkorb Sebo, Priyanka Krishnamurthi, and Brian Scassellati. "I don't believe you" investigating the effects of robot trust violation and repair. In *2019 14th ACM/IEEE International Conference on Human-Robot Interaction*, pages 57–65. IEEE, 2019.

55 Burr Settles. Active learning literature survey. Technical report, University of Wisconsin-Madison Department of Computer Sciences, 2009.

56 Solace Shen, Petr Slovak, and Malte F Jung. Stop. I see a conflict happening: A robot mediator for young children's interpersonal conflict resolution. In *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*, pages 69–77. ACM, 2018.

57 Elaine Short and Maja J Mataric. Robot moderation of a collaborative game: Towards socially assistive robotics in group interactions. In *2017 26th IEEE International Symposium on Robot and Human Interactive Communication*, pages 385–390. IEEE, 2017.

58 Tianmin Shu, Michael S Ryoo, and Song-Chun Zhu. Learning social affordance for human-robot interaction. *arXiv preprint arXiv:1604.03692*, 2016.

59 Annika Silvervarg and Kristian Månsson. How do you introduce an agent? The effect of introduction type on how a teachable agent is experienced by students. In *Proceedings of the 18th International Conference on Intelligent Virtual Agents*, pages 29–34. ACM, 2018.

60 Kristen Stubbs, Pamela J Hinds, and David Wettergreen. Autonomy and common ground in human-robot interaction: A field study. *IEEE Intelligent Systems*, 22(2):42–50, 2007.

61 Hiroki Tanaka, Hideki Negoro, Hidemi Iwasaka, and Satoshi Nakamura. Embodied conversational agents for multimodal automated social skills training in people with autism spectrum disorders. *PloS one*, 12(8):e0182151, 2017.

62 Hamish Tennent, Solace Shen, and Malte Jung. Micbot: A peripheral robotic object to shape conversational dynamics and team performance. In *2019 14th ACM/IEEE International Conference on Human-Robot Interaction*, pages 133–142. IEEE, 2019.

63 Edwin J Thomas and Clinton F Fink. Effects of group size. *Psychological bulletin*, 60(4):371, 1963.

64 Marynel Vázquez, Elizabeth J Carter, Braden McDorman, Jodi Forlizzi, Aaron Steinfeld, and Scott E Hudson. Towards robot autonomy in group conversations: Understanding the effects of body orientation and gaze. In *Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction*, pages 42–52. ACM, 2017.

65 Marynel Vázquez, Alexander May, Aaron Steinfeld, and Wei-Hsuan Chen. A deceptive robot referee in a multiplayer gaming environment. In *2011 International Conference on Collaboration Technologies and Systems (CTS)*, pages 204–211. IEEE, 2011.

66 Susan A Wheelan. Group size, group development, and group productivity. *Small group research*, 40(2):247–262, 2009.

67 Michele Williams. In whom we trust: Group membership as an affective context for trust development. *Academy of Management Review*, 26(3):377–396, 2001.

4.5 Working Group on Scenarios for Human-Agent and Human-Robot Groups

Patrícia Alves-Oliveira (ISCTE-IUL & INESC-ID, PT & Cornell University, US, patricia_alves_oliveira@iscte-iul.pt)
Mary Ellen Foster (University of Glasgow, GB, MaryEllen.Foster@glasgow.ac.uk)
Hatice Gunes (University of Cambridge, GB, hatice.gunes@cl.cam.ac.uk)
Sviatlana Höhn (University of Luxembourg, LU, sviatlana.hoehn@uni.lu)
Franziska Kirstein (Blue Ocean Robotics, DK, fk@blue-ocean-robotics.com)
Andrea Marrella (Sapienza University of Rome, IT, marrella@diag.uniroma1.it)
Mark Neerincx (TU Delft & TNO, NL, mark.neerincx@tno.nl)
Michio Okada (Toyohashi University of Technology, JP, okada@tut.jp)
Ana Paiva (INESC-ID & University of Lisbon, PT, ana.paiva@inesc-id.pt)
Elizabeth Sklar (University of Lincoln, GB, elizabeth.sklar@kcl.ac.uk)
Hannes Högni Vilhjálmsson (Reykjavik University, IS, hannes@ru.is)

4.5.1 Introduction

As social agents and robots become more prominent in our lives, they also increasingly become members of groups and teams, and new scenarios emerge. This document reflects upon the different types of scenarios that are being built and investigated where humans and agents act in groups. It provides a framework for scenarios analysis looking at different dimensions that affect the way the scenarios are conceived, analysed and evaluated.

4.5.2 The ENACTED Framework for scenarios analysis

The *Environment Norms Autonomy Composition Task Embodiment Duration (ENACTED)* framework is designed to provide an organised structure around the definition of scenarios for human-agent and/or human-robot group interaction.

Environment and Domain

The environment conditions the way the groups interact and collaborate. Different environments with the same agents, robots and task may shape and condition the way the group dynamics unfolds. For example, the robots may share the same physical environment as the humans (i.e., *proximal* interaction, such as a robot collecting/delivering items from/to a human in an office, on a farm or in a warehouse); the interaction may involve physically touching each other (e.g., robot lifting a human in a healthcare or rescue operation); the actors may be physically separate from each other (i.e., *remote* interaction, such as a robot surveilling the structure of a nuclear power plant); the "robot" may not have a physical body (i.e., *chatbot*, such as a helper on a web site); the "robot" may be embodied in a virtual platform (i.e., *virtual agent*, such as an agent demonstrating how to conduct a physical task). In situations where the group members are not all physically co-located, there might be multiple settings defined which together comprise the scenario's environment. For example, in the search-and-rescue scenario, you might have a human first-responder who is located in a van outside of a collapsed building while you have a group of robots inside the building searching for victims and providing sensor data to the human team member.

Table 2 Versions of autonomy and delegation.

	Human decides	Robot decides
Human acts	100% human, no robots	Robot transfers a task to a human, e.g when the robot reaches its limits.
Robot acts	Human operator determines what the robot needs to do and when.	100% robot's autonomy, no human needed, but can be included as a controlling instance.

Norms, Social Interactions and Culture

Groups follow certain social norms that shape the way their members interact. In fact, particular groups can even create new norms as the group interactions unfold. Some questions to consider in the scenarios include:

- How are norms defined for behaviour?
- Are norms associated with particular roles/tasks/responsibilities?
- What is the alignment between normative behaviour and expected behaviour when agents and humans act in a group setting?
- Are there actions that are within the norms but are unexpected?

Often, it is considered that the norms for social interaction human groups are also conditioned by cultural learning[7]. What happens if the actions of the agents in a group fall outside what is acceptable?

Degree of Autonomy

There are scenarios in which a robot handles completely autonomously (e.g. killer drone), in others it accomplishes particular tasks that a human explicitly asked it to do (arrange a meeting with a project partner and send a calendar invite), or a robot decides to transfer a task to a human when it reaches its limits (transfer a call to an operator if a chat robot cannot satisfy the caller's request).

We can represent the dimension of autonomy as a combination of two features: autonomy of decisions and autonomy of actions. Table 2 summarises their combinations.

Group Composition

When considering group composition, we need to look at the concepts of outgroup and ingroup first. An outgroup is any group that the entity perceives as not belonging to, while an ingroup is a group that the entity psychologically identifies as being a member of.

Number of entities: The number of entities in the group directly affect the composition of the group. Differently from dyadic interactions, the number of entities in the group need to be considered in relation to other compositional aspects, as well as the other factors included in the ACTED Framework.

Questions that need to be addressed include the following:

- When going from three to four to many number of entities (n), at what point (n=?) the group is formed and the dynamics do not change if one entity leaves/changes or another entity joins the group?

[7] https://journals.sagepub.com/doi/full/10.5772/57260

- How do the dynamics of the group change when we have one vs many agents/robots w.r.t one vs many humans? E.g., One drone approaching a human vs many drones? How does the perception differ if they all look different vs. they all look the same?

Roles: The notion of roles relates to top-down vs. bottom-role assignment or acquisition of a role in a group or team setting. There will be notable differences in assigned team roles compared to acquired team roles. Assigned role is likely to start with a static role (e.g., in a group-based card game setting) and may or may not change (to a different role depending on the group interaction and the dynamics). Instead, acquired role is expected to be formed through the actual interaction taking place. Additionally, we need to consider the roles of the robot and the humans in terms of augmentation or replacement of human abilities.

Questions that need to be addressed here include the following:

- What are the responsibilities that the roles bring along?
- With dynamic roles, what is the means of change (changing roles)? Does change evolve/emerge as the group interacts? Or is there some external factor or norm that dictates the change?
- Does "change" refer to different actors (human, robot) switching between role assignments? Or could it also refer to how the norms/behaviours/responsibilities/actions of the role changes (i.e., the assignment of named role to actor does not change but the actions associated with that role change)?

Homogeneity: The homogeneity vs. heterogeneity of the entities is another factor that directly affects the composition of the group. This in itself is multi-dimensional and would consider various aspects that are used in defining ingroup versus outgroup membership. Examples include not only the commonly used gender (all-female groups, all-robot groups, etc.) and ethnicity (all-white groups vs. all-Asian groups) criteria but also other aspects such as colour (the colour of the robot, the colour of the shirt worn by the team members to indicate belonging), rhythm (how fast/slow one talks, walks, interacts etc.) or other categorisation aspects such as personality (a group consisting of extroverts only vs. a group consisting of both extroverts and introverts). Therefore, the perception of homogeneity or heterogeneity of a group or team can be easily modified using such criteria (e.g. same/similar colours, sounds, movements – fast/slow, rhythm). Stereotypes usually emerge from the tendency to see members of an outgroup as similar (outgroup homogeneity) and members of one's ingroup as different from each other (ingroup heterogeneity).

Questions that need to be addressed include the following:

- How is the composition of the team defined in terms of homogeneity vs. heterogeneity? Does it change? Is it flexible?
- Which factors define hetero/homogeneity beyond roles? I.e., category (robot vs. human), appearance (e.g., orange robot vs. blue robot), behaviour/personality, acceptance (making the human feel in-group or out-group).

Task

One of the dimensions of analysis is the type of task and the environment where the task is being carried out. The task, is related with the type of problem the group is solving: What is the problem that you are trying to solve? Are you addressing a human need? Are you answering fundamental question about human interaction? Are you addressing an identified business need? Are you pursuing a common goal or an individual goal? Are you pursuing a social or task-based goal or a mixture of both?

Types of tasks:

- Informative Tasks (e.g., robot as museum guide or shopping mall guide)
- Competitive Tasks (e.g., robot playing a game with the purpose of winning)
- Collaborative Tasks (e.g., working together to solve a problem or teach a language or complete an action in a factory)
- Action-oriented Tasks (e.g., robot as a butler, or robots in factories)
- Creative Tasks (e.g., playful interactions by acting and reacting)
- Planning Tasks (e.g.,)
- Mixed-motive Tasks (e.g.,)
- Judgemental Tasks (e.g., a jury making a judgement of good vs. bad without acting)

Embodiment and interaction Resources

A presence of a physical robotic body in a group does not automatically mean that this body uses its all capacities in interaction with other group members. In this section we describe the identified types of embodiment and the types of interaction resources that can be associated or made available with the given type of embodiment. For instance, users can talk to a Pepper robot, but it is also possible to use WhatsApp to interact with it.

From chatbots to humanoid robots: Chatbots (text or voice-based) usually do not have a special "body", their interface is usually tailored to process input/output signals. Therefore, for text-based chatbots usually "live" in messengers or webchat, and voice-based chatbots need a microphone for user input and a speaker for output. However, a humanoid robot such as NAO or Pepper can be also used as interfaces for other services, for which the full embodiments is actually not needed.

Interactional Resources: The choice of a particular type of embodiment does not necessarily prescribe to use a specific communication channel. Sometimes it is even desired for the purpose of a particular research work to change or to limit the use of particular communication channel. Therefore, we see it as a separate feature.

Relative position: The relative position of the robot(s) is determined by their size and task. We intuitively found the following four but other may exist:

1. Human is inside a robot (autonomous car, autonomous space shuttle, smart-home);
2. Robot shares space with a human (humanoid robot in a lab, robotic arms in production halls and similar);
3. Robot inside a human (smart medical devices, artificial eye etc.).

Duration of the interaction

We distinguish between long-term and short-term interaction. Categories such as episodic and continuous interaction are included into sort and long-term scenarios.

Short-term interaction: Short-term interaction may include one or multiple *independent* sessions to accomplish short, usually well-defined tasks. Example: robotic information desk in a shopping mall, robot in a hospital lifting a patient.

Even a large number of short subsequent interaction can is classified a short-term as long as these interactions do not form one process or are not parts of a more complex task. As soon as sequential dependencies between subsequent tasks occur, we speak about long-term interaction.

	1: MuMMER	2: Patient Lifting Robot	3: Strawberry Harvesting Robot	4: Promoting Creativity with Robots	5: PAL	6: Gaming / Entertainm. / Science Comm	7: Robots for Emergency Management	8: Social Training	9: Text a Robot
AUTONOMY	Full (respond to approach, initiates)	Semi (controlled by nurse)	Full (Collab. with "farmer-in-the-loop" to learn)	Full (collaborate with children)	Full	Full (Detect gestures and interpret who wins)	Full (Reactive to orders, proactive in task)	Full (react to turn-taking, initiate turns)	Full
COMPOSITION	One robot, many humans, dynamic	One robot, two humans (nurse, patient)	Multiple robots, one human (farmer)	One or more robots, one or more humans (children)	One robot + avatar, many humans (patients, parents, HCP)	One robot, one to five humans (children)	Many robots, many humans. Hetero-geneous	3-5 agents, one human	One robot, many humans
TASK	Guidance (small talk in shopping mall)	Patient-lifting (in hospital)	Strawberry identification	Story creation (creative, but goal-oriented task)	Child diabetes self-managem ent (learning and behavioral change)	Game (playing rock-paper scissor)	Operation in dangerous environments (evaluate damage,remo ve debris)	Improvement of social skills (turn-taking); (human joins group of agents)	Guidance of people, patrolling building
EMBODIMENT	Physical humanoid (Pepper)	Physical non-humanoid (custom robot platform)	Physical non-humanoid (robot platform in poly tunnel)	Physical non-humanoid (toys)	Physical and virtual humanoid (both Nao)	Physical humanoid (Nao)	Physical non-humanoid (tracked wheels and digger)	Virtual humanoid (human-sized)	Physical humanoid + text-based chat
DURATION	Long series, short inter. (<3 min)	Infinite series, short inter. (throughout day)	Long series, cont. inter. changes in composition	Short interaction	Long series (6 mon), short inter. daily (1 - 60 min)	On day series (7 hours), one off short inter. (5 mins)	Infinite series, short interaction	Course series (3 weeks), short inter. (<15 min)	Long series, short interactions (several mins)

Figure 8 Exercise in applying the a chosen set of 5 fairly well defined dimensions to 9 different social group scenarios.

Long-term interaction: Long-term interaction includes multiple sessions with sequential order and sequential dependencies. The order and dependencies can be defined through:
- Information gathered in earlier sessions that is used in later sessions (e.g. names, daily duties of team members, preferences of team members...);
- Level of social proximity (politeness, intimacy) expressed through use of particular interactional resources (invitations, thank-yous and effort for interaction management);
- Conflict and relationship management (e.g. the need to interact for many weeks and a conflict identified in week 1).

4.5.3 Scenarios Evaluation

The scenarios evaluation is available on Figure 8.

4.6 Working Group on Social Behaviours for Group Interactions between Humans and Social Agents and Robots

Elisabeth André (Universität Augsburg, DE, andre@informatik.uni-augsburg.de)
Dirk Heylen (University of Twente, NL, d.k.j.heylen@utwente.nl)
James Kennedy (Disney Research, US, james.r.kennedy@disney.com)
Ana Paiva (INESC-ID & University of Lisbon, PT, ana.paiva@inesc-id.pt)
Catherine Pelachaud (Sorbonne University, FR, catherine.pelachaud@upmc.fr)
Hannes Högni Vilhjálmsson (Reykjavik University, IS, hannes@ru.is)

4.6.1 Introduction and Definitions

When considering social agents, it is first important to have a collective understanding of what is being studied. Breazeal et al. [2] define a social robot as a "robot designed to interact with people in a natural, interpersonal manner". Accordingly, the following criteria for a social agent was established for this working group: "a social agent must have been designed as such". That is to say, the agent must have capabilities for social interaction. An object can be part of a social group, yet it may not be a social agent. Therefore, a vacuum cleaning robot would not be considered social as it has not been explicitly designed with social interaction in mind. However, if the same robot were equipped to understand human conversation and avoid vacuuming when a conversation is taking place, it would then be considered social.

Social agents must communicate using verbal and nonverbal social signals. Such signals include facial expressions, posture, vocal prosody, or spoken language. We then define social behaviours as a higher-level task, like handing over an object or acting towards a goal. A social behaviour can be produced from the combination of multiple social signals. When a group of agents starts to have the same goal, and agrees to this goal, they become a team. The outcome of the actions of the individuals affects all members of the team. In this context, there is a distinction between the social behaviours of the individuals and their behaviours toward the task goal.

4.6.2 Open Problems & Current Challenges

Many research papers have been published that explore social signal and social behaviour analysis between humans (e.g., [3]), and between humans and agents (e.g., [5, 7]). Figure 9, based on [6], provides an overview of the complexity that exists when considering behaviours in groups. Social behaviours form part of this, creating many challenges in considering these processes during design and evaluation of social agents in group settings.

From our understanding of the field gleaned through experience and published works, we identify the open problems and challenges listed below. In particular, we focus on issues for groups, where groups consist of more than two agents (whether human or artificial).

- **Awareness** of the group is important. Which cues signal awareness? Considering the group as an entity, do social behaviours depend on the group identity?
- **Timing** of the cues is extremely important in social behaviours. Sometimes it is not so much about the behaviour itself but *when* to do it. It does not matter what you do (to a certain extent!) provided that you do it at the right moment.

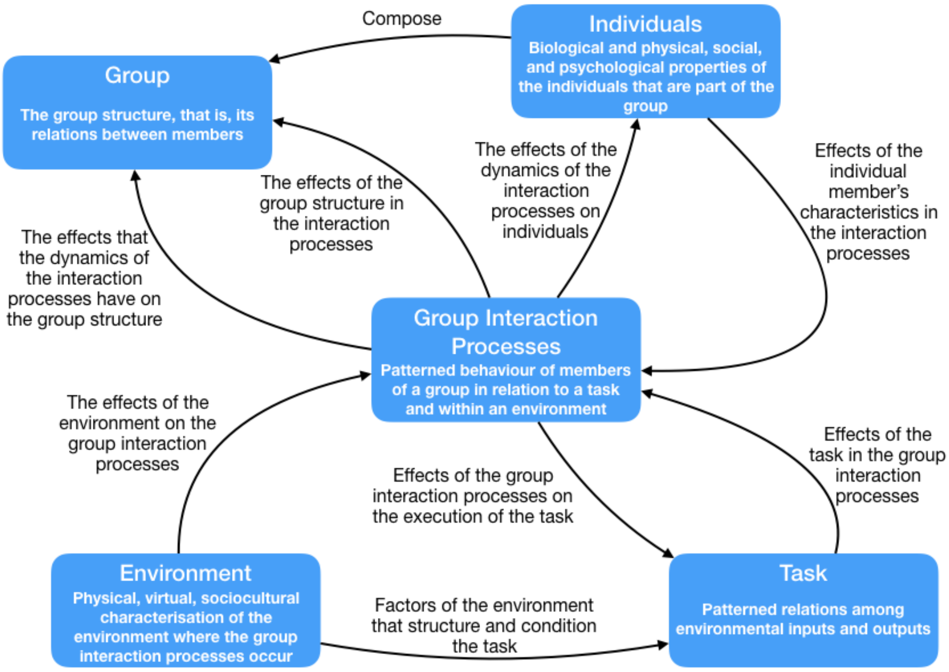

Figure 9 Group Processes based on [6].

- **Social Roles**. Multiple simultaneous roles are possible within a group. Social group hierarchy has an influence on the role.
- **Type of Task** such as creative, competitive, etc. Tasks can be broken into segments, and each has a 'Purpose'. Changing the 'purpose' might also change the social roles.
- **Types of Groups and Context** – Groups are placed in a formation in the world (physical or virtual), but members of the group may additional group structure and membership, like for example a family. A "family" is a group but may not be placed physically in the same F-formation in a given scenario. The influence of this group may still be present in an interaction; how do we model this?
- **Adaptation** – How social behaviours are used and people adapt the social behaviours. How should agents adapt in a group: to the group as a whole, or to individual members of the group? For example, if a robot is a speaker to a class, the members of the class can be considered as a whole. It is not the individual members that matter but rather the whole group. Understanding how to adapt in a way that the identity (personality, character, etc) of the agent is maintained is unclear.
- **Cohesion of a Group** – task or social? How to make a group cohesive? And what is a difference with a team?
- **Emotional Behaviours** affect the team cohesion and performance. Balancing both the positive and negative emotions of individuals is important.
- **Stereotypes** should the behaviours of agents and robots follow and promote certain stereotypes. How do we evaluate when social behaviours may be aligning to stereotypes, and how might we avoid this?
- **Transparency** of the social behaviours. If an agent is being used to influence people, should it explain the behaviours it is using to do so? If it did, how would this affect the influence it has? If the influence is reduced, this may be a negative consequence for agents that aim to positively change people's lives (e.g., a weight loss coach [4]).

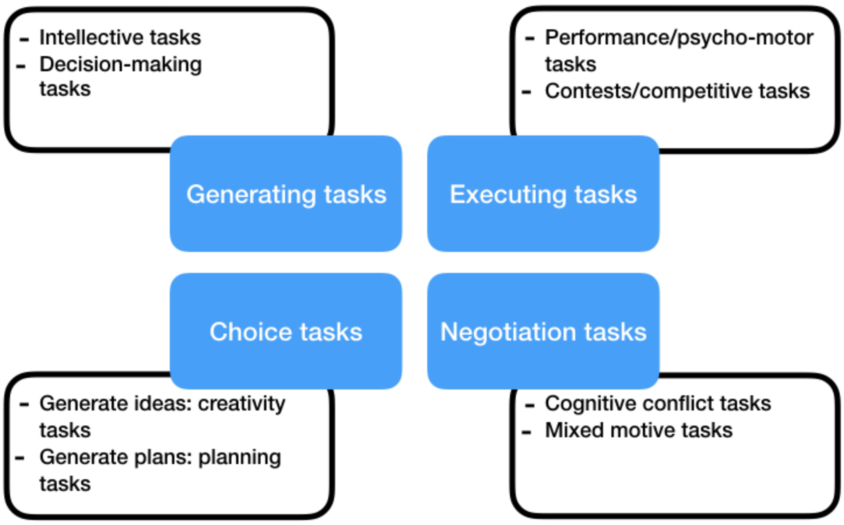

Figure 10 Types of tasks in group interactions based on [6].

4.6.3 Methodologies for Studying Groups

Social skills should enhance interactions with people, to improve the group. Finding the right set of social abilities should be tied to the function that the agent has in the group; we should not be designing just for the sake of being social. Group interactions can be classified differently according to task on demand, see for instance Figure 10. Such frameworks may guide the design of social behaviours for agents in group settings.

Many annotation schemes currently exist for dyadic interactions, i.e., interactions between two agents. These have been used extensively in research to understand social signals and social behaviours between people. This understanding is often transferred to our design of social agents, where the annotation schemes can subsequently be used to evaluate the behaviour of the agent in interactions. However, many complexities arise in applying these same schemes to groups, as any factors that encode behaviours in relation to other group members will increase exponentially with the addition of group members. This is problematic when annotating data manually due to the significant impact on the amount of resources required. Automated analysis as seen in [1] may provide part of the answer, but exploration of alternative schemes is also needed. Many existing schemes either do not capture concepts that arise when studying groups, or are limited in their coding scheme for labelling certain phenomena (for example, synchrony and mimicry in a group is difficult to capture outside of pairwise labelling between individuals). We see the development of annotation schemes and metrics specifically for groups as a key step toward greater understanding of group social behaviour, potentially leading to improved agent interaction capabilities.

References

1 Tobias Baur, Gregor Mehlmann, Ionut Damian, Florian Lingenfelser, Johannes Wagner, Birgit Lugrin, Elisabeth André, and Patrick Gebhard. Context-Aware Automated Analysis and Annotation of Social Human–Agent Interactions. ACM Transactions on Interactive Intelligent Systems (TiiS) 5(2):1–31. ACM, 2015.
2 Cynthia Breazeal, Kerstin Dautenhahn, and Takayuki Kanda. Social Robotics. In *Springer Handbook of Robotics*, pages 1935–1972. Springer, 2016.

3 Kevin El Haddad, Sandeep Nallan Chakravarthula, and James Kennedy. Smile and Laugh Dynamics in Naturalistic Dyadic Interactions: Intensity Levels, Sequences and Roles. In *Proceedings of the International Conference on Multimodal Interaction*. ACM, 2019.

4 Cory D. Kidd, and Cynthia Breazeal. A robotic weight loss coach. In *Proceedings of the National Conference on Artificial Intelligence*, 22(2):1985. AAAI Press, MIT Press, 2007.

5 Iolanda Leite, Carlos Martinho, and Ana Paiva. Social robots for long-term interaction: a survey. International Journal of Social Robotics 5(2):291–308. Springer, 2013.

6 Joseph Edward McGrath. Groups: Interaction and Performance. Vol. 14. Prentice-Hall, 1984.

7 Alessandro Vinciarelli, Maja Pantic, Dirk Heylen, Catherine Pelachaud, Isabella Poggi, Francesca D'Errico, and Marc Schroeder. Bridging the gap between social animal and unsocial machine: A survey of social signal processing. IEEE Transactions on Affective Computing 3(1):69–87. IEEE, 2011.

Participants

- Patrícia Alves-Oliveira
ISCTE-IUl, INESC-ID – Lisboa, PT & Cornell University – Ithaca, US
- Elisabeth André
Universität Augsburg, DE
- Tony Belpaeme
Ghent University, BE
- Cindy Bethel
Mississippi State University, US
- Merijn Bruijnes
Delft University of Technology, NL
- Filipa Correia
INESC-ID – Porto Salvo & University of Lisbon, PT
- Mary Ellen Foster
University of Glasgow, GB
- Kobi Gal
University of Edinburgh, GB
- Hatice Gunes
University of Cambridge, GB
- Dirk Heylen
University of Twente, NL
- Sviatlana Höhn
University of Luxembourg, LU
- Malte Jung
Cornell University – Ithaca, US
- James Kennedy
Disney Research – Glendale, US
- Franziska Kirstein
Blue Ocean Robotics – Odense, DK
- Stefan Kopp
Universität Bielefeld, DE
- Iolanda Leite
KTH Royal Institute of Technology – Stockholm, SE
- Andrea Marrella
Sapienza University of Rome, IT
- Samuel Mascarenhas
IST – Lisbon, PT
- Christoforos Mavrogiannis
University of Washington – Seattle, US
- Mark Neerincx
TU Delft, NL
- Catharine Oertel
TU Delft, NL
- Michio Okada
Toyohashi University of Technology, JP
- Ana Paiva
INESC-ID – Porto Salvo & University of Lisbon, PT
- Simon Parsons
University of Lincoln, GB
- Catherine Pelachaud
Sorbonne University – Paris, FR
- André Tiago Abelho Pereira
KTH Royal Institute of Technology – Stockholm, SE
- Rui Prada
INESC-ID – Porto Salvo & University of Lisbon, PT
- Laurel Riek
University of California – San Diego, US
- Sarah Sebo
Yale University – New Haven, US
- Julie Shah
MIT – Cambridge, US
- Elaine Short
Tufts University – Medford, US
- Elizabeth Sklar
University of Lincoln, GB
- Marynel Vázquez
Yale University – New Haven, US
- Hannes Högni Vilhjálmsson
Reykjavik University, IS