Trade-Offs in Human-AI Interaction

Human-AI Interaction

Luigi De Russis

POLITECNICO DI TORINO

AI+HCI

- "From the earliest times in the development of computers, activities in the field of Human-Computer Interaction (HCI), and Artificial Intelligence (AI) have been intertwined. But as subfields of Computer Science, AI and HCI have always had a love-hate relationship."
- "Together, the community can make user interfaces a little less stupid and frustrating than they are today."
- Henry Lieberman, "User Interface Goals, AI Opportunities", AI Magazine, 2009, https://doi.org/10.1609/aimag.v30i4.2266

Can AI and People Really Work Together?

- Should we design and develop technology that
  - automatizes people's actions, replaces humans (*Artificial Intelligence*)
  - or augments them (*Intelligence Augmentation*)?
- Is there a clear winner?

- Which kind of human-AI collaboration can we envision?
  - full human control… no human involvement?
  - full AI autonomy… no AI?
  - in which cases?

Artificial Intelligence vs. Intelligence Augmentation?

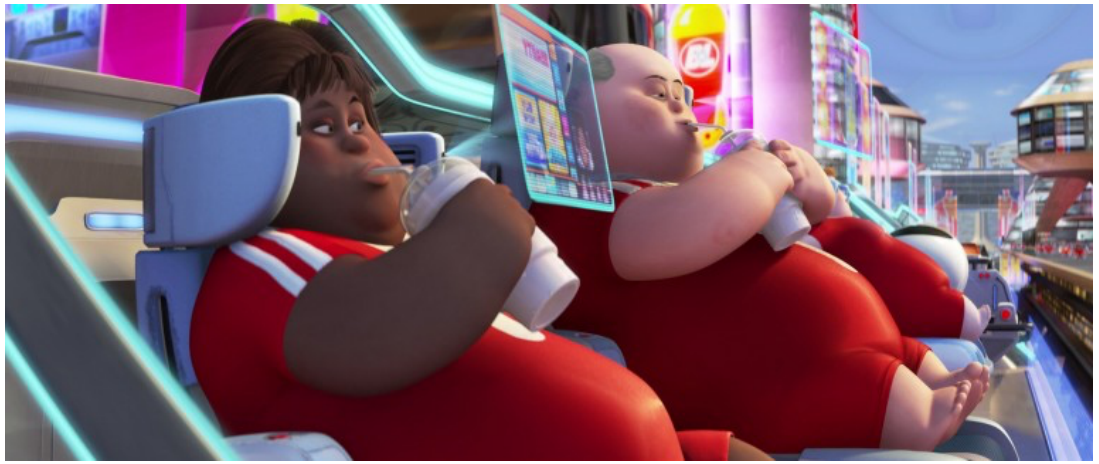

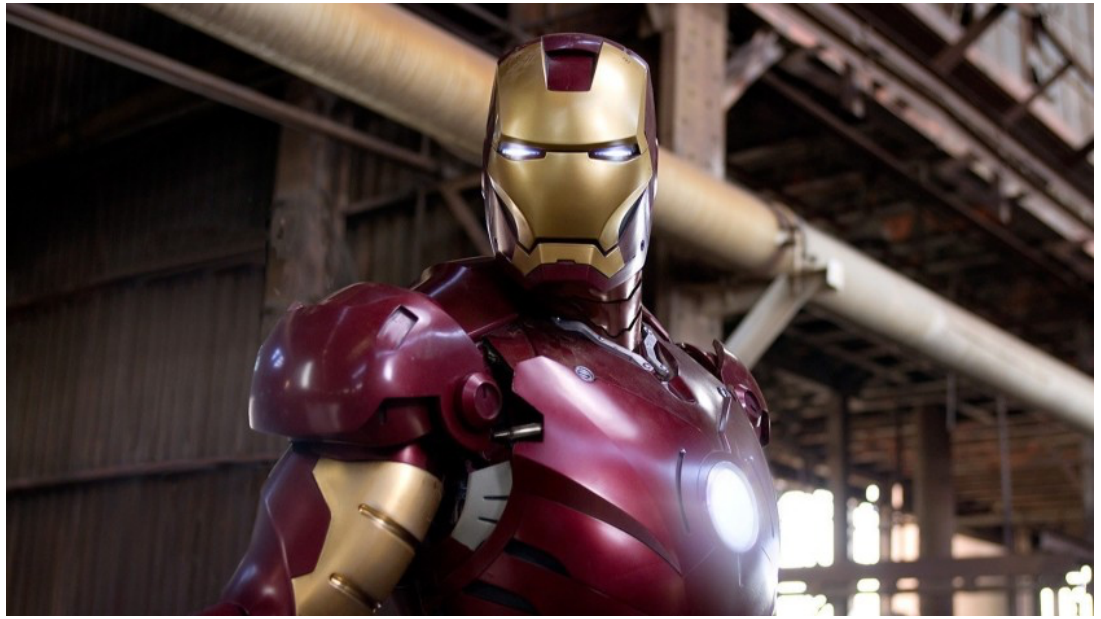

Artificial Intelligence

- "To proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that **a machine can be made to simulate it**" (John McCarthy, 1956)
- "Machines will be capable, within twenty years, of **doing any work** that a man can do" (Nobel Laureate Herbert Simon, 1965)
- Defining AI is not trivial
  - building a human-like intelligence?
  - in general? for specific kind of task, only?

Intelligence Augmentation

- "**Elegant combination** of reasoning machinery and direct manipulation" (Horvitz's paper, 1999)
- "Computing machines will do the routinizable work that must be done **to prepare the way** for insights and decisions in technical and scientific thinking" (J. C. R. Licklider, Man-Computer Symbiosis, IRE Transactions on Human Factors in Electronics, 1960)
- People have been realizing technologies to augment their intelligence (and activities) for millennia…

Let's Put This In Context…

- 1930-1940: early start of modern computer science
- 1956: direct keyboard input to computers
  - for the first time, in research, at MIT
- ASCII was invented in 1963
- 1966: HP introduces its first computer

As We May Think

- "Consider a future device… in which an individual stores all his books, records, and communications, and which is mechanized so that it may be consulted with exceeding speed and flexibility. It is an enlarged intimate supplement to his memory."
  - Vannevar Bush, 1945
    - https://www.theatlantic.com/magazine/archive/1945/07/as-we-may-think/303881/
- Video (from 2009):
  - https://www.youtube.com/watch?v=c539cK58ees

As We May Think

- Is it AI? Or is it augmenting people?
- Does it "look" to be related to any current (piece of) technology?

Put-That-There

Richard A. Bolt, Put-That-There: Voice and Gesture at the Graphics Interface, SIGGRAPH, 1980, https://www.youtube.com/watch?v=RyBEUyEtxQo

Put-That-There: Assumptions

- About speech recognition:
  - why do you need to point?
  - why do you need to talk?
  - what might go wrong?
- About people:
  - How would this augment them?
  - How might it fail?

CoBot

- "CoBot robots follow a novel **symbiotic autonomy**, in which the robots are aware of their perceptual, physical, and reasoning limitations and proactively ask for help from humans, for example for object manipulation actions."
  - source: http://www.cs.cmu.edu/~coral/projects/cobot/

- Video: Autonomy and Human-AI Interaction, Manuela Veloso, CMU
  - https://www.youtube.com/watch?v=OhoTFGjUPQs

CoBot

- Is this AI or IA?

- Assumptions
  - What do people do?
  - How do they learn how to interact with the robot?
  - How do the robot learn?
  - What can go wrong?
  - What are we missing?

The Apple Knowledge Navigator (1987)

source: https://www.youtube.com/watch?v=JIE8xk6Rl1w

Are We There?

- What parts of augmentation have we achieved?
- What parts have we not? Why?
- Should we fully reach "there"?

Humans or AI: Who Should Have The Conn?

- It is essential that people *feel* in control of their lives and surroundings
- When we "put intelligence" in things, people should:
  - be **comfortable** with the actions made by the intelligent system
  - **understand** why some actions are happening
  - **trust** the intelligent system
- "Automation" is typically met with *resistance*
  - however, it can reduce the workload and allow to complete dangerous tasks
- We should avoid (and consider) over-exaggerated expectations
  - *claimed*: "we have reach human-parity in speech recognition!"
  - *pre-assumed*: "I can speak with it, it understands my words, THUS it has full language understanding"

Amazon Go

- What do you think? Cool? Not cool? Is everything totally fine? Why is this "accepted"?

source: https://www.youtube.com/watch?v=NrmMk1Myrxc

"Full" Self-Driving Cars

- Who should be in control? When? What can go wrong? Why?

source: https://www.youtube.com/watch?v=tlThdr3O5Qo

Panels

Some instructions

Organization

- Panelists will seat in the first row of the room
- Panelists will have around 30-40 minutes, overall, to present themselves, their positions, and answer to the moderator's questions
- Audience will have around 20-30 minutes, at the end of the previous step, to pose questions to one or more panelist, i.e., a question can be for a single panelist or for a "group" of them
- Audience may ask questions either by voice or by using slido: https://app.sli.do/event/iqm4fgju
- At the end of both panels, we will have a 5-10 minutes summing-up session

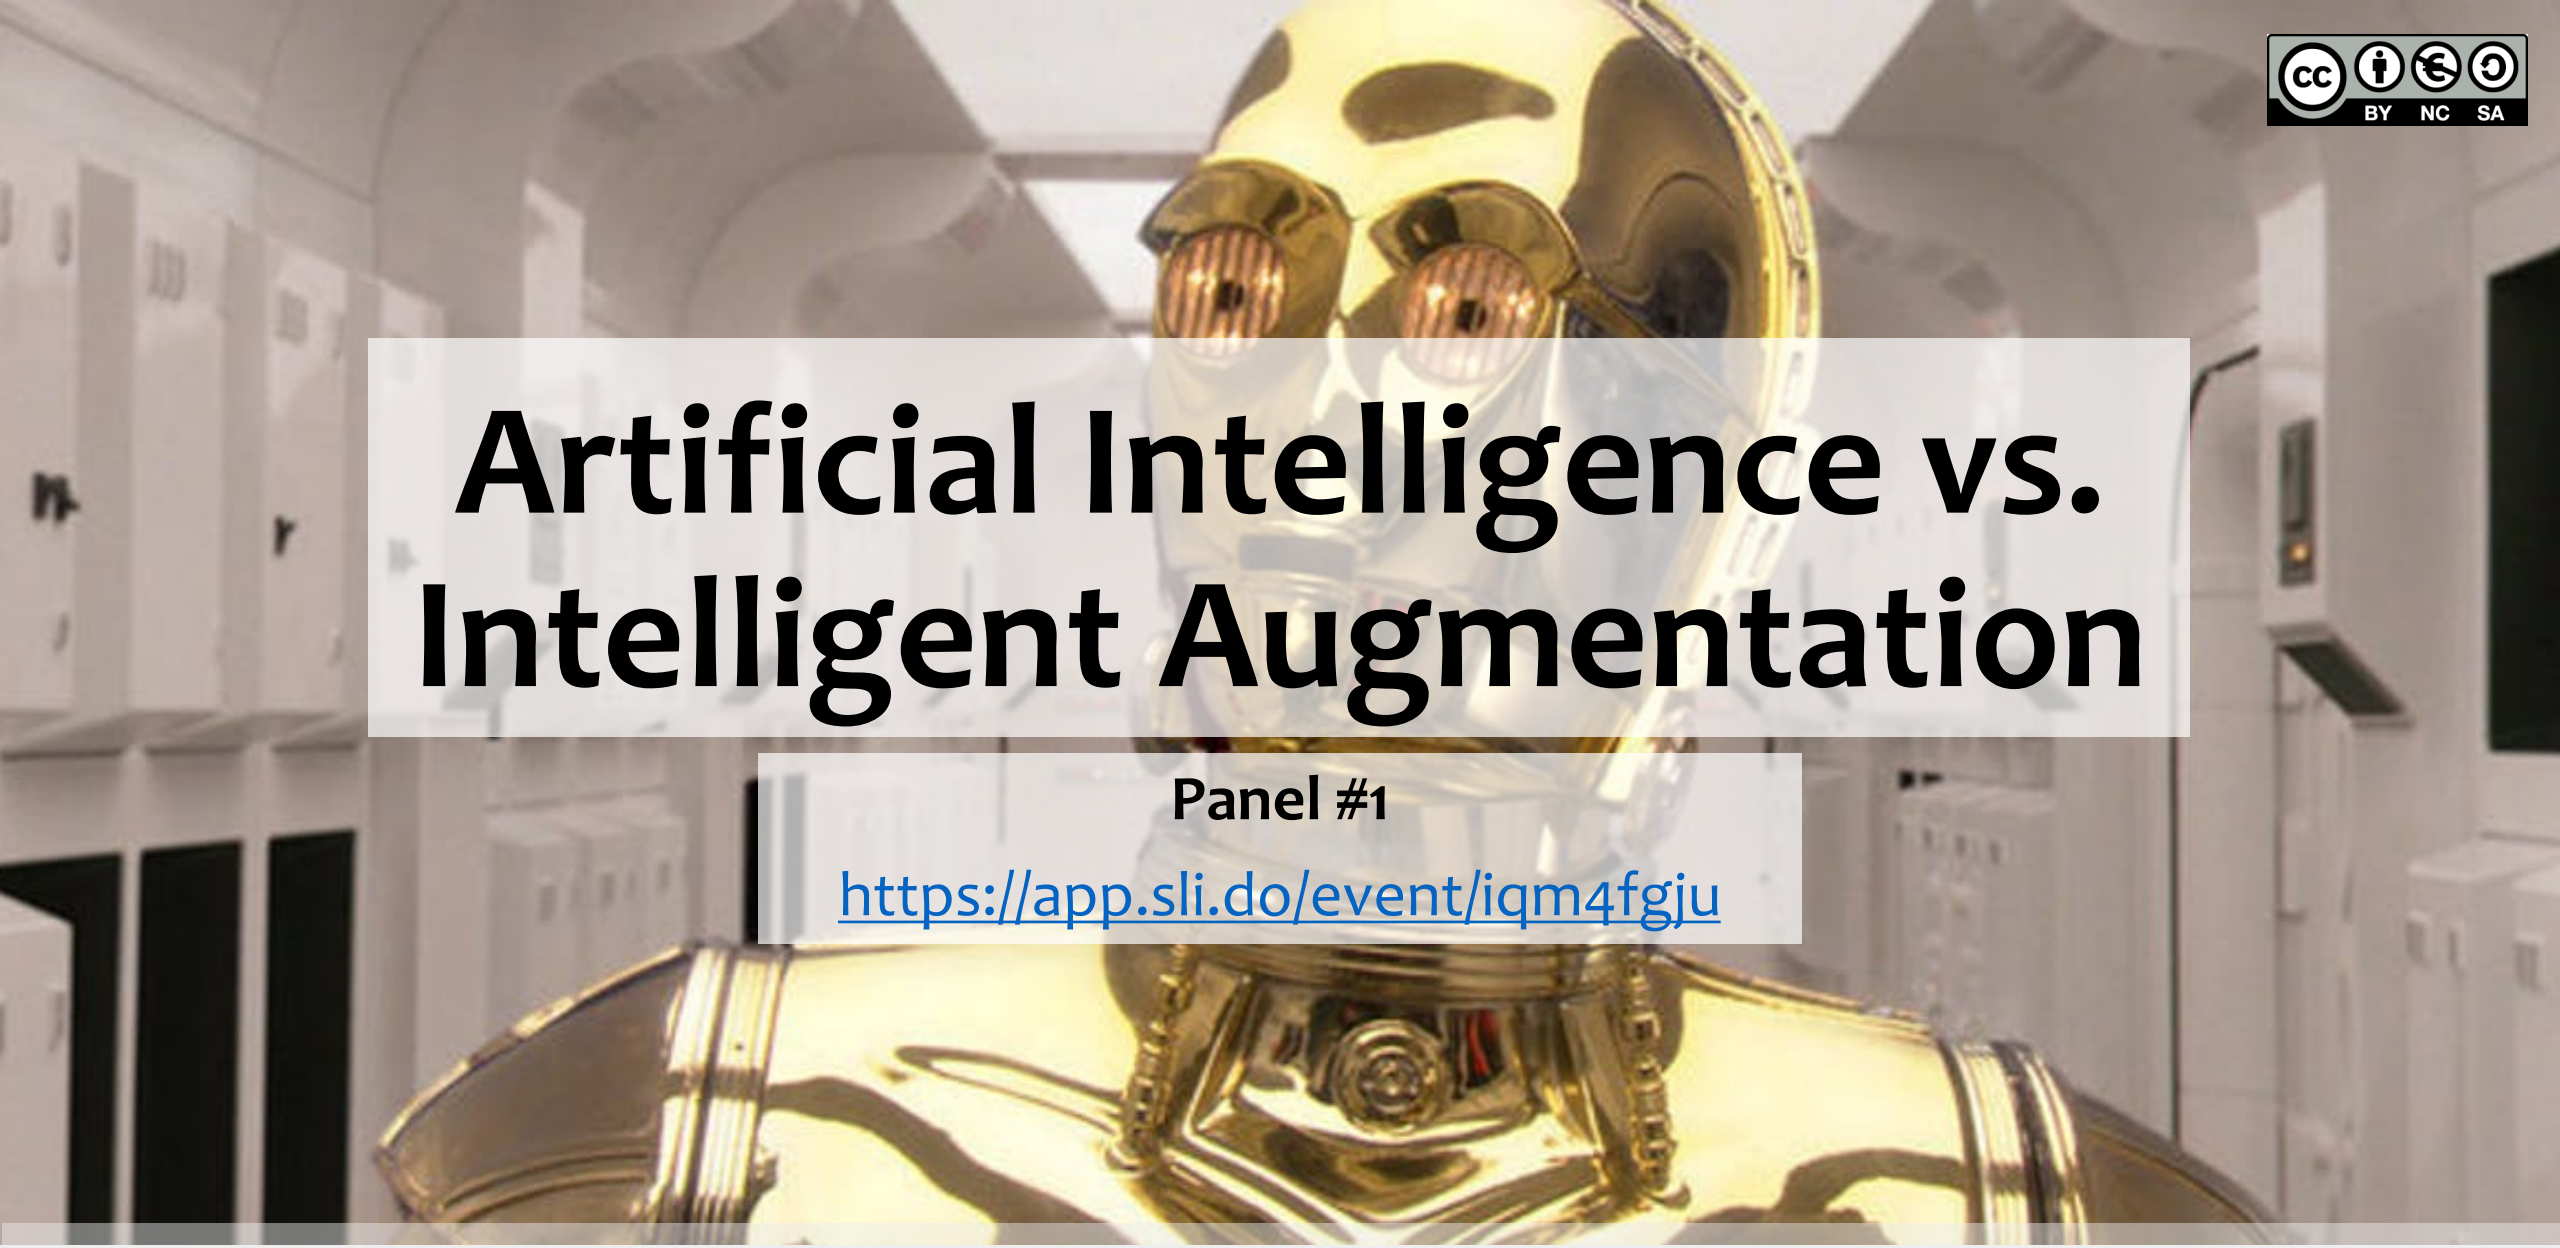

Artificial Intelligence vs. Intelligent Augmentation

Panel #1

https://app.sli.do/event/iqm4fgju

Questions

From the audience

Is it better to have control or accuracy?

Panel #2

https://app.sli.do/event/iqm4fgju

Questions

From the audience

After-panels Discussion

Wrapping up…

Questions

- Is there a clear winner between AI and IA?
  - in general? in some particular cases?
- Is IA a good strategy to build intelligent systems?
  - where might it be problematic?

Questions

- Is there a clear winner between human control and AI autonomy?
  - in general? in some particular cases?
- Do we need to balance them?
  - how and when?

License

- These slides are distributed under a Creative Commons license "**Attribution-NonCommercial-ShareAlike 4.0 International (CC BY-NC-SA 4.0)**"

- **You are free to:**
  - **Share** — copy and redistribute the material in any medium or format
  - **Adapt** — remix, transform, and build upon the material
  - The licensor cannot revoke these freedoms as long as you follow the license terms.

- **Under the following terms:**
  - **Attribution** — You must give appropriate credit, provide a link to the license, and indicate if changes were made. You may do so in any reasonable manner, but not in any way that suggests the licensor endorses you or your use.
  - **NonCommercial** — You may not use the material for commercial purposes.
  - **ShareAlike** — If you remix, transform, or build upon the material, you must distribute your contributions under the same license as the original.
  - **No additional restrictions** — You may not apply legal terms or technological measures that legally restrict others from doing anything the license permits.

- https://creativecommons.org/licenses/by-nc-sa/4.0/