IN5480 Lecture notes 12 september 2019

Updated: 12 september 2019, 08:55
Note: This document will be revised for legibility and clarity. Please do not hesitate to send feedback about the document, info about errors, suggestions for improvement and so forth. I hope to revise and improve this text iteratively (of course without extending it too much - or adding new sections), in the same way as you will do with your individual assignment and the group assignment.

Today we continue with module 1 of the course, after these three introductory sessions with module 1, module 2 and module 3.

In your individual assignment you are asked to write about these three topics:

- Concepts, definition and history of interaction with AI.
- Human Robot Interaction.
- Universal Design and Interaction with AI.

The objective for module 1 is to:

- Get background knowledge and understanding about central concepts within AI.
- Get introduced to the Human Robot Interaction field
- To be able to discuss and critically think through Universal Design and Interaction with AI.

In this session, we start off with describing the six mandatory articles for module 1. Then we will attempt to unpack some concepts within Interaction with AI. After this we will work with Human Robot Interaction with focus on human movement and machine movement.

Articles in Module 1

1. Grudin, Jonathan. AI and HCI: Two Fields Divided by a Common Focus. AI magazine 30, no 4 (September 18, 2009).
2. Dautenhahn, K., 2018. Some Brief Thoughts on the Past and Future of Human-Robot Interaction. ACM Trans. Hum.-Robot Interact. 7, 4:1–4:3.
3. Thrun, S., 2004. Toward a Framework for Human-robot Interaction. Hum.-Comput. Interact. 19, 9–24.
4. Schulz, T., Herstad, J., & Torresen, J. (2018). Classifying Human and Robot Movement at Home and Implementing Robot Movement Using the Slow In, Slow Out Animation Principle. International Journal on Advances in Intelligent Systems, 11, 234–244.
5. Norman, D (1990). The problem of automation: Inappropriate feedback and interaction, not over-automation. Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences, Vol. 327, No. 1241, Human Factors in Hazardous Situations (Apr. 12, 1990), pp. 585-593

6. Verne, G, Bratteteig, 2018, Does AI make PD obsolete?; exploring challenges from Artificial Intelligence to Participatory design.

This is a mix of articles that are selected so that we can learn some of the central concept within human centered AI and Robots.

One of the challenges when doing studies, learning and writing is to keep track of literature. To have an overview of what articles you have read carefully and annotated and written about, and articles that are in the pipeline is important. If any of you want to make a library in Zotero for this class - and share it as a group library - that is superfine.

I will say a few words about each article here, one by one. Then, I suggest that each group pick one of these articles, and present in more detail next week. Each group select one article - and present what the article is about next week? Up to 5 minutes for each group - including questions and answers for the article?

What we talk about when we talk about…. AI and robots

You are asked in the individual essay to present different definitions on AI and of Robots. When writing, discussing and doing this, it will become evident that various research communities interpret and understand these concepts in different ways. There are no unified definition, explanation or understanding of neither AI nor Robots today. It very much depend on which scientific tradition or community the concepts are used in; like psychology, engineering, computer science or philosophy. However, this does not mean that we should avoid attempting to define it, make sense of it and come to some understanding of what is meant by these terms.

However, the AI field has a history - or histories - both pragmatic, technical history and a theoretical, philosophical history. AI has been understood in many different ways; as technologies, as simulation of activities, knowledge or represented knowledge, AI as models of biological processes, or theoretical and philosophical perspectives and understanding of AI for example.

Within HCI, there is a long history of modelling interaction based on an understanding of what is going on when humans communicate, work together, cooperate or collaborate. When using computers, be it with command based interfaces, graphical interfaces or physical buttons - it is after all similar with what is going on when a person is communicating with another person. For example as illustrated in the following example between two people living together, sitting at the breakfast table at home an early Wednesday morning:

Kim: Are you working today?
Sue: Yes.

This topic of this conversation that happen between two people is work. Ten minutes later when Sue is receiving a telephone call from a colleague at the office, this happens:

Colleague: Are you working today?
Sue: No. I am fixing the things at home.

Ths transcript is an example of what is often called the qualification problem, and is closely linked to the framing problem. Rommetveit says: “Mutual understanding on the part of conversation partners is achieved by an attunement to the attunement of the other, by which states of affairs are brought into joint focus of attention, made sense of, and talked about from a position momentarily adopted by both of them” (Rommetveit, 1990). The situation, the is specific. The lifeworlds of Sue, Kim and the colleague are negotiated in everyday life, and mutual understanding may happen.

In the example with Kim, Sue and the colleague. What is the correct, true or valid interpretation and understanding of these transcripts? Is Sue working or is she not working? Can both be true? Is she lying? It all depends on what the situation is, and who is doing the talking and listening, in what situation. One way to interpret this is that Sue is working on something at home with Kim that day, a renovation project they have for example. For her colleague at work, she is working. Or, Kim and Sue might have a joke together about work, and that Sue lies to Kim or the Colleague? The interpretation and the following understanding is constituted and evolve during the conversations.

We learn from and simulate, model and represent dialogues between people and then make similar dialogue systems between people and machines. Let us take another example of a telephone conversation between Kim and Sue, who are now at different locations.

Kim: Hello
Sue: Hello, Kim - good to hear your voice

This opening dialogue is used as inspiration for interfaces and interaction mechanisms between humans and machines in for example command based systems, graphical interfaces and voice based interfaces. Conversational analysis and interaction analysis are the discipline that investigates the talk going on between people; the turn taking, the handover, the start and ending sequences and so forth. Inspiration for modelling interaction and interfaces to computers has been communication between people since the early days in HCI. Terry Winograd and Lucy Suchman are examples of researchers who has been influential in this.

Concepts like *intelligence, thinking, knowledge, representation, information, communication, autonomy, automation, understanding and learning* are used when discussing interacting with AI. These concepts are ambiguous or polymorphous. They can be used to say something in one situation, by one person, - and something very different in another situation. They are not technical terms, with a general rigid, fixed explanation within computer science. However, these concepts might have specific definitions within certain research fields. For example learning might be defined quite specifically within branches of pedagogy.

Compare the concept above to well defined, and unified concepts like the second (SI unit), a TerraByte (2 to 40th byte) or frequency (measured in Hz). These last examples all have more or less clear and concise definitions in the sciences as well in everyday life, and explanations of it can be shared and understood in a “common way” or “mutually understood”.

There is no escape from the everyday use of central concepts within AI. Interestingly we all use these words and concepts in everyday life when we are speaking, listening, writing and reading. And this conversations are going on regardless if we are analysts, programmers, engineers, architects end users or shopkeepers. AI, and interaction with AI is on the agenda in many research communities. The Stanford institute for human centered AI (HAI) was recently established to foster multidisciplinar work for interaction with AI.

In the following, I will use the concept “understanding”, a concept that is often used for discussing AI systems. It is quite common in everyday language to talk about *mutual understanding*, or common understanding between people. Usually, we talk about mutual understanding between two or more people, for example in a project. That we understand a phenomenon “in the same way”. An example of this is about gravity. I know that you also know that when I drop a pencil, it will fall. We have a mutual understanding about this phenomenon in everyday life. We can further explain to each other details about friction, forces, velocity, acceleration, mass and so forth - and by explaining to each other, get to an understanding that is shared - or common.

Within literature on AI, “understanding” and “mutual understanding” pops up also when we talk about how computers, robots, algorithms and programs “understand”. One example is in the following sentence: “Robots will need to base their behaviour on human needs by understanding and anticipating them” (from IEEE technology and society magazine) - or the statement “on the importance of mutual understanding between machines and computers.”

What does this mean for us, students of Interaction with AI?
First, that computers has “understanding something in the first place”. This is also described as the computer has knowledge, as in Knowledge based systems.
Secondly, that this understanding is “shared” or “in common” with humans.

One interpretation of this way of using the word *understanding* is that it is “just a way of talking, and the authors do not really mean that computers “understand anything”. Understanding, historically, has meant that we, humans, make sense of something and know something. Historically, it has not made any sense to say that a machine “understand” anything. With tools like hammers, typewriters or measurement sticks, it would make no sense to say that “the hammer understands the nail” would it? Or that the pencil understand what is written on the page with it?

However, with computers as tools and medium, this is a modern way of talking. That “they” for example need to understand our behaviour, as in the example above. We, humans, are the ones who understand and interpret something, be it a text, some numbers or causality.

Traditionally, this has been a human activity. Two branches of philosophy that has concerned itself with this is hermeneutics and phenomenology.

One language researcher who tried to say something about this confusion is Searl, in the Chinese Room Argument. This is a thought experiment where he attempts to challenge what is done when manipulating symbols in a machine like manner. He wants to differentiate the simulation of understanding with genuine human understanding. Imagine a test where person sitting inside a box, where she receives notes with chinese symbols from a person outside the box, and she sends notes with chinese symbols out from the box. The notes she sends from the box, she has written herself by transforming the symbols according to rules. The lady in the box does not know chinese.

This situation may fulfill all the requirement for the Turing test, in that the person outside the box cannot differentiate the interlocutor in the dialogue with a human. She receives meaningful responses, on hand written notes. Is it then correct to say that the machine (here wizard of oz, i.e. simulated) understand anything?

The illustrated system simulates understanding. The same holds for any manipulation of symbols, regardless of how complex this is done by computing machinery. The argument is that genuine understanding will not be possible by the machine. What is possible is *simulated understanding* (or intelligence, intelligent behaviour). The core issue here is that understanding (and knowledge) seems to be a precondition for the theories that attempt to explain understanding (or knowledge) in the first place.

This argument is discussed since it was posted in 1984, and critiques of the experiment is addressed in a article in Scientific American magazine in 1990 for those who is interested and curious perhaps.

In what ways does this matter? If we say that the “computer understand this or that”, and that it “know about this or that”? In what way does it matter?

DISCUSSION 5 minutes.

In some situations in everyday life we treat computers and media like people!! We are definitely aware that this is not a human, it is a machine, but we behave, scream towards it…You stupid computer!!

Two often used concepts to describe what's going here is anthropomorphizing and zoomorphism. The etymology of anthropomorphizing says that it is “involving the attribution of human qualities to divine beings”. However, today we currently do this with mundane things as well. In what way is it possible to “not” anthropomorphize?

Children play with dolls and toys, and make them full of life. And some people do not stop to talk with their teddy bear even when they are grown ups. It is for example said that Arne Næss talked with his teddy bear every day, discussing different matters. However, then it is a “play” and a “game”. When doing this, children and grown ups know that it is not for real

that the teddy bear understand anything. Our imagination of what can be, and exploration of possible futures is “helped” by these teddy bears - perhaps much in the same way as dialogues with other machines. However, challenges happens when we take the teddy bear for real. That it has real knowledge and real understanding of anything.

Understanding is furthermore always of something. In order to understand something, we very often try to get an overview - by sorting things out and categorizing “something”. We find patterns, and organize the material in different ways to get to know something better. Be it numbers, physical shapes, colours or things that we come up against. Based on the sorting, we can organize into categories, decide what should be included and what should be excluded, and order it into appropriate sequences.

What are the limits of mutual understanding with other people? Is it possible for a man of 40 to truly understand what it is like to be 90+ years walking stairs? Or is it possible for a man to genuinely understand what a menopause is? During the latest CHI2019 conference in Glasgow, there were a number of HCI studies about design and evaluating of systems that supported women during their menopause. In what way is it possible for a man, who has never experienced this, and will not experience this, to understand and design for this? This is an open question?

The computer is extraordinary in some of this processes going on, as compared with us humans. Especially sorting and classifying. Yesterday I heard that the most advanced smartphone in 2019 has a processor that can do 1 trillion operations every second. It had 8.5 billion transistors in its silicone design. This is more transistors than there are humans her on this planet, in one chip!.

Exercice “attuning to the computer”

Sorting out. Classifying. Categorizing.
Playing a computer!
Get to know a bit of the activities going on in machines.

Activities and Understanding - what we can do?

What it is that we do? Behaviour and physical activities can be observed, described, analyzed and thought about. Our understanding cannot really be “seen” directly. We come to believe that the machinery is like the human, i.e. intelligent and understanding, whereas there are some differences between the two. What we can do is to observe, analyze and make sense of human activities, with and without machines.

Human Robot Interaction

This section is based on a fresh phd thesis from our institute by Trenton Schulz, with the title: *Exploration of Moving Things in the home.* This thesis was written as part of a research project, MECS here at the department of informatics.

HRI is the sister discipline of HCI - Human Computer Interaction. The R is for Robot, understood as a physical thing that can move.

In HCI, the computer is always somewhere, in a specific context or situation, be it a desktop computer or a smartphone. The interface to the user, and the activities taking place when using mobile or stationary computers happens somewhere, always at a specific location or place, i.e it is situated.

What is somewhat different in HRI is that the “thing” can move around, either “by itself” autonomously, or remotely controlled by operators or people other places. This kind of movement is usually called locomotion, the change of location over time. Some things do not change place, but has local movement of arms, heads or other things. This is in HRI terms often called “reconfiguration”. This can be exemplified with the Barney here. Often, in various research robots, there is both locomotion and reconfiguration taking place.

Example: Movement of arms and head.

How should it move? How to start and stop? How close should the thing be to a person? How will it tell what direction it is moving, and what about the “intentions”? These are some questions that are investigated within HRI, and its subfield Social Human Robot Interaction (SHRI).

Human movement and familiarity

Humans, we, are familiar with movement. From birth, even before birth, we start to move, and experience movement. First local movement of heart, blood and thereafter of whole limbs, before rolling, falling, jumping and walking and other ways of locomotion. We are familiar with movement. We are not only familiar with movement, we are these movements. And we are aware of the movements - and can talk about it in various ways.

Interestingly, movement always happen somewhere, in a specific location. In mathematics, we can represent equations to describe velocity, direction and acceleration in abstract terms. However, when it comes to actual movement, of human bodies, or things - they happen somewhere. At least to the best of my knowledge.

A brief exercise: 2 minutes of observing these two paperbirds. What movement do you observe? Write down, and we will talk about it afterwards.

Discussion.

A specific location is one of the condition for the possibility for movement. Domestic locations like dwellings, office locations, public places, streets. In what ways is this important? The context of use matters also here, as it does for traditional HCI and mobile HCI.

What happens when “computers” start to move around, in our domestic settings or in public settings? There are many different things going on when introducing self-moving things in the environment, and we will here try to unpack and describe some of the things that is going on.

Outside the window, there is one moving computer, the lawn mower. An autonomous thing, that is programmed to move around the fenced lawn, cut the grass with a rotating knife and then go charger station. What is the interaction with this device? Who is interacting with it, and in what situations? Everyday use of the lawn, sunbathing? This is a public lawn, and it is indeed interesting to listen to people watching the movement taking place, and the ways in which it is talked about. Still, the phenomenon of self-moving things is new, so it is highly visible to us as something unfamiliar and even strange. How do you think this will be in 20 year from now?

This year, downtown at Akershuskaia, there is a trial going on when a public bus is self-driving. Have you been there and observed - or had a ride?

Exercise: Discuss some autonomous moving things that you are familiar with, or that puzzle you. And movement of things that seems like it is “moving by itself”?

We said that one basic condition for the possibility of movement is place. A second condition for human movement is “orientation”. We, humans, do this all the time, more or less aware of what we are doing. At the moment, we know that this is up, this is down, the wall is over there. The table is there, and my colleagues are there. In order for us to move swiftly around, we use this awareness of the surrounding in order to orient ourselves and move about. And this is true, both for local movement (configuration) and for locomotion. An example of the former is when sitting at the desk, I am oriented towards this tabletop, with the items being sorted here in some way. Something is close to the hands, some other things are further away. There is a side there, and something is behind me. Many things are not on the table. This is nothing extraordinary, we “know” this by circumspection, or omsikt in Norwegian.

Autonomous robots needs either a fully represented model of the environment to operate, and/or it require sensors to detect physical conditions and then make a representation of the environment. In order for self driving cars, or autonomous lawn movers to move safely around, sensors are needed to detect or sense things in order to avoid obstacles and to navigate.

Here we have used the concept autonomous a few times. Auto-nomous, what does that mean? What are the differences between human autonomy and machine autonomy? In what ways can autonomy of humans and autonomy of machines be negotiated?

Sensors.

In the MECS project, we have worked with different kind of sensors to map out the environment. RGB camera, Infrared, ultra wideband. These sensors “see” nothing, but they

get information from the environment about different shapes, colours, distances, temperatures and so forth - from the position of the sensor system.

So the sensors are there for two main reasons:

- to get information about the environment. This in order to be able to navigate in the place, find its way and avoid obstacles.
- to get information from the user who want to interact with it, steer it.

Is it a necessary condition for human users interacting with robots to know what sensors the robot is equipped with? When two people communicate, we know roughly what the other interlocutor sees, hears, smells etc. When the interlocutor is a robot, do we then need to know what it "sees"? This is a question that is worth investigating.

Exercise

The telephone table. Discuss various ways in which the table can

a) sense the environment
b) sense input from the user
c) provide information to the user about where it is going.

Wrapping up:

In this session we have done two things.
First, we have attempted to address the question "what do we talk about when we talk about AI and robots". This matters, because words matter.
Then, we have introduced Human Robot Interaction, and especially investigated human movement and the movement of things.