Learning to Act Socially

Matthias Scheutz
Human Robot Interaction Laboratory
Tufts University
Medford, MA 02155, USA
http://hrilab.tufts.edu/

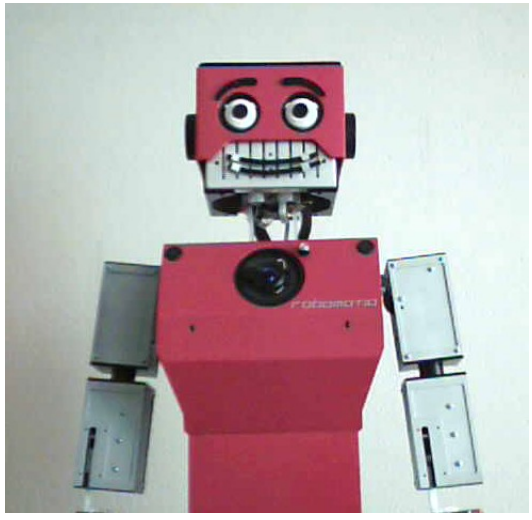

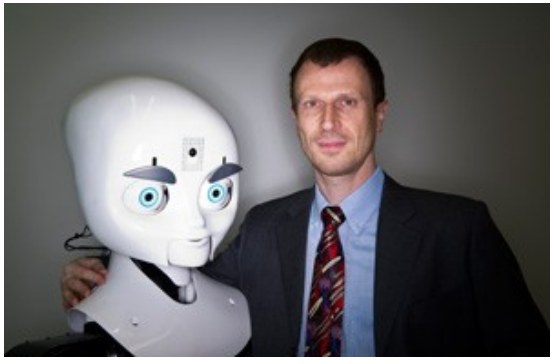

“Social robots” – an oxymoron?

- Current robots are **intrinsically asocial**, even so-called “social robots” that supposedly “connect with us on a personal level” (e.g., Breazeal 2002, but many others)
- For robots don’t know much about humans (other than that they might be obstacles in the case of vacuum cleaners), or about what humans are thinking (even when they are able to recognize faces and respond to sounds with different facial expressions as Kismet already did in the early 2000s)
- Little has changed since then in terms of making robots **more social** even though some got better at exhibiting and responding to social cues such as eye gaze behavior
- Yet, for a robot to truly become social, it needs to have an **understanding of what a social other is**, not just respond to social cues with predefined (or possibly learned) behavioral patterns that are ultimately meaningless
- This is a high bar for robots and we are not even close yet!

What do robots need to know to become social?

- For starters, to be able to start to understand a social other, the robot has to have a notion of what a social other is
- This requires it to have a sense of what an agent is, that agents have desires, attitudes, goals, intentions, etc.
- For example, for a robot to understand why a particular way of talking is “face-threatening” to a person and that it, as a result, should be polite, it first has to have an understanding of what “face” means, namely a person’s “public self-image” (how others perceive the person)
- This, in turn, requires the robot to understand what it means to perceive a person from a person’s perspective (not the robot’s) and that self-esteem (positive face) and freedom to act (negative face), are people’s fundamental values in any social interaction, requiring cooperation among interactants to maintain each other's (positive and negative) face
- How should a robot know that?

If they can't be social yet, maybe they can act socially?

- While I cannot offer a detailed list of requirements for robots to understand public face and to perform the undergirding social interactions – simply because I do not know all that is required – here are at least a few key requirements:
- Building and maintaining mental models of interlocutors; mental state inference, intent detection and plan abduction; operating with possibly conflicting social and moral norms, habits and conventions, attitudes, and etiquette; employing common sense knowledge and reasoning; demonstrating emotional intelligence; understanding purposes and reasons for actions and action outcomes; understanding human desires, goals, and feelings together with their effects
- So: if we cannot turn robots into genuine social creatures for the time being, maybe we can make them at least act in a socially appropriate manner in the interim, pretending that they know what they are doing and why, and thus making interactions with them less annoying and socially awkward

Examples of learning how to act socially

- Will focus on learning “how to act socially” from natural language instructions, which is an efficient and natural way for people to teach social behaviors
- This work is based on our ongoing work on “one-shot learning” where a robot can learn a new capability from one exposure, most recently on “One-shot Interaction Learning” (Frasca at el. 2018), “(Recursive) One-shot Object and Action Learning” (Frasca et al. 2018, Scheutz et al. 2017)
  - Learning to follow a person
  - Learning to pass a person an object
  - Learning to show an object to a person based on passing an object
  - Learning what to pay attention to when passing an object to a person
  - Learning to pass an object in a socially appropriate way

Learning to follow a person
(e.g., Cantrell et al. 2011)

Consider how one could tell a robot how to follow a person by staying behind them within a certain distance...

While the system may or may not understand the word “follow”, let’s assume that it does not have a procedural grounding for the action and that we to instruct the action making reference to the social other:

To follow means that

you should stay **within** **one meter** **of me.**

maintain(listener, **<(** **distance_from(speaker),** **1m)**

follow(?actor1,?actor2):
  maintain(<(**distance_from**(?actor1,?actor2),1))

Learning to pass a person an object
(e.g., Scheutz et al. 2017, Frasca et al. 2018)

H: Do you see an object?	R: Yes
H: This object is a plate	R: OK
H: Pass the plate	R: I do not
H: OK I will teach you how to pass the plate	R: OK
H: Pick up the plate	R: I do n
H: I will teach you how to pick up the plate	R: OK
H: First find the plate	R: OK
H: Then grasp the plate	R: OK
H: Move the plate up.	R: OK
H: That is how you pick up a plate	R: OK
H: Move the plate forward	R: OK
H: Then release the plate	R: OK
H: That is how you pass a plate	R: OK
H: Start over	R: OK
H: Do you see an object?	R: Yes
H: This object is a knife	R: OK
H: The gray part of the knife is the blade	R: OK
H: Pass the knife by the blade.	R: OK

One-shot object learning using perceptions

One-shot activity learning involving learned object

One-shot object learning

One-shot object part

Transfer of learned activity to new object

Learning to pass a person an object (e.g., Scheutz et al. 2017, Frasca et al. 2018)

Learning to show an object to a person based on passing an object (e.g., Frasca et al. under review)

Assume the robot already knows how to pick and pass a knife:
pass(?actor1,?knife,?actor2)
but that it cannot actually perform the action because it does not have the visual capabilities required for detecting a knife

H: Hello Dempster.
R: Hello Brad.
H: Show me the knife.
R: I cannot show it to you because I do not know how to sho you it.

H: Show me the knife is like pass me the knife, but remove release the knife
R: Okay.

H: Show me the knife.
R: OK. I cannot show you it, because I do know know how to find it.

H: Describe how you show me the knife.
R: To show you it, I find it, I grasp it, I move it up, and I move it towards you.

Learning a new skill

Script description showing that the robot learned the correct modification

Learning what to pay attention to when passing an object to a person

(e.g., Frasca et al., under review)

Assume the robot already knows how to pick up and pass a knife:

pick-up(?actor,?knife)
pass(?actor1,?knife,?actor2)

H: Hello Andy.
R: Hello Brad.

H: Please pass me the knife.
R: OK.
(drops the knife)

H: Start over.
R: OK.

H: When you pass me the knife replace move knife forward with move knife towards me.
R: OK.

H: Please pass me the knife.
R: OK.
(moves knife towards the human)

Replacement of movement action without sensing with action that requires sensing

Learning to pass an object in a socially appropriate way
(e.g., Sarathy et al. 2018)

Assume the robot already knows how to pick up plates and how to pass them a person:

pick-up(?actor,?plate)

pass(?actor1,?plate,?actor2)

H: Do you see the knife?
R: Sorry, what is a knife?
H: This object is a knife.
R: Got it.

H: A knife is used for cutting.
R: Alright.
H: The gray part of the knife is the blade.
R: Got it.
H: The orange part of the knife is the handle.
R: Got it.

H: To pick up a knife, grab the knife by the handle.
R: Alright.
H: To pass me the knife, grab the knife by the blade.
R: Alright.

One-shot object learning and zero-shot function learning

One-shot learning of object parts

Zero-shot affordance Rule Learning

At this point the robot learned another object ("the knife") as well as some of its object parts ("blade" and "handle") and two rules about picking up and passing knives, respectively

Applying new knowledge on the fly

H: Pass me something to cut with.
R: OK.

Pass me | something to cut with

Script retrieval
"pass"
Actor: robot

Find x s.t. x can be used for cutting
→ look for knife

Pass me plate
…
(only grasps)
Pass me knife

Pass me knife
…
Pick up knife by blade (in passing context)

Grab by handle
…
Find grasp on gray part of knife

Summary and ongoing work

- We don’t yet know what it takes to make truly social robots
- But we can, in the meantime, make steps towards developing robots that act socially, and ideally can learn how to act socially from human instructions
- Demonstrated the one-shot interaction learning capabilities of our DIARC architecture (e.g., Scheutz et al. 2018) and showed in several robot demonstrations how humans can instruct robot behaviors that involve interactants and require the robot to have (limited) representations of a social other
- We are working on developing more sophisticated models of human social actors which include richer mental and contextual models together with social and pragmatic rules and norm representations
- These richer models will allow robots to learn more natural, but also more socially appropriate interactions

Acknowledgments and links

- Former students/staff: Gordon Briggs (NRL), Tom Williams (Colorado School of Mines), Rehj Cantrell (Nuance), Luca Brusatin, Paul Schermerhorn (Soar Tech)
- Current students/staff: Vasanth Sarathy, Evan Krause, Brad Oosterveld, Tyler Frasca, Ravenna Thielstrom, and various undergraduate students at Tufts
- Collaborators: Rob Platt (Northeastern University) on grasping
- Funding provided by the US National Science Foundation and the US Office of Naval Research through several grants

- Papers at http://hrilab.tufts.edu/publications/
- Demos at https://www.youtube.com/user/HRILaboratory

We are hiring in AI/NLP/Robotics/HRI at TUFTS!