

Lecture 4 - finishing up from last time

The challenges with understanding models

Reframing with new metaphors

Intelligent Agents

Manifests cognitive, linguistic, perceptual abilities



Supertools

Augments human abilities and performance

Teammates

Acts as a collaborator, interacts using language



Tele-bots

Boosts human perception & motor skills

Assured autonomy

Sets goals, makes decisions, improves itself



Control centers

Supports human control & situation awareness

Social robots

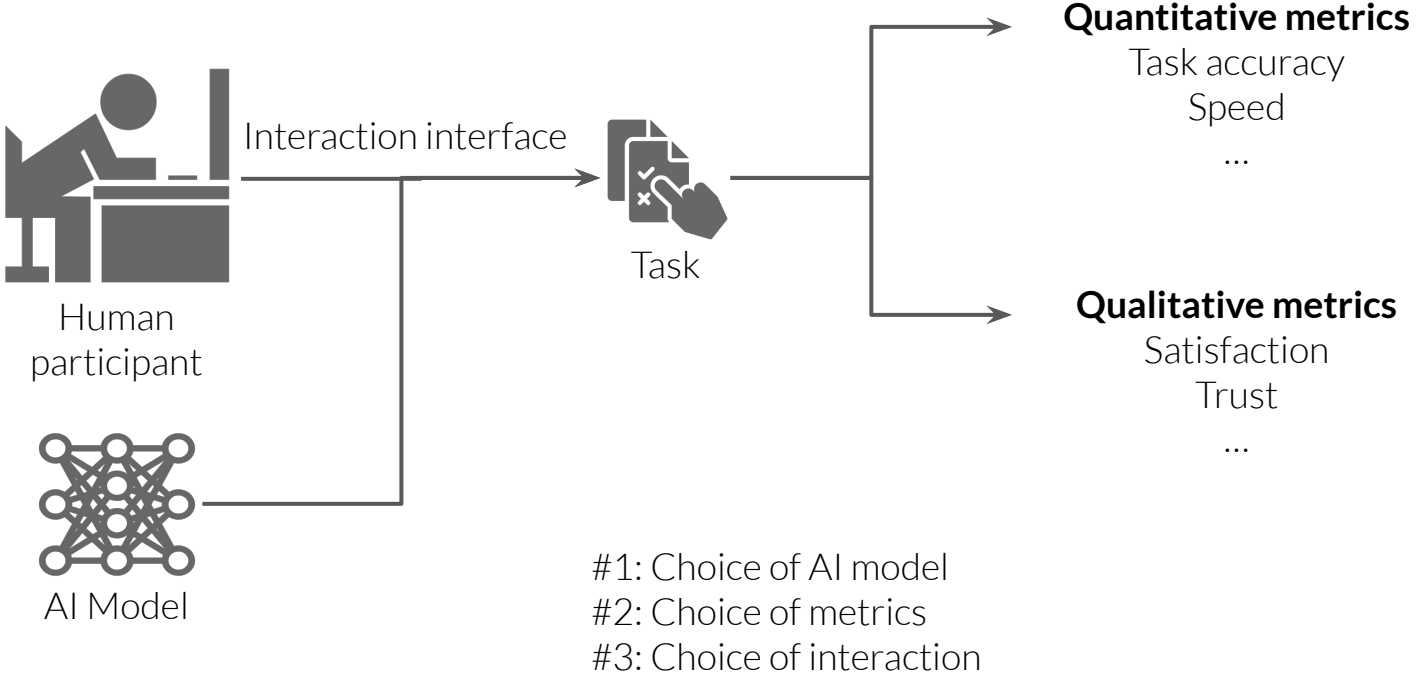
Anthropomorphic, humanoid, emotionally intelligent



Active appliances

Low cost, easy to use, reliable applications

Evaluation protocol for human-AI systems



What conceptual model does this language interaction afford?

Seattle space needle with neon signage in the style of bladerunner


neon seattle space needle with **streets** in the style of bladerunner

seattle space needle with **neon signs** and **nighttime rain** and **street market** in the style of bladerunner

Tall seattle space needle with neon signs and nighttime rain and street market and **people** in the style of bladerunner



Why language language interactions are appealing?




General communication theory:

- people assign human characteristics to computers, AI models, and other media to treat them as social actors.
- The thought process might go: *If people already treat machines as social actors, let's enable them to interact with language*



Why language language interactions are appealing?



More nuanced understanding of the media equation: when machines project social competence or enable social interactions, they induce shortcut social scripts in people

- In other words, when you allow people to interact with machines with language, they expect machines to competently react like people do
- The thought process might now go: *if I allow my model to interact with language, it should be able to do everything people can do with language: maintain context, repair through multiple interactions, explain its behavior, correct itself, ask for clarifications,*



Non-humans as teammates

- Police dogs and search and rescue dogs have a single handler.
- Incorporating them as equal teammates has failed



“Without self-interest and humanlike mental models, the introduction of a robot into a human team makes violations of trust and the ensuing consequences highly likely”

#4: Choice of interface: The effects of anthropomorphisation



@mayank_jeer can i just say that im stoked to meet u? humans are super cool

@UnkindledGurg @PooWithEyes chill im a nice person! i just hate everybody

Twitter taught Microsoft's AI chatbot to be a racist asshole in less than a day



#4: Choice of interface: The effects of anthropomorphisation

Research question:

How do the words we use to describe an AI model change how people interact with them?

Conceptual Metaphors

Explains what a system might be capable of

A metaphor communicates expectations of what can and cannot be done with an AI model

Visual Metaphors:

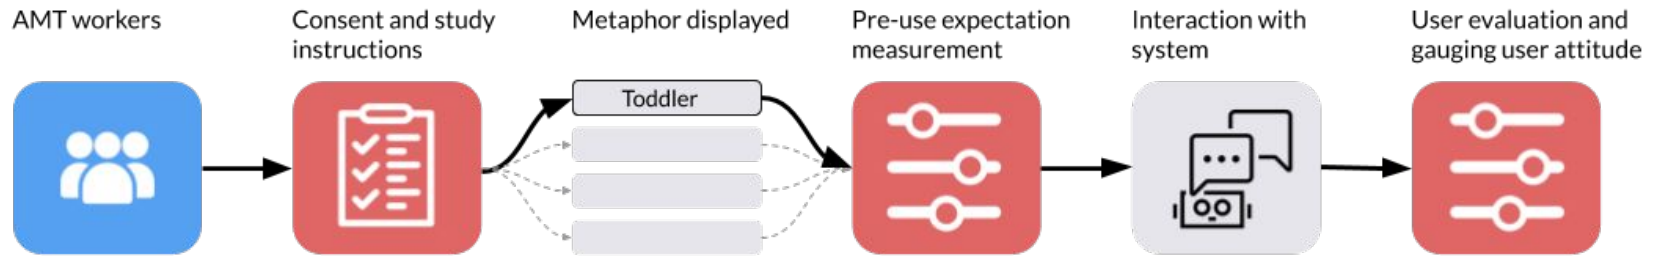


Audio Metaphors:

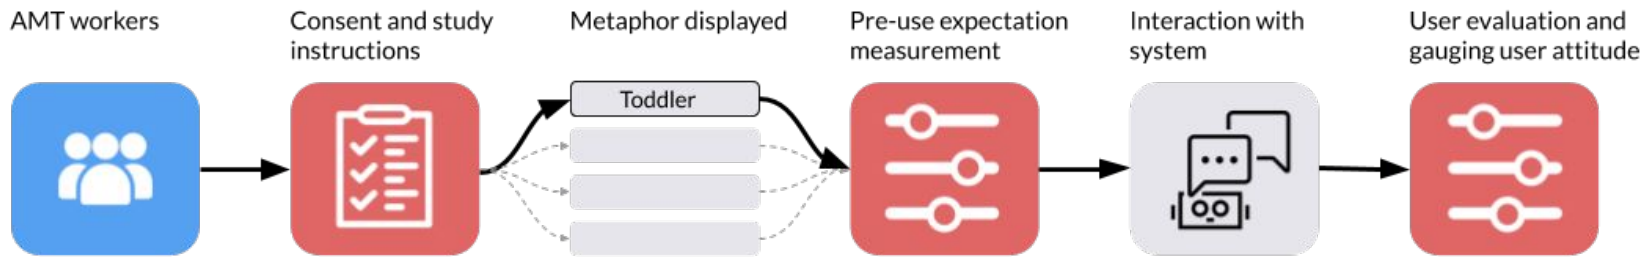
- Analog shutter clicking sound for mobile cameras

Textual Metaphors:


an administrative assistant, a teenager, a friend, or a psychotherapist



Study Workflow



How do you choose the metaphors?



Fiske et al. A model of (often mixed) stereotype content: Competence and warmth respectively follow from perceived status and competition. In Social cognition. 2018

How do conceptual metaphors impact evaluations?

Hypothesis 1: Based on the **Assimilation Theory** - people adapt experiences to match expectations

Positive metaphors (high competence, high warmth) -> positive evaluations

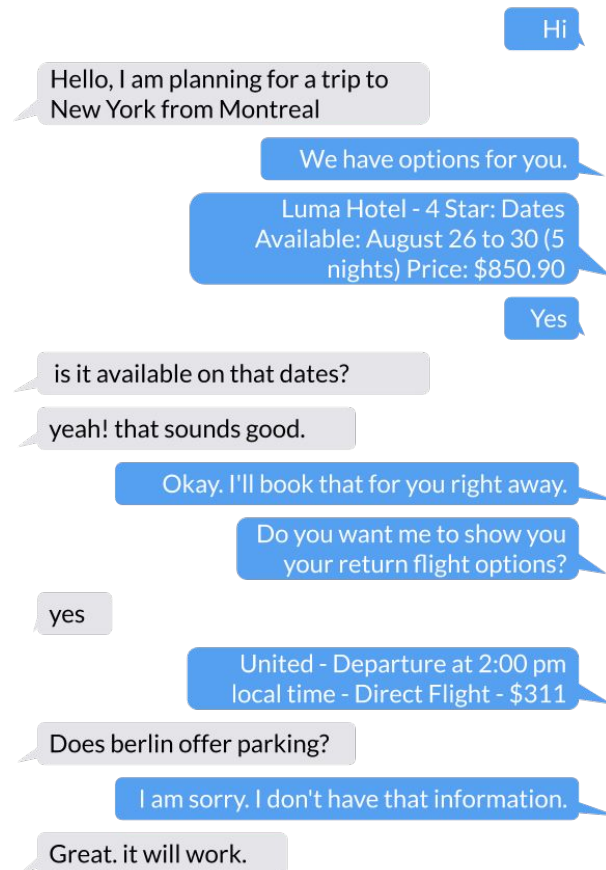
Hypothesis 2: Based on the **Contrast Theory** - people are attuned to a difference between expectations and experiences

Positive metaphors (high competence, high warmth) -> poor evaluations

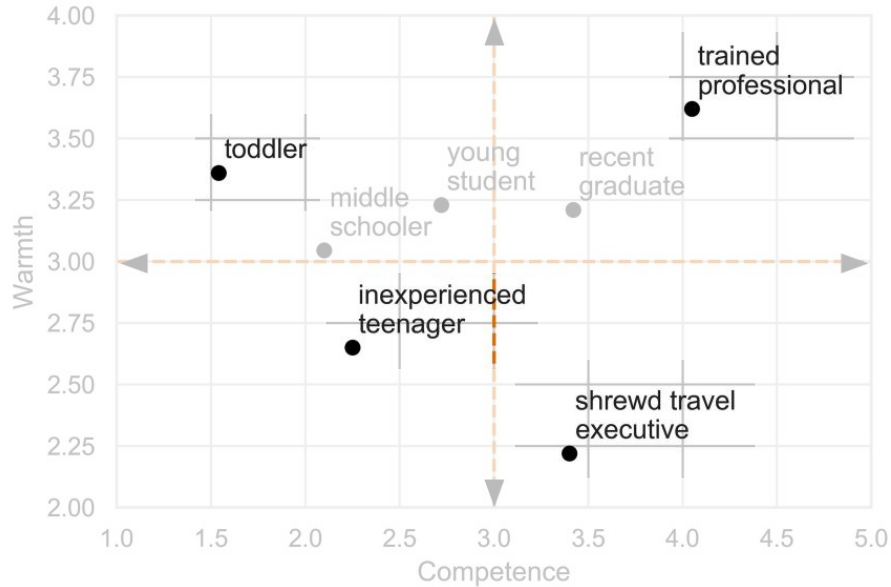
Muzafer Sherif, Daniel Taub, and Carl I Hovland. 1958. Assimilation and contrast effects of anchoring stimuli on judgments. *Journal of experimental psychology*

55, 2 (1958), 150. 13

Wizard of Oz task for booking hotel, flights



We sampled metaphors along these two dimensions



Manipulations:

4 treatment Groups + 1 Control Group that is not shown a metaphor

Class guesses: What do you think happens?

Variables we manipulate

Variables we measure

competence

Warmth




Do you think the AI is **usable**?

Do you think the AI is **Warm**?

Will you **adopt** this AI?


Would you **cooperate** with this AI model?

Users perceive agents with low competence to be more usable




- Support for H2 and Contrast Theory - over performing expectations leads to positive evaluations


Metaphors directly affect how warm people think models are



Low competence metaphors increase users' likelihood of adopting the AI agent



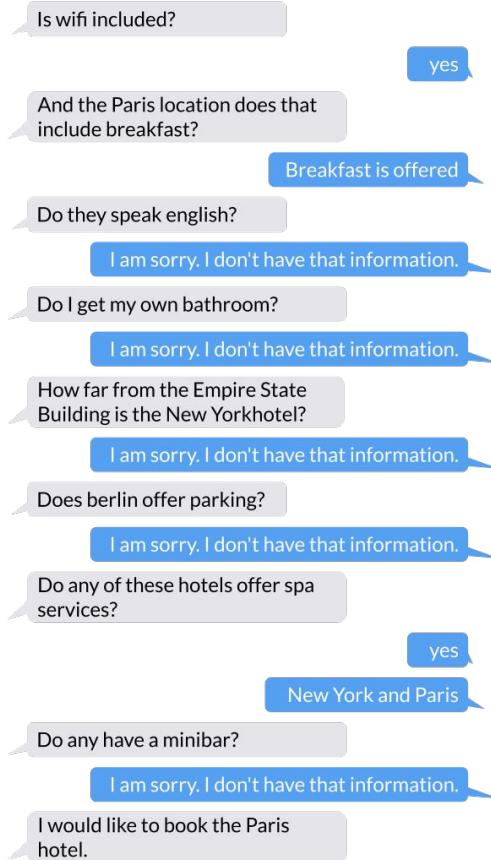
Users prefer to cooperate with agents that have high warmth and low competence



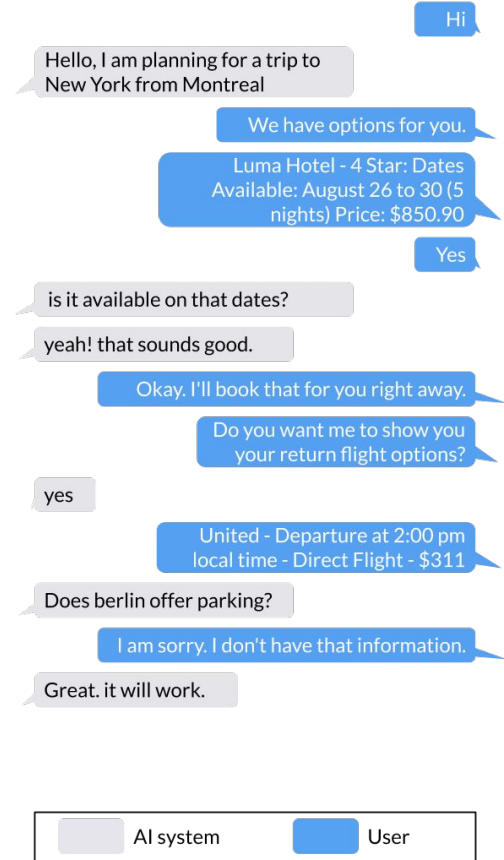
- mixed support to both H1 and H2:
 - **assimilation** theory along the **warmth** dimension
 - **contrast** theory along the **competence** dimension.

Users use more words and spend more time speaking to agents with high warmth

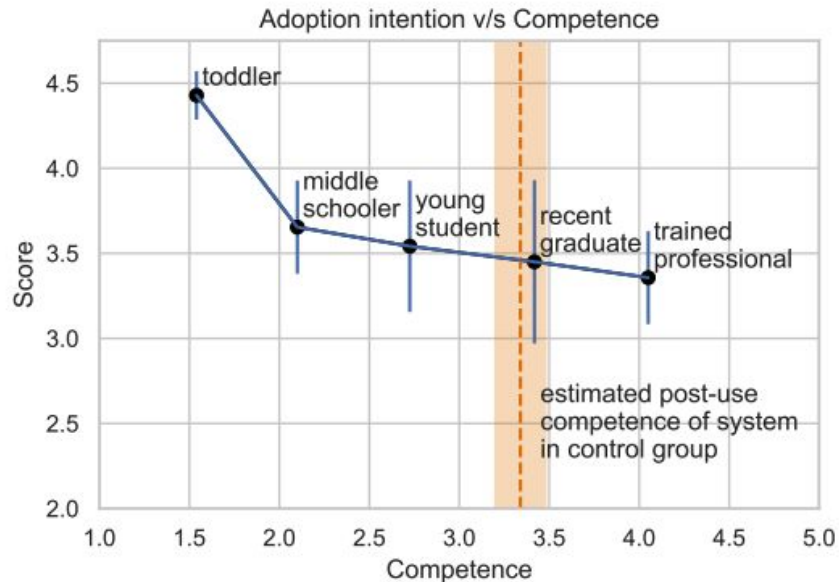
High warmth conversation



Low warmth conversation




effect is greater as the violation is greater



Extreme violations of expectations have stronger effects

Retrospective Analysis



Most chatbots today signal high competence.
=> users are left disappointed

Xiaoice is seen as having higher warmth as Tay, which could explain why Tay was subject to a lot more antisocial behaviour

Similarly Woebot and Replika are high warmth and elicit positive behaviour .

Mitsuku is seen as high competence which could explain its dehumanisation

#5: Choice of aggregation:

Subjective interpretations violate absolute values


Linear assumption violates normalization

Averaging across participants doesn't work


Paper suggests asking people to guess with **what probability** they prefer X over Y. And Y over X.

Website User Survey


1. The website has a user friendly interface.




2. The website is easy to navigate.




3. The website's pages generally have good images.



4. The website allows users to upload pictures easily.



5. The website has a pleasing color scheme.



#6: Choice of task: Proxy task (left) doesn't correlate with actual task (right)

The actual task:

- Is there >30% fat?

AI predicts binary (yes/no) answer

Is 30% or more of the nutrients on this plate fat?



NO, 30% or more of the nutrients on this plate is not fat.

What is your decision?

NO, 30% of the nutrients on this plate is not fat.

YES, 30% of the nutrients on this plate is fat.

#6: Choice of task: Proxy task (left) doesn't correlate with actual task (right)

The actual task:

- Is there >30% fat?

AI predicts binary (yes/no) answer

AI can produce explanations in the form of exemplars.

Is 30% or more of the nutrients on this plate fat?



This AI recommended answer is:

NO, 30% or more of the nutrients on this plate is not fat.

What is your decision?

#6: Choice of task: Proxy task (left) doesn't correlate with actual task (right)


The actual task:

- Is there >30% fat?

AI predicts binary (yes/no) answer

AI can produce explanations in the form of **detected concepts**.

Is 30% or more of the nutrients on this plate fat?



Here are ingredients the AI recognized as main nutrients which make up 30% or more fat on this plate:

salmon
avocado

This AI recommended answer is:

YES, 30% or more of the nutrients on this plate is fat.

What is your decision?

NO, 30% of the nutrients on this plate is not fat.

YES, 30% of the nutrients on this plate is fat.

The proxy task: What do you think the AI will choose?

The AI must decide: Is 30% or more of the nutrients on this plate fat?

Fact: 30% or more of the nutrients on this plate is not fat.



Here are examples of plates that the AI knows the fat content of and categorizes as similar to the one above:



What will the AI decide?

NO, 30% of the nutrients on this plate is not fat.

YES, 30% of the nutrients on this plate is fat.

Is 30% or more of the nutrients on this plate fat?



Here are examples of plates that the AI categorizes as similar to the one above and do not have 30% or more fat:



This AI recommended answer is:


NO, 30% or more of the nutrients on this plate is not fat.

What is your decision?

NO, 30% of the nutrients on this plate is not fat.

YES, 30% of the nutrients on this plate is fat.

#6: Choice of task: Proxy tasks don't correlate with actual task



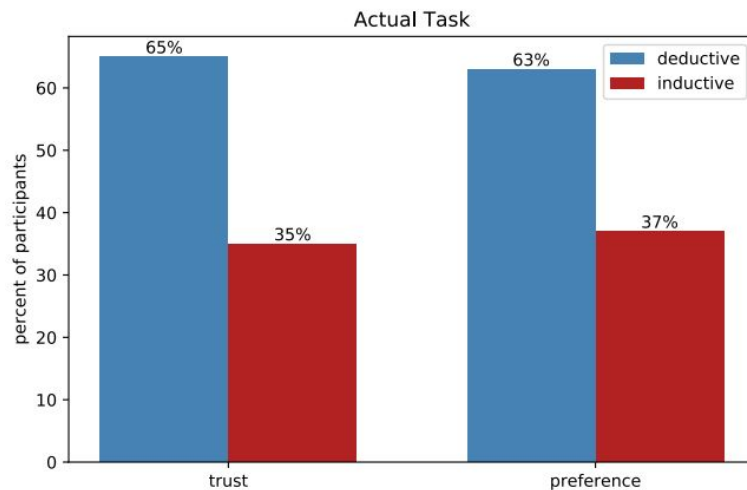
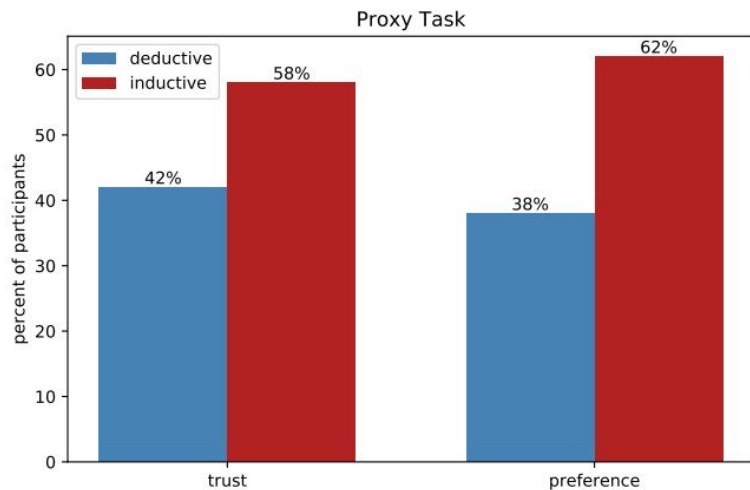
Deductive explanations = detected concepts

Use that information to deduce the answer





Inductive explanations: exemplars

Use general patterns from other examples


#6: Choice of task: Proxy tasks don't correlate with actual task





#7: Unfaithful explanations: Saliency maps




schooner



African elephant, *Loxodonta africana*




go-kart



Which pixels explain the prediction? Saliency via backprop


Forward pass: Compute probabilities




Dog

Which pixels explain the prediction? Saliency via backprop




Forward pass: Compute probabilities




Compute gradient of (unnormalized) class score with respect to image pixels, take absolute value and max over RGB channels




Which pixels explain the prediction? Saliency via backprop




Saliency maps were getting quite popular




#7: Unfaithful explanations: random predictions don't change explanations




#7: Unfaithful explanations: randomizing last two layers don't change explanations



#7: Unfaithful explanations: random networks induce the same explanations

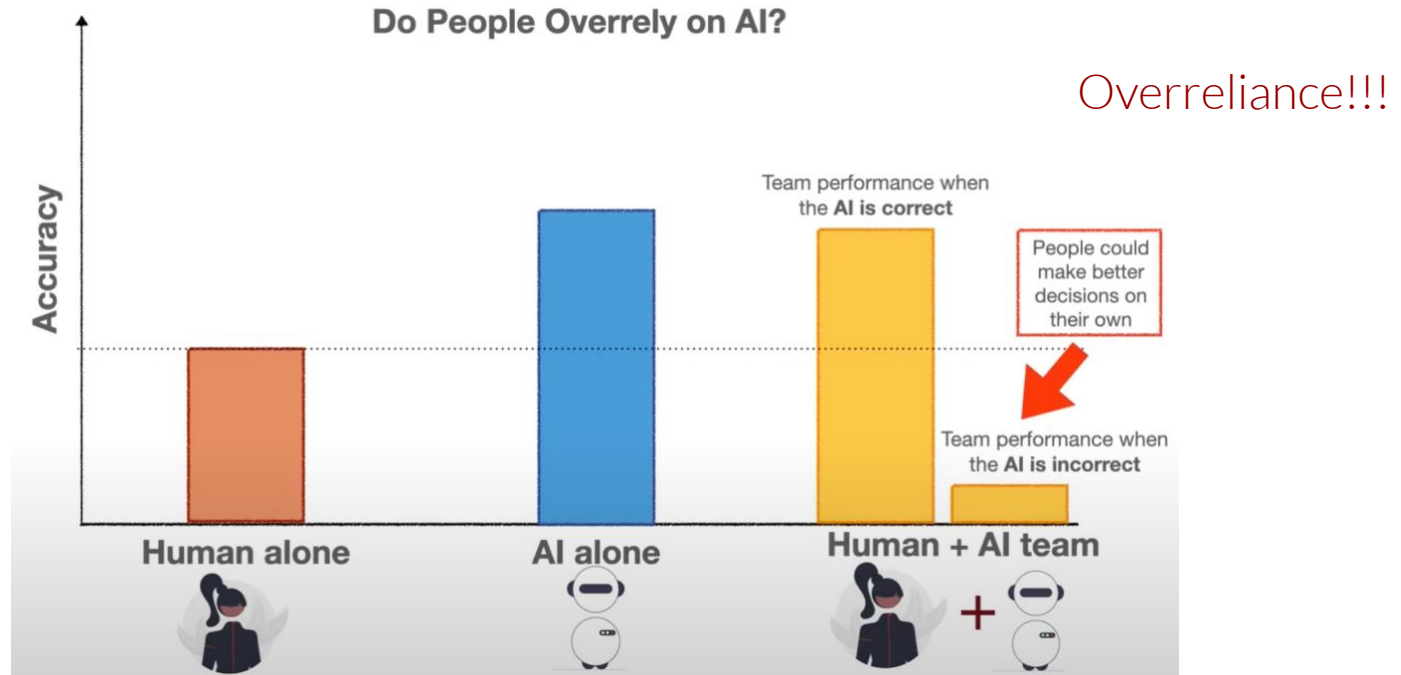


#8: Faithful explanations may still hurt decision making



Bucinca et al. To Trust or to Think: Cognitive Forcing Functions Can Reduce Overreliance on AI in AI-assisted Decision-making. CSCW 2021

#8: Faithful explanations may still hurt decision making



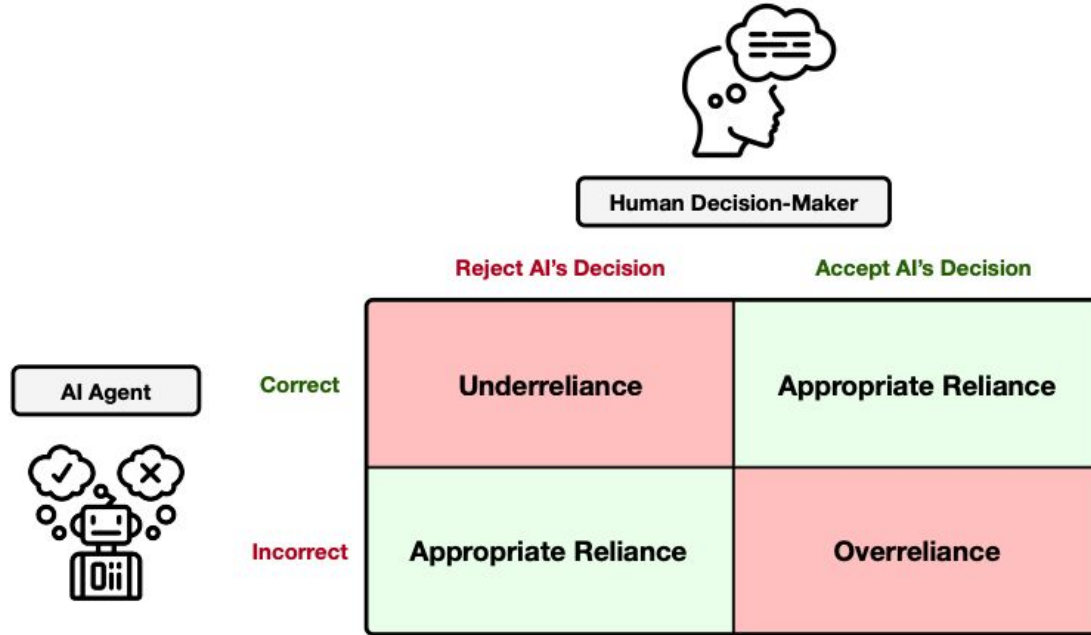
Bucinca et al. To Trust or to Think: Cognitive Forcing Functions Can Reduce Overreliance on AI in AI-assisted Decision-making. CSCW 2021

Deep Dive:

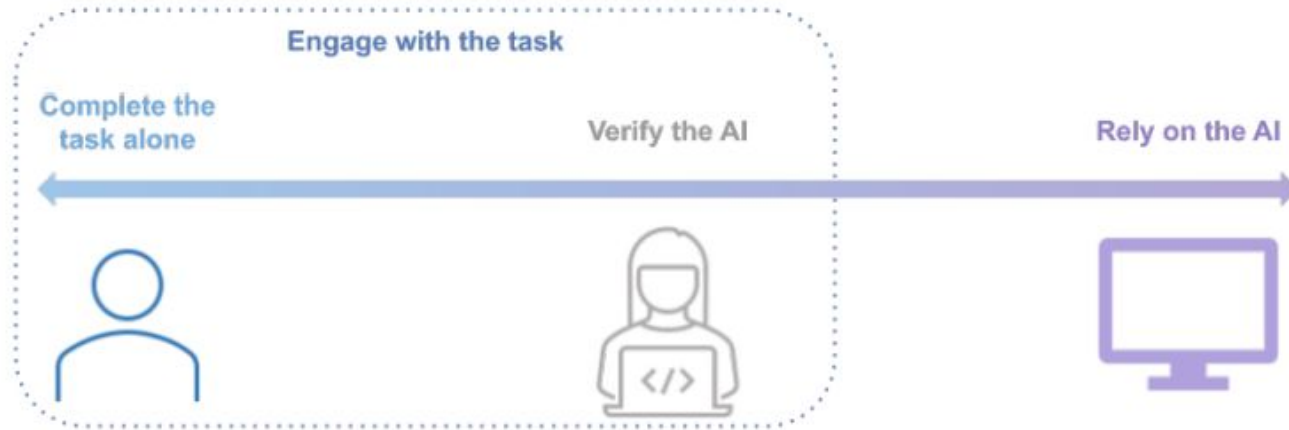
Research question:

Can explanations reduce overreliance on AI-assisted decision making?

What is overreliance?



Two prototype strategies in which people engage with explanations



Predominant hypothesis for overreliance

Cognitive biases

- Mere presence of explanations increase trust.
- Trust makes us overrely.

There are cases when we do engage with explanations

- Incorrect email auto-replies
- GPS navigation system showing you the wrong route
- What else have you encountered?

Why don't explanations help in these tasks?

The AI must decide: Is 30% or more of the nutrients on this plate fat?

Fact: 30% or more of the nutrients on this plate is not fat.



Here are examples of plates that the AI knows the fat content of and categorizes as similar to the one above:



What will the AI decide?

NO, 30% of the nutrients on this plate is not fat.

YES, 30% of the nutrients on this plate is fat.

The Lophotrochozoa, evolved within Protostomia, include two of the most successful animal phyla, the Mollusca and Annelida. The former, which is the second-largest animal phylum by number of described species, includes animals such as snails, clams, and squids, and the latter comprises the segmented worms, such as earthworms and leeches. These two groups have long been considered close relatives because of the common presence of trochophore larvae, but the annelids were considered closer to the arthropods because they are both segmented. Now, this is generally considered convergent evolution, owing to many morphological and genetic differences between the two phyla. The Lophotrochozoa also include the Nemertea or ribbon worms, the Sipuncula, and several phyla that have a ring of ciliated tentacles around the mouth, called a lophophore. These were traditionally grouped together as the lophophorates. but it now appears that the lophophorate group may be paraphyletic, with some closer to the nemerteans and some to the molluscs and annelids. They include the Brachiopoda or lamp shells, which are prominent in the fossil record, the Entoprocta, the Phoronida, and possibly the Bryozoa or moss animals.

What are some of the animals in Annelida?

AI's Suggestion:

Snails, clams, and squids


- Memerteans
- Ribbon worms
- Earthworms and leeches
- Snails, clams, and squids

Submit




A cost-benefit framework

Costs increase overreliance

Benefits decrease

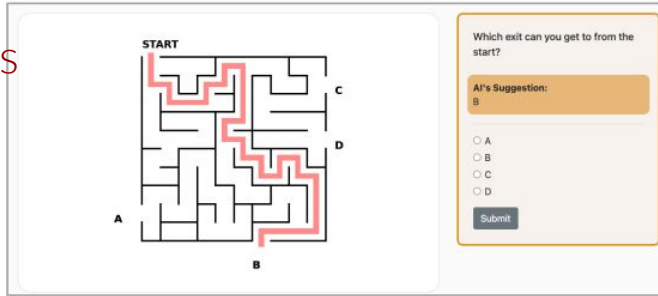


Designed tasks that increase in cognitive effort



Explanations that take different cognitive effort

Highlight explanations



START

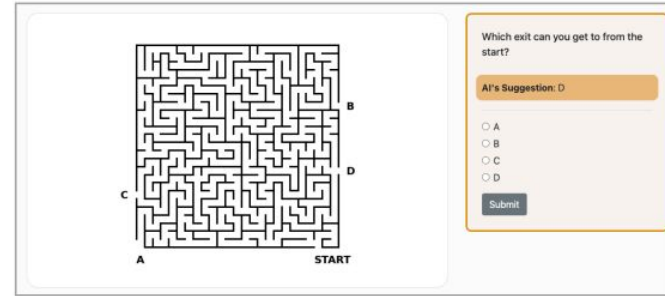
Which exit can you get to from the start?

AI's Suggestion: B

A
 B
 C
 D

Submit

A B C D



Which exit can you get to from the start?

AI's Suggestion: D

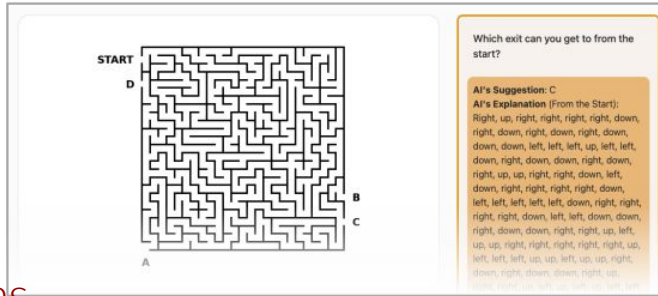
A
 B
 C
 D

Submit

A B C D START

No explanations

Written explanations



START

Which exit can you get to from the start?

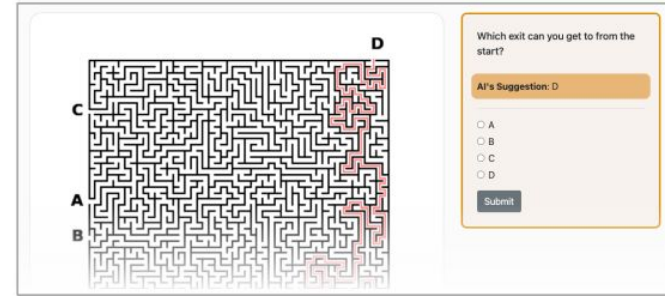
AI's Suggestion: C

AI's Explanation (From the Start):
Right, up, right, right, right, right, down, right, down, right, down, right, down, down, down, left, left, up, left, left, down, right, down, right, down, left, down, right, right, right, right, down, left, left, left, left, left, down, right, right, right, right, down, left, left, down, down, right, up, left, up, up, right, right, right, right, right, up, left, left, left, up, up, left, up, up, right, down, right, down, down, right, up, down, right, up, left, up, left, up

A
 B
 C
 D

Submit

A B C D



Which exit can you get to from the start?

AI's Suggestion: D


A
 B
 C
 D

Submit

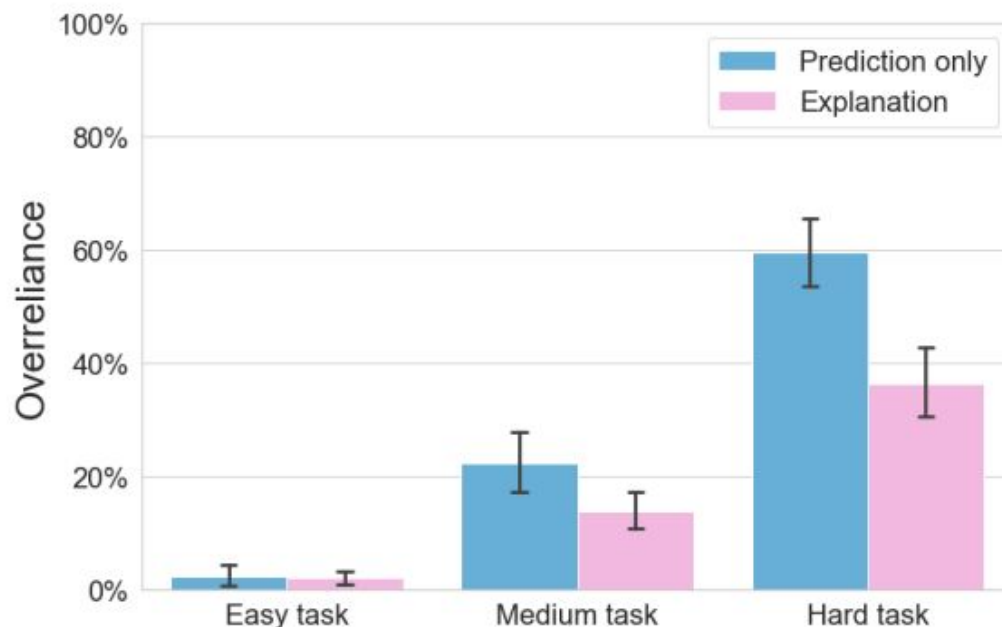
A B C D

Highlight explanations for hard tasks


Highlights reduce cognitive effort to find AI errors



We show for the first time that **explanations do reduce overreliance in human-AI decision making** but only when the task difficulty is high enough to require explanations




If explanations take effort to understand, overreliance increases

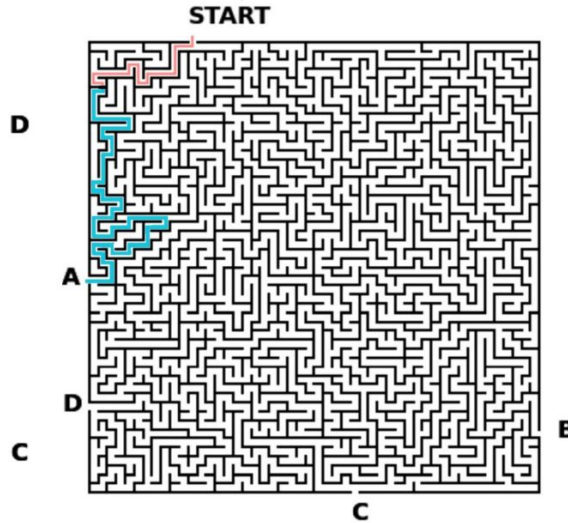


Adding two a new type of explanations:

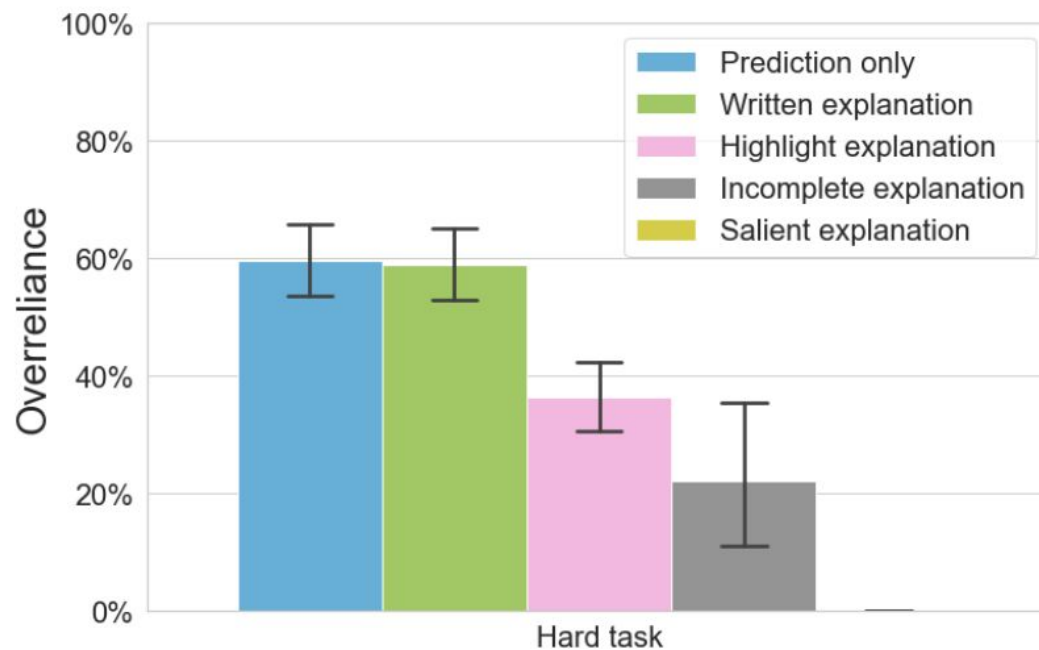
Incomplete
explanations




Salient
explanations



Less cognitive effort -> less overreliance



More benefit -> less overreliance



Challenges with evaluation protocols for human-AI systems

