Communicating with AI

Prof. Dr. Jan Borchers Media Informatics and Human–Computer Interaction RWTH Aachen University

hci.rwth-aachen.de



About i10



HCI Computer Science



Talk about communicating with Al today But wrong picture of Al = robots So start with history But will come back to this conflict later



First message: AI is not magic, far from it



AI = development of computational systems to imitate human behavior.

Intelligent systems: practical implementation of AI for useful stuff



Created by Alan Turing in 1950, the Turing test was designed to determine if a computer is "intelligent".

An AI is said to pass the Turing Test if it can reliably fool human interrogators.

Weak AI relies on fairly specific tricks and techniques to solve a particular problem – think Siri, or predictive texting, or Google's news clustering algorithms.

Welc	ome to						
	EEEEE	LL	IIII	ZZZZZZZ	AA	AAA	
	EE	LL	II	ZZ	AA	AA	
	EEEEE	LL	II	ZZZ	AAA	АААА	
	EE	LL	II	ZZ	AA	AA	
	EEEEE	LLLLLL	IIII	ZZZZZZZ	AA	AA	
Eliza	a is a mock Rogerian	psychoth	erapis	t.			
Th -							1000
ine (original program was	describe	ea by J	oseph Wei	zenb	aum in	1966.
This	implementation by N	lorbert La	andstei	ner 2005.			
	THE CHICK OF CHICK OF L						
CLIZA:	Is something troubl	ing you ?	?				
YOU:	Men are all alike.						
ELIZA:	What is the connect	ion, do y	you sup	pose ?			
YOU:	They're always bugging us about something or other.						
ELIZA:	Can you think of a	specific	exampl	.e ?			
YOU:	Well, my boyfriend	made me c	come he	re.			
ELIZA:	Is it important to	you that	your b	oyfriend	made	you co	ome here
YOU:	He says I'm depress	ed much o	of the	time.			
ELIZA:	I am sorry to hear	that you	are de	pressed.			
YOU:	It's true. I am unh	appy.					
ELIZA:	Can you explain what made you unhappy ?						
YOU:		10		FF1 -			
100.							

Weizenbaum 1964

Just simple syntax analysis and echoing But response...



1959 Dartmouth Summer Conference, organizer John McCarthy.

LEFT: Marvin Minsky: top-down (symbolic)

RIGHT: Others: bottom-up (neural networks) - this will give us problems

- Machine learning is learning from data. Example: Training a spam filter (supervised learning). Spam tends to have more misspellings, more complex URLs, and less personal addressing.

- Unsupervised: Given lots of information, system has to find patterns without any feedback. Typical mode of Deep NNs. Can discover patterns (create new knowledge) that the trainer did not know.



To explain rule-based system:

Bruce Willis and Samuel L Jackson in "Die Hard With A Vengeance": 5 gal and 3 gal jugs, measure 4 gal (to defuse a bomb)



General "solver" + specific rules

Solution:

1. Fill the 5-gallon jug and pour that water into the 3-gallon jug until the 3-gallon jug is full, leaving 2 gallons in the 5-gallon jug.

2. Empty the 3-gallon jug and pour the 2 gallons of water from 5-gallon jug into the 3-gallon jug.

3. Refill the 5-gallon jug and pour that water into the 3-gallon jug until the 3-gallon jug is full, leaving 4 gallons in the 5-gallon jug.



SRI 1966: Shakey. Half an hour to calculate each move. It's just code.

R1/XCON(DEC 1980): Expert system to configure computer deliveries, saving DEC \$25m/year, 2500 rules

1988 Statistics supplements rules: IBM publishes "A statistical approach to language translation". Introduces PROBABILITY into the rule-driven field of machine learning. A move towards mimicking cognitive processes of the human brain.

1997: IBM's Deep Blue beats Garry Kasparov. Rule-Based.

- Watson (IBM 2011) beats the two best performers of all time on the US quiz show jeopardy. ("IBM supercomputer destroys humans in jeopardy")

- Watson (IBM 2011) beats the two best performers of all time on the US quiz show jeopardy. Wins \$1M

- Apple Special Event (2011) announcing Siri, developed by SRI and Nuance, became a standard feature on the iPhone 4S in 2011.
- June 2012: Stanford and Google publish unsupervised learning based on deep neural network with **1 billion connections.** This avoided the task of manually labeling data for machine learning.

- Generative Adversarial Networks (2014). thispersondoesnotexist.com

A generative adversarial network (GAN) is a class of machine learning systems invented by lan Goodfellow and his colleagues in 2014.[1] Two neural networks contest with each other in a game (in the sense of game theory, often but not always in the form of a zero-sum game). Given a training set, this technique learns to generate new data with the same statistics as the training set. For example, a GAN trained on photographs can generate new photographs that look at least superficially authentic to human observers, having many realistic characteristics. Though originally proposed as a form of generative model for unsupervised learning, GANs have also proven useful for semi-supervised learning,[2] fully supervised learning,[3] and reinforcement learning.[4] In a 2016 seminar, Yann LeCun described GANs as "the coolest idea in machine learning in the last twenty years".[5]

- 2015 Machines "see" better than humans. The annual ImageNet challenge (algorithms compete in recognizing and describing images from large library), now at >90%.

- 2016 Alpha Go, created by Google Deep Mind, defeated the world Go champion over five matches. The number
 of game variations required **neural networks** to study the game and learn as it played. Experts found that it
 discovered moves that hadn't been thought of before.
- 2016 New doctor program takes symptoms and suggests possible diagnosis.

- 2018 self-driving cars (suggested by GM at the 1939 World's Fair), 2018 Google spinoff Waymo self driving taxi service launched in Phoenix. Human driver still behind wheel for emergencies.

- 2019 designer Philippe Starck and Kartell unveil the A.I. Chair, using Autodesk generative design software.
 Software learns from each iteration about the designer's intent.
- The Turing test was passed about 5 years ago... with few people noticing, because it does not seem that all encompassing anymore.

Now you have a 50,000-ft overview of how AI has developed

The right side is more powerful but also problematic, more similar to how we think

LEFT: Marvin Minsky: top-down (symbolic)

RIGHT: Others: bottom-up (neural networks) - this will give us problems

- Machine learning is learning from data. Example: Training a spam filter (supervised learning). Spam tends to have more misspellings, more complex URLs, and less personal addressing.

- Unsupervised: Given lots of information, system has to find patterns without any feedback. Typical mode of Deep NNs. Can discover patterns (create new knowledge) that the trainer did not know.

XAI example (2014)

XAI is not enough.

What is a useful explanation?

Back to communicating with AI

Key Message 2: We don't know how we even want to communicate with AI.

For the last 40 years, HCI used to consider the computer as a tool. Human initiative, predictable,... Guidelines. Human-Centered, not Human.

But: Doesn't work with autonomous systems.

But people treat computers as social actors anyway, so...?

Treat AI as human?

No. It's superhuman in some respects (patience). Should we thank it? Diminishing your own contribution to a task (unlike tool). Should also be able to reprogram quickly.

Augmenting Human Senses and Cognition

Is this OK? Story of little girl & Alexa

Weird fact. Ask anybody to imagine a future with AI...

- they always imagine evil overlords. Why? Only movies?

Today, Narrow AI already sociopathic.

With what kind of AI do we want to live?

I want at least *pleasant* entities to interact with. Honest, supportive,...

CHI'19 MSR Guidelines

Amershi et al. (Microsoft Research): Guidelines for Interacting with AI (CHI'19)

A Call to Arms

Some Pointers

- Martin Maguire (U Loughborough): History of Al
- Marc Hassenzahl (U Siegen): How to perceive Al
- Albrecht Schmidt (TU Munich): Human Augmentation vs. Al
- Amershi et al. (Microsoft Research): Guidelines for Interacting with AI (CHI'19)

RWTHAACHEN UNVERSITY

36 Name: Topic