

Introduction

CSEN266
Artificial Intelligence

What is AI?

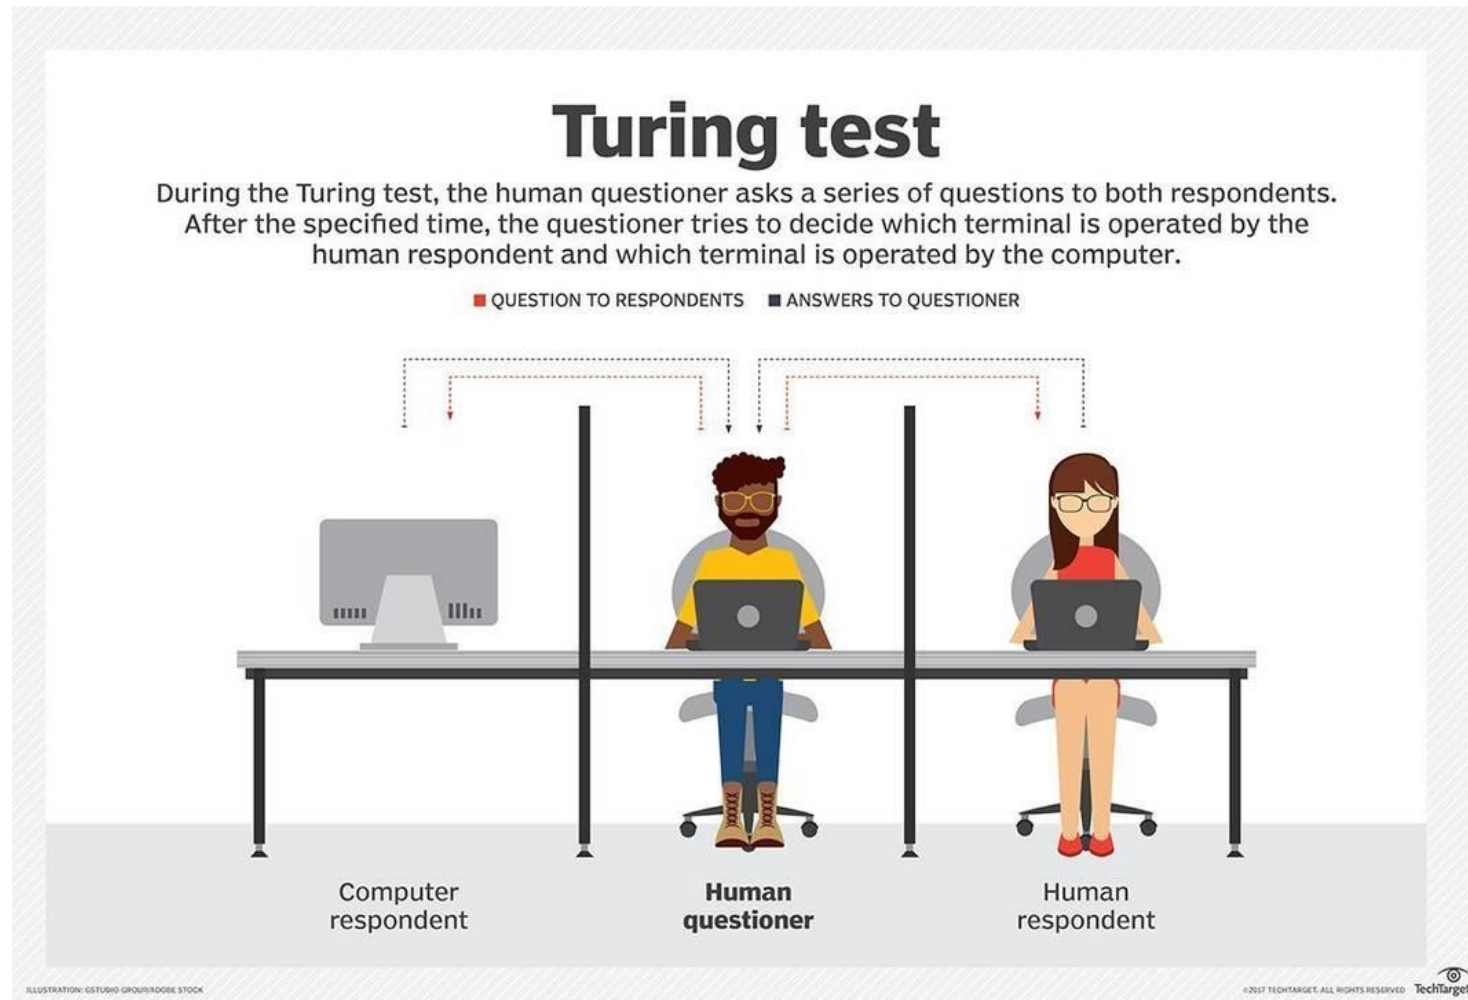
- What is Artificial Intelligence?
- Artificial?
 - Computers, machines, ...
- Intelligence?

Turing Test

- How to evaluate that a machine system is intelligent?
- Alan Turing created the Turing Test (1950-1952)
- **Basic idea**
 - I don't have a good definition of “intelligence”
 - However, I know humans are intelligent
 - If an entity (e.g. computer, machine) can masquerade as human so well that I can't distinguish the difference, I will say this entity is intelligent



Turing Test



If 30% of human questioners fail to distinguish a machine from humans, then this machine is considered to have passed the test

Turing Test

- [The Loebner Prize](#): the annual competition on Turing Test (launched in 1990)
- Prediction: By 2000, 30% chance of a machine fooling a person for 5 mins
- First AI that passed the Turing Test
 - Eugene Goostman (2014)



Loebner Prize:
awarded annually
to best AI program

Examples of Artificial Intelligence

- Examples of AI
 - Computer Games
 - Robotics
 - Smart Devices
 - ...
- What capabilities do they have?

Game Playing: AlphaGo

- AlphaGo beat Korean grandmaster Lee Sedol (2016)
 - Google DeepMind
 - 4 games to 1
- Capabilities:
 - automatic reasoning, decision making



Chatbots

- ELIZA (1966): the first chatbot in the history
 - To mimic psychotherapist
 - Convert declarative sentences to questions
 - Does not understand the semantics

Conversation with ELIZA:

You: I feel bad these days.

Eliza: Do you often feel bad these days?

You: Yes, but I don't know how to make myself happy.

Eliza: Why do you not know what to make yourself happy?

You: Will you tell me something?

Eliza: Oh, I tell you something.

You: Tell me then.

Eliza: Can you elaborate on that?

You: No, I just feel unhappy. Maybe I am worried about something.

Eliza: How long have you been worried about something?

- Capabilities:
 - Natural language understanding

Chatbots



Valerie

- Valerie: CMU Robot Receptionist
- Capabilities:
 - Natural language understanding
 - Perception

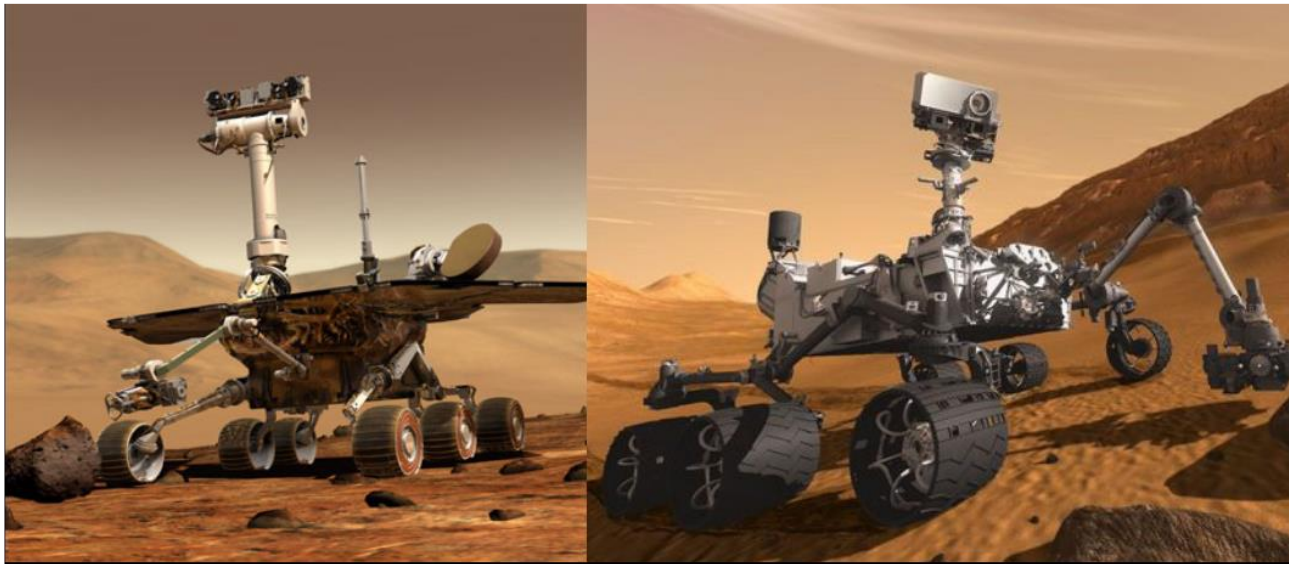
Cleaning Robots

- iRobot
- Roomba: Autonomous home vacuum cleaner
- **Capabilities:**
 - mobility/robotics (can move about), perception (achieved by infra-red sensors to detect walls and obstacles), planning



Space Exploration Robots

- Mars-rovers: motor vehicles that travel across the surface of Mars, searching for evidence of ancient life, water, and minerals
- **Capabilities:**
 - visual perception, robotics (manipulate objects and move about), planning, learning (adapt to new circumstances)



Machine Translation

- The use of software to translate text or speech from one language to another.
- The famous Russian saying: “The spirit is willing but the flesh is weak.”
 - Actual meaning: someone has good intentions but fails to live up to them.
- Early days machine translation:
 - The vodka is good but the meat is rotten.
- Capabilities:
 - Natural language understanding

Virtual Assistant

- Apple Siri, Amazon Alexa
- **Capabilities:**
 - Speech recognition, natural language understanding



Face ID

- Unlocks the phone using infrared and visible light images to uniquely identify your face
- **Capabilities:**
 - Visual perception, image processing, learning (learns a pattern such that it can recognize a newly captured face image)



What Capabilities Does AI Have?

- Capabilities of AI

- Reasoning, decision making
- Learning, planning
- Visual recognition
- Robotics
- Speech recognition
- Natural language understanding

- How to summarize these capabilities?

- Artificial Intelligence is the science that makes machines to think humanly, act humanly, think rationally, and act rationally.

Thinking Humanly

- Make machines with “minds”
 - Such that they can make decisions, solve problems, and learn things
- Examples
 - Play Go games, solve math problems, visual recognition (image classification, object detection, etc.), speech recognition, natural language understanding

Thinking Humanly

- **Strong AI vs Weak AI**
 - **Strong AI:** artificial intelligence can have a mind in exactly the same sense human beings have minds.
 - **Weak AI:** machines can demonstrate intelligence, but do not necessarily have a mind, mental states or consciousness.

Acting Humanly

- AI: build machines that perform human tasks automatically
- Tasks that people do? e.g.
 - Sing and dance
 - Play soccer
 - Do surgeries
 - Navigate in a building
 - Assemble parts
 - ...

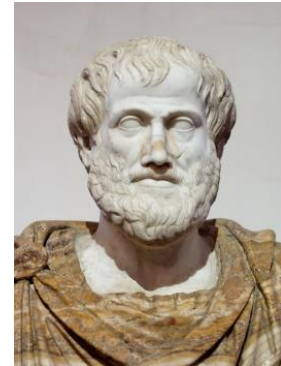
Thinking rationally

- Make machines that use logical reasoning and/or computation to solve problems, make decisions, and learn things.

- Example of logical reasoning:

- Aristotle's syllogisms:

“Socrates is a man; all men are mortal;
therefore, Socrates is mortal”.



Acting rationally

- Rational behavior

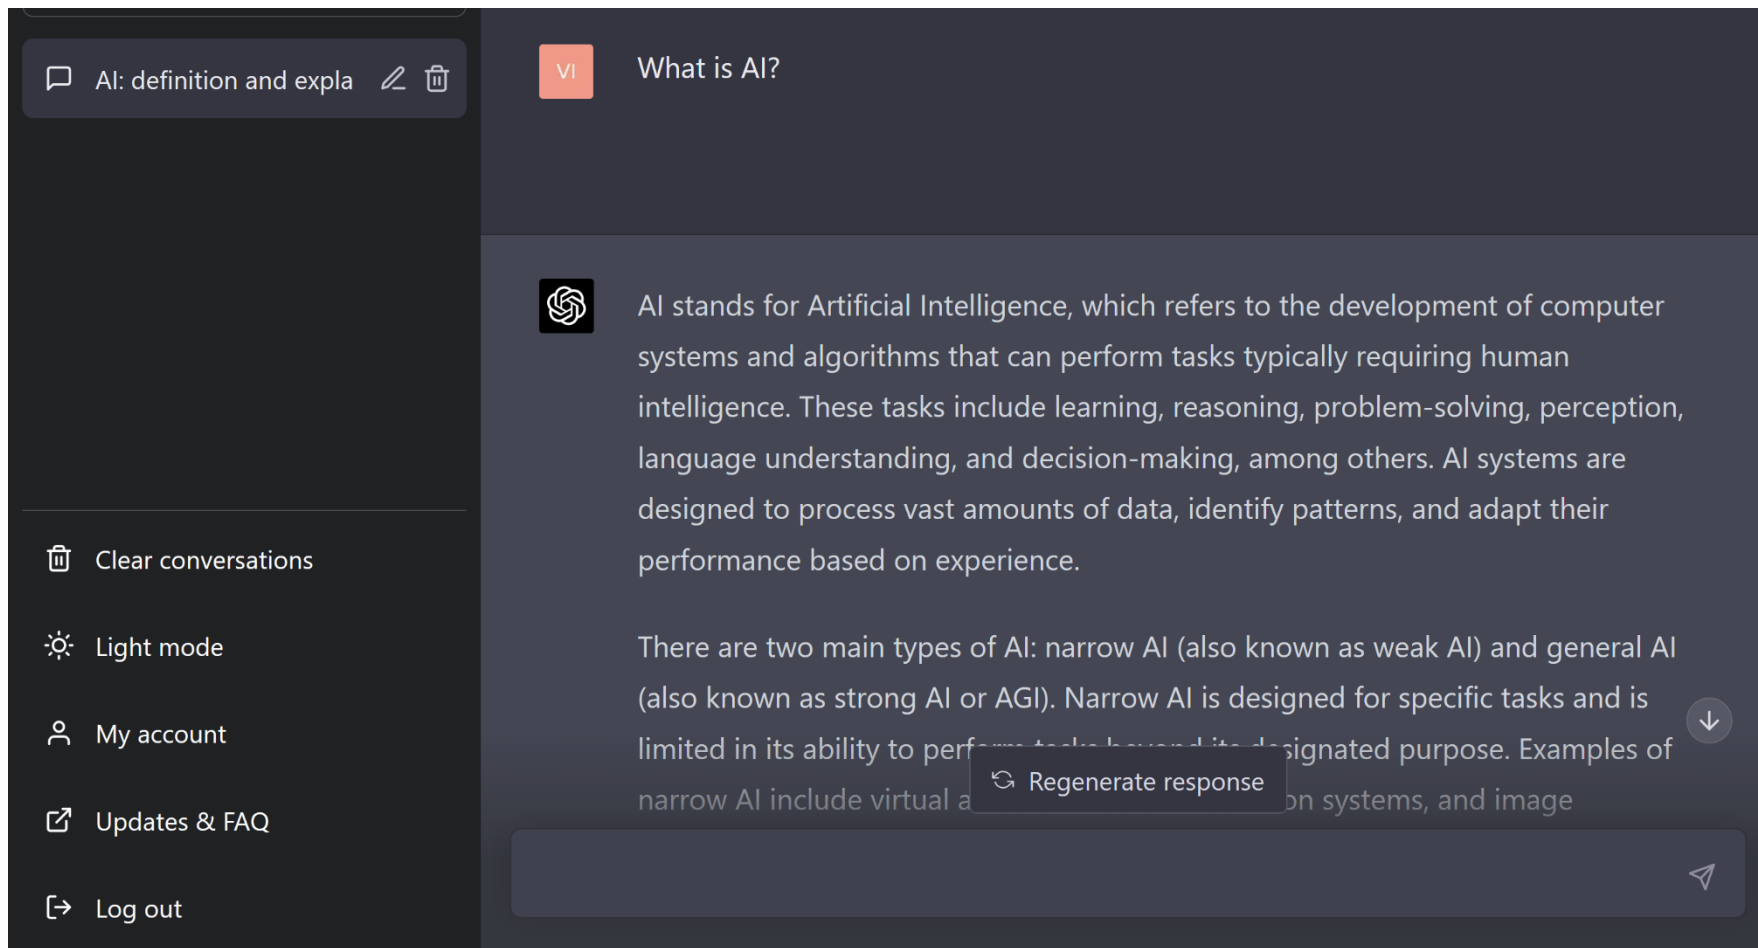
- The behavior/action that achieves the goal

Example: a route-planning app can achieve the goal of selecting the route from one city to another.

- Achieve the best performance, given available information

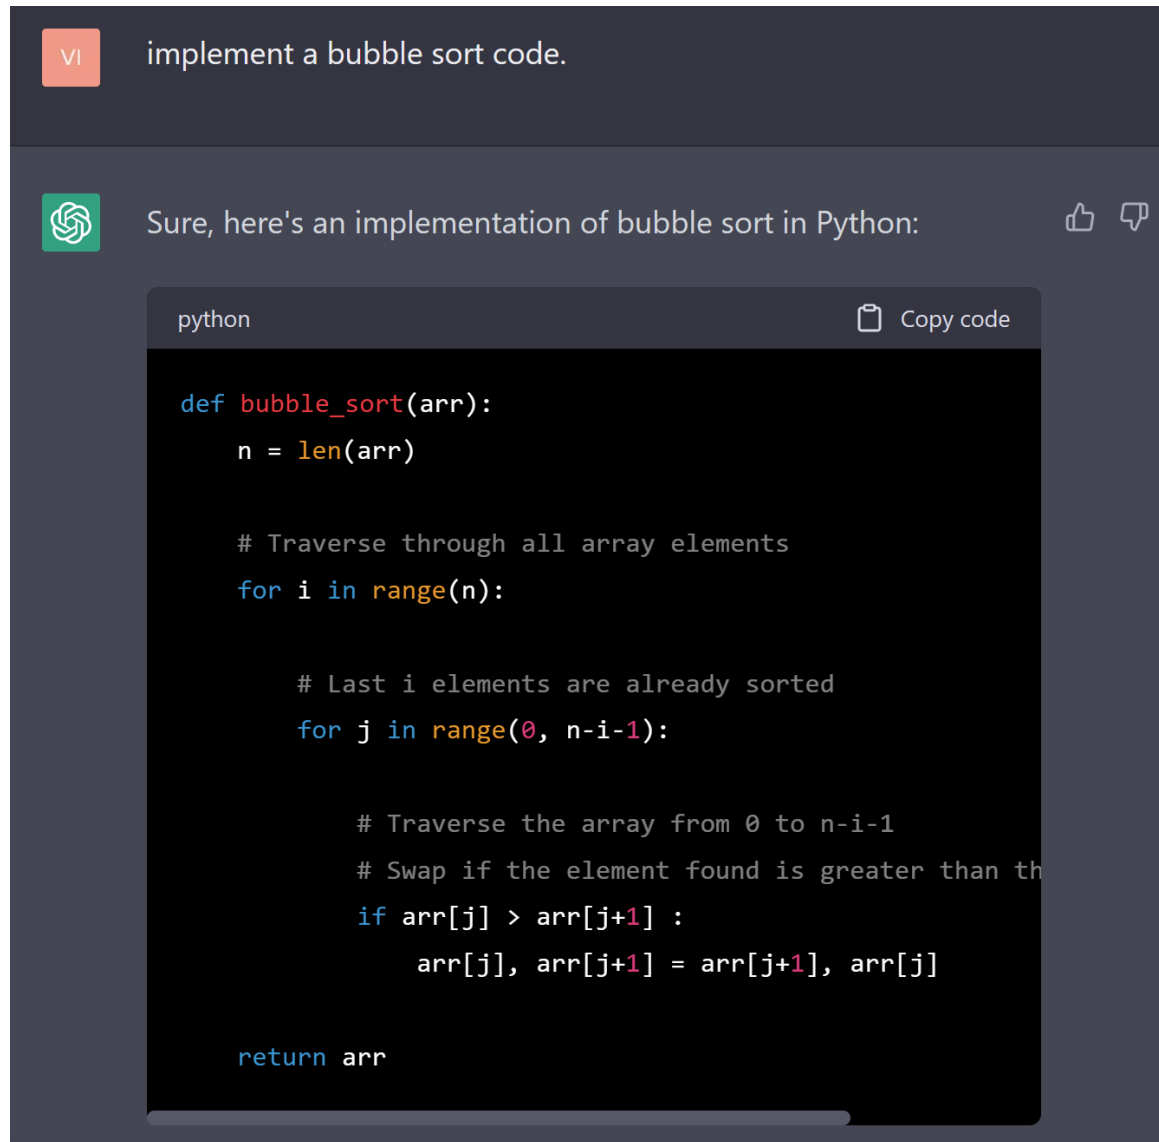
Example: When selecting the route, the app can minimize the path cost

The State of the Art: ChatGPT






The screenshot displays the ChatGPT web interface. On the left is a dark sidebar with navigation options: 'Clear conversations', 'Light mode', 'My account', 'Updates & FAQ', and 'Log out'. The main chat area has a dark background. At the top, a user prompt reads 'What is AI?'. Below it, the AI assistant's response is shown, starting with the OpenAI logo and the text: 'AI stands for Artificial Intelligence, which refers to the development of computer systems and algorithms that can perform tasks typically requiring human intelligence. These tasks include learning, reasoning, problem-solving, perception, language understanding, and decision-making, among others. AI systems are designed to process vast amounts of data, identify patterns, and adapt their performance based on experience.' A second paragraph begins: 'There are two main types of AI: narrow AI (also known as weak AI) and general AI (also known as strong AI or AGI). Narrow AI is designed for specific tasks and is limited in its ability to perform tasks beyond its designated purpose. Examples of narrow AI include virtual a...'. A 'Regenerate response' button is visible over the second paragraph. At the bottom of the chat area is a text input field with a send button on the right.

The State of the Art: ChatGPT



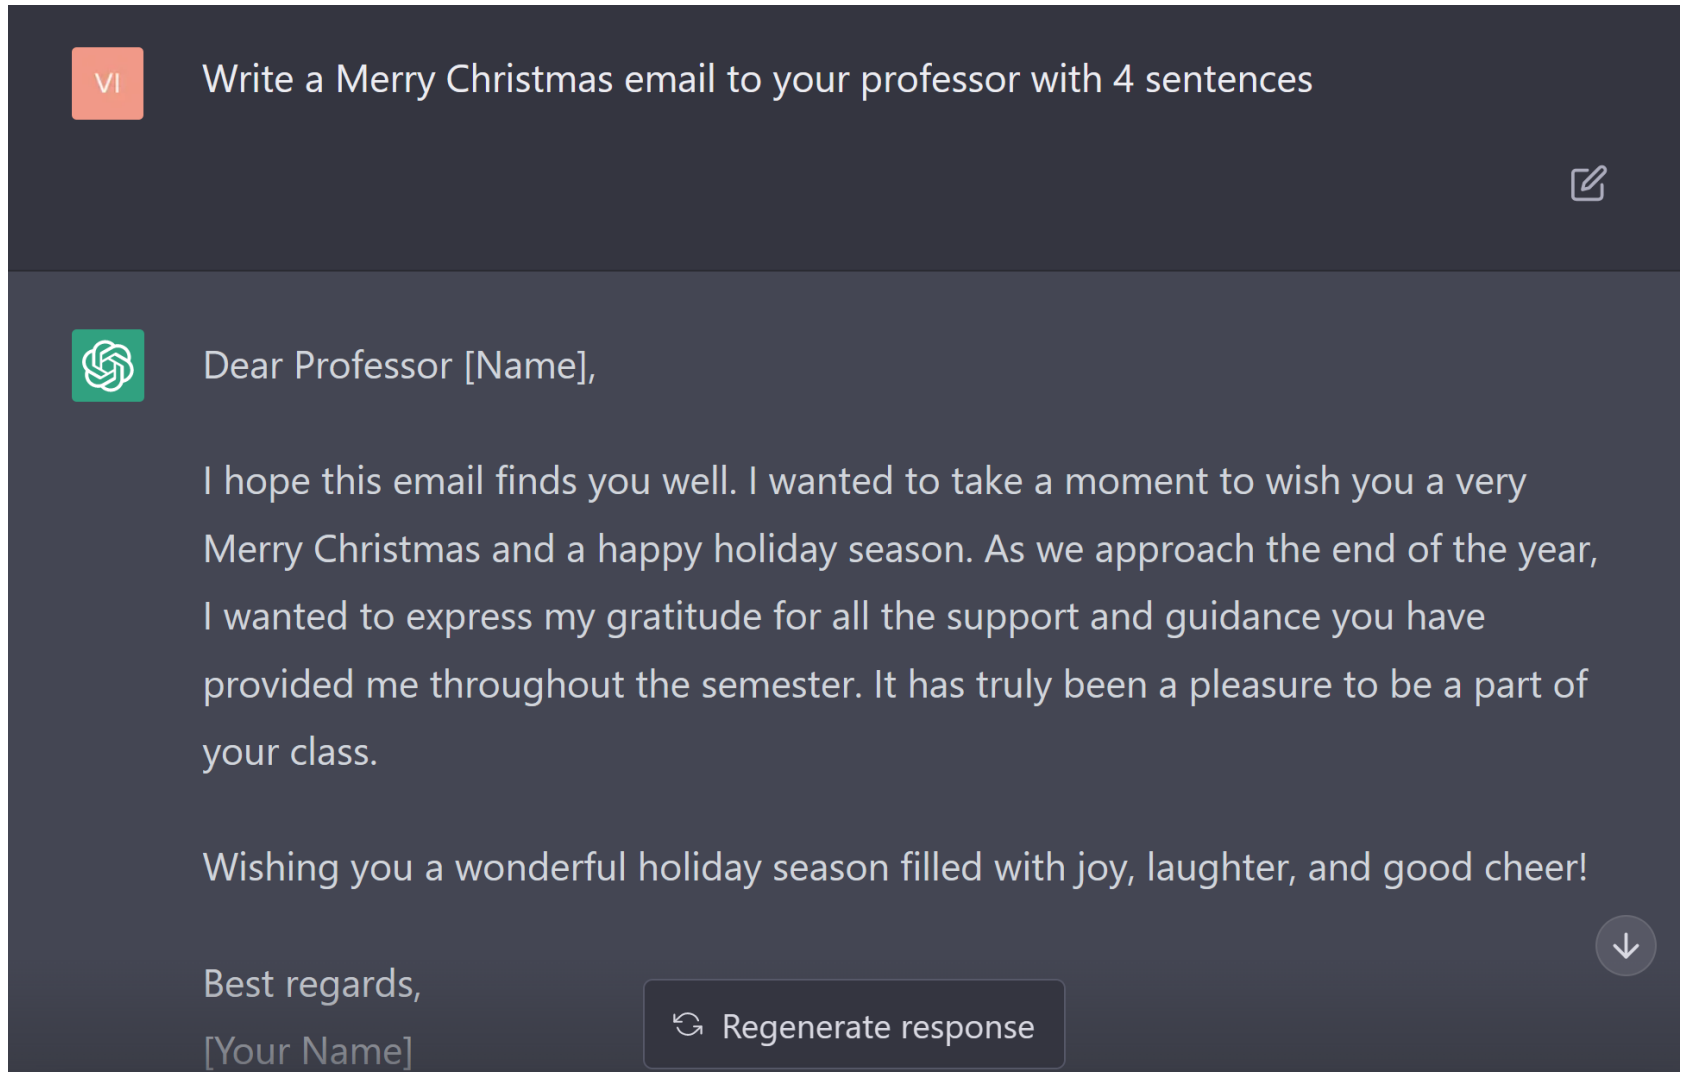
VI implement a bubble sort code.

 Sure, here's an implementation of bubble sort in Python:  

```
python Copy code
```

```
def bubble_sort(arr):  
    n = len(arr)  
  
    # Traverse through all array elements  
    for i in range(n):  
  
        # Last i elements are already sorted  
        for j in range(0, n-i-1):  
  
            # Traverse the array from 0 to n-i-1  
            # Swap if the element found is greater than th  
            if arr[j] > arr[j+1] :  
                arr[j], arr[j+1] = arr[j+1], arr[j]  
  
    return arr
```

The State of the Art: ChatGPT



The screenshot displays a ChatGPT interface with a dark background. At the top, a user prompt is shown in a dark grey box with a red square icon containing the number 'VI'. The prompt reads: "Write a Merry Christmas email to your professor with 4 sentences". To the right of the prompt is a white edit icon. Below the prompt, the AI's response is shown in a lighter grey box, starting with a green square icon containing the OpenAI logo. The response is a four-sentence email. At the bottom of the response area, there is a "Regenerate response" button with a circular arrow icon and a downward arrow icon in a circle.

VI Write a Merry Christmas email to your professor with 4 sentences

Dear Professor [Name],

I hope this email finds you well. I wanted to take a moment to wish you a very Merry Christmas and a happy holiday season. As we approach the end of the year, I wanted to express my gratitude for all the support and guidance you have provided me throughout the semester. It has truly been a pleasure to be a part of your class.

Wishing you a wonderful holiday season filled with joy, laughter, and good cheer!

Best regards,
[Your Name]

Regenerate response

The State of the Art: ChatGPT

ChatGPT

Input

Andrew is free from 11 am to 3 pm, Joanne is free from noon to 2 pm and then 3:30 pm to 5 pm. Hannah is available at noon for half an hour, and then 4 pm to 6 pm. What are some options for start times for a 30 minute meeting for Andrew, Hannah, and Joanne?

Output

Here is a list of their availability:

- Andrew: 11 am to 3 pm
- Joanne: noon to 2 pm, and 3:30 pm to 5 pm
- Hannah: noon to 12:30 pm, and 4 pm to 6 pm

Based on their availability, there is a 30-minute window where all three of them are available, which is from 4 pm to 4:30 pm. So, the meeting can be scheduled at 4 pm.

GPT-4

Input

Andrew is free from 11 am to 3 pm, Joanne is free from noon to 2 pm and then 3:30 pm to 5 pm. Hannah is available at noon for half an hour, and then 4 pm to 6 pm. What are some options for start times for a 30 minute meeting for Andrew, Hannah, and Joanne?

Output

Andrew: 11 am - 3 pm

Joanne: 12 pm - 2 pm, 3:30 pm - 5 pm

Hannah: 12 pm - 12:30 pm, 4 pm - 6 pm

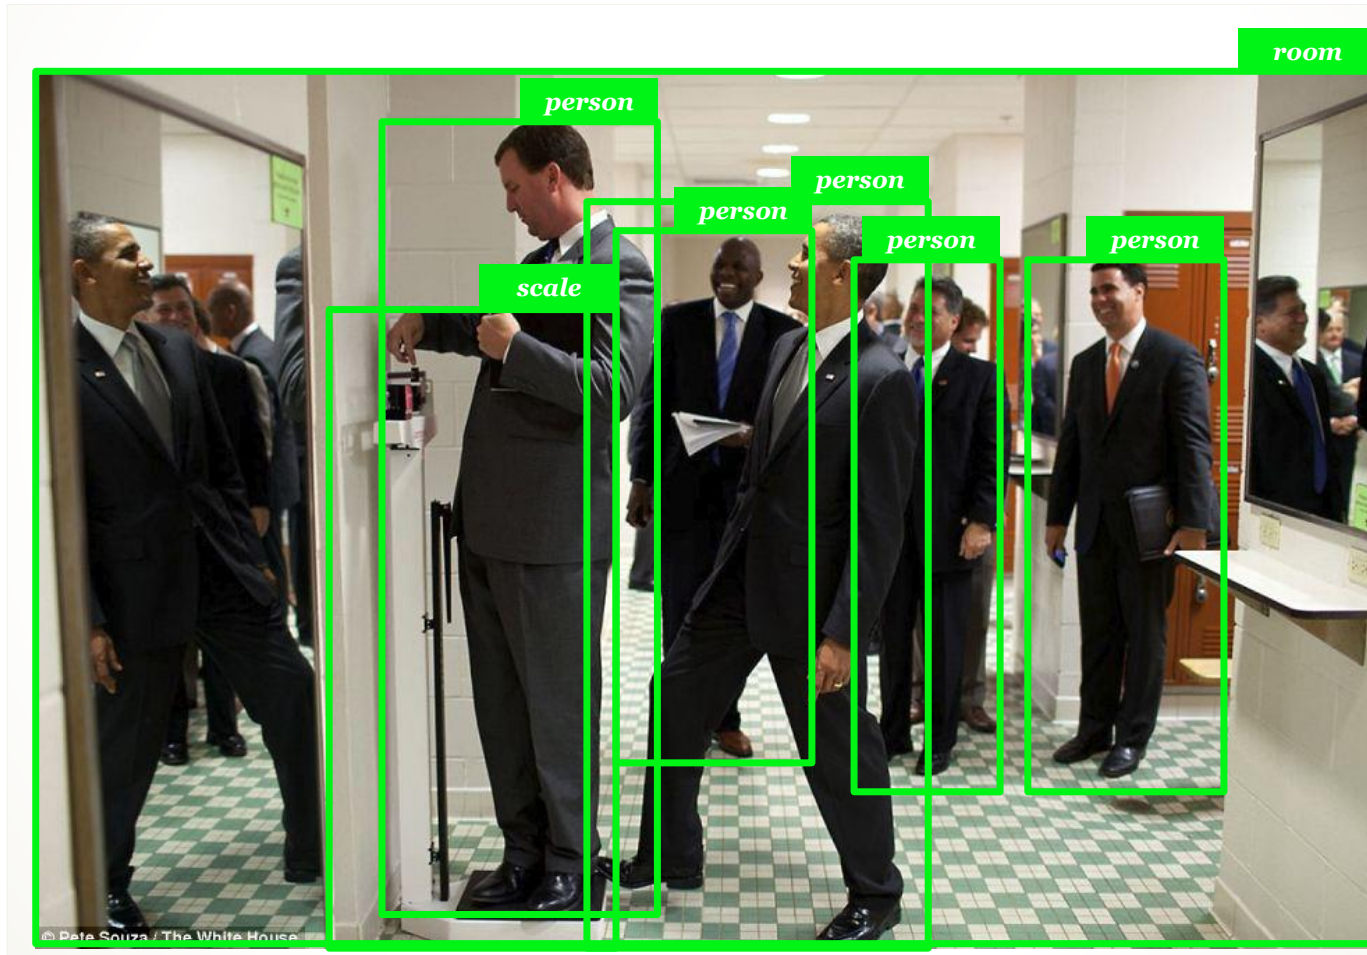
Common availability for a 30-minute meeting: 12 pm - 12:30 pm

The State of the Art: Visual Recognition

- ImageNet
 - A large image database for visual recognition research, such as image classification and object detection.
- ImageNet Large Scale Visual Recognition Challenge (ILSVRC)
 - An annual contest (2010-2017)
 - Breakthrough in 2012 (AlexNet): the beginning of the deep learning revolution

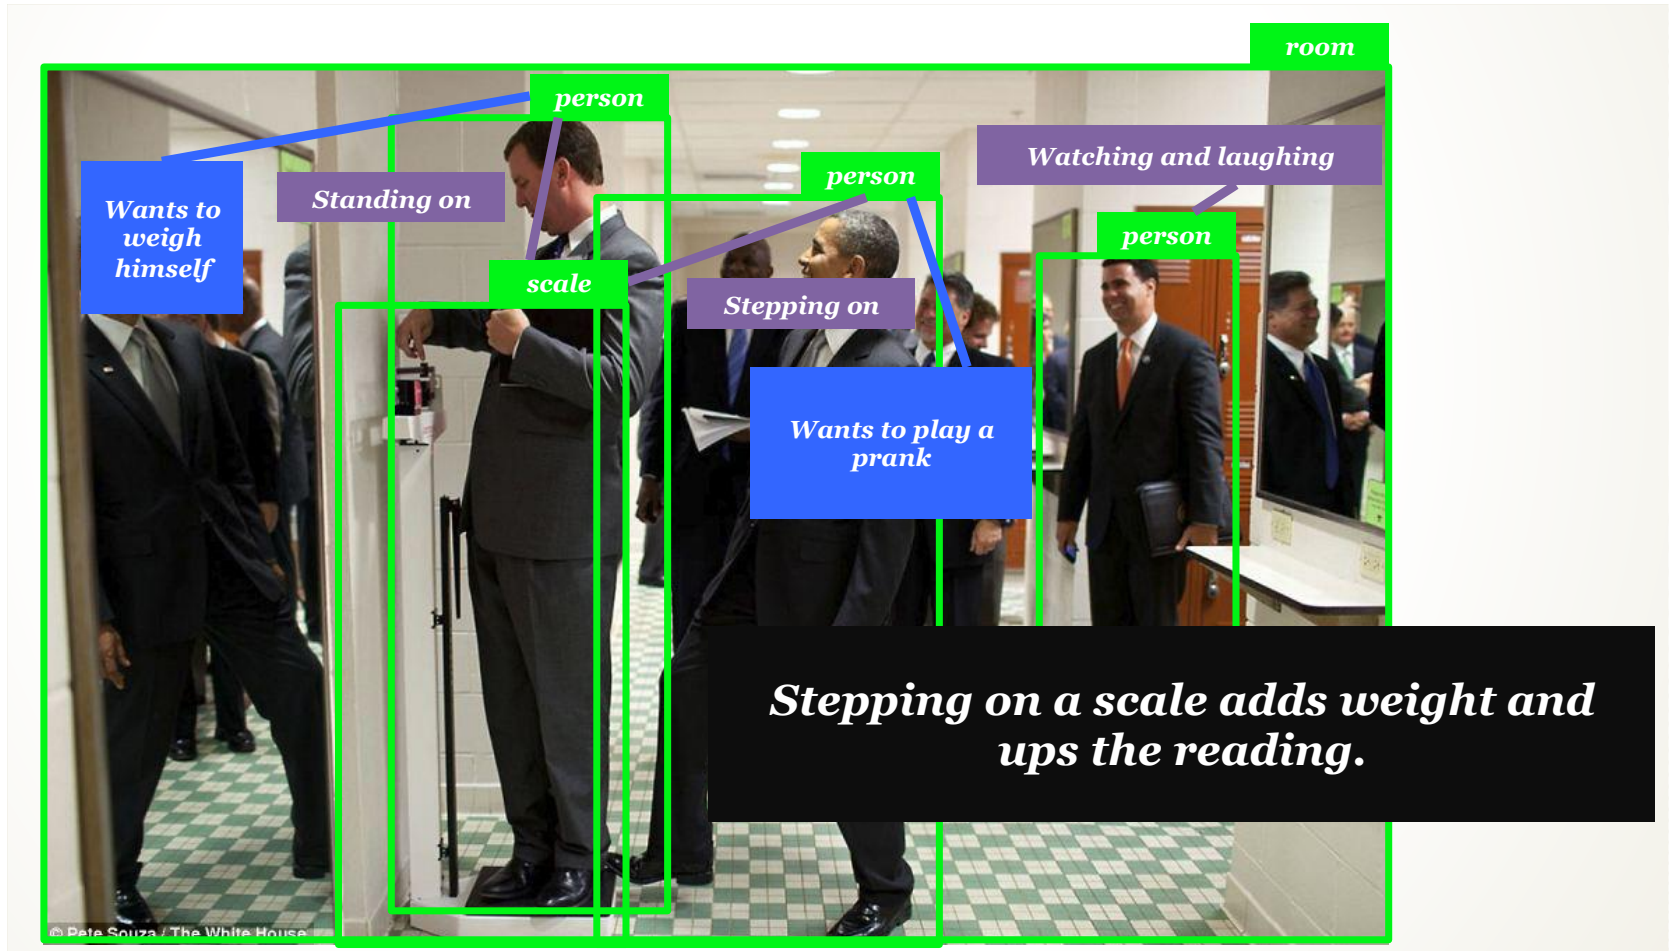
The State of the Art: Visual Recognition

- Moving from object detection...



The State of the Art: Visual Recognition

- ... to human-level understanding



The State of the Art: Visual Recognition

- Image Captioning (Success)



A cat sitting on a suitcase on the floor



A cat is sitting on a tree branch



A dog is running in the grass with a frisbee



A white teddy bear sitting in the grass



Two people walking on the beach with surfboards



A tennis player in action on the court



Two giraffes standing in a grassy field



A man riding a dirt bike on a dirt track

The State of the Art: Visual Recognition

- Image Captioning (Failure)



A woman is holding a cat in her hand



A person holding a computer mouse on a desk



A woman standing on a beach holding a surfboard



A bird is perched on a tree branch



A man in a baseball uniform throwing a ball

The State of the Art

- Autonomous Vehicles
- Level 0
- No automation
- The driver performs all driving tasks



The State of the Art

- Autonomous Vehicles
- Level 1
- Driver assistance
- Vehicle is controlled by the driver
- Driving assist features
 - Anti-lock Breaking System (ABS)
 - Adaptive Cruise Control (ACC)
 - Lane Keeping Assistance...



The State of the Art

- Autonomous Vehicles
- Level 2
- Partial automation
- Vehicle has combined automated functions
 - Acceleration
 - Steering
- Driver must remain engaged with the driving task and monitor the environment at all times



The State of the Art

- Autonomous Vehicles
- Level 3
- Conditional automation
- Driver is a necessity, but is not required to monitor the environment.
- The driver must be ready to take control of the vehicle with notice.



The State of the Art

- Autonomous Vehicles
- Level 4
- High automation
- The vehicle is capable of performing all driving functions **under certain conditions**
 - e.g. only in limited spatial areas ([geofenced](#))
 - e.g. under special circumstances, like traffic jams
- No take over request
- The driver need to control the vehicle outside these areas or circumstances



The State of the Art

- Autonomous Vehicles
- Level 5
- Full automation
- Driverless
- The vehicle is capable of performing all driving functions under all conditions
- e.g. a robotic taxi



Discussion

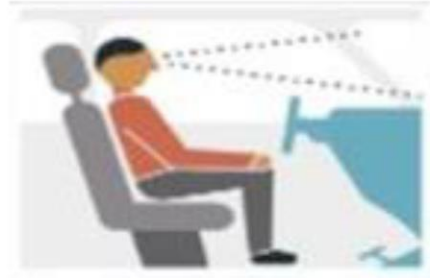
- What's the relation between **rational thinking** and **rational acting**? Is **making decisions** more important, or **taking actions** more important?
- Which **level transition** of autonomous vehicles do you think is the most important? (the transition from which level to which level?)



Level 0 – No automation



Level 1 – Driver Assistance



Level 2 – Partial Automation



Level 3 – Conditional Automation



Level 4 – High Automation



Level 5 – Full Automation

The State of the Art

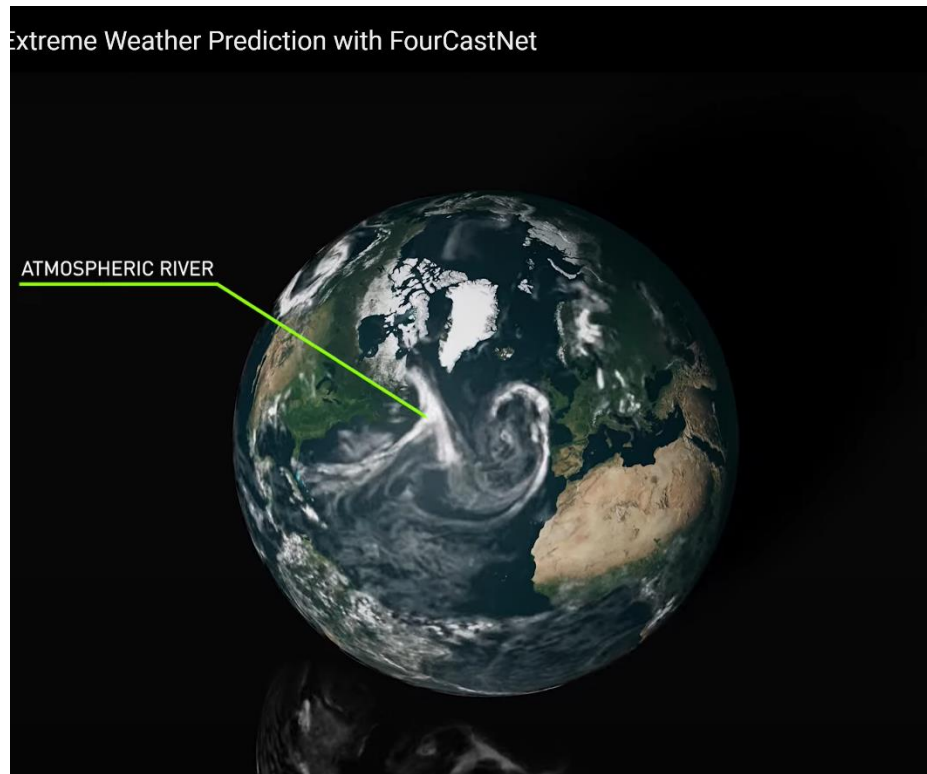
- Currently, Level 4 is achieved
- Level 5 still has a long way to go
- Example



The State of the Art: Weather Forecasting

- FourCastNet

- https://www.youtube.com/watch?v=nuT_U1AQz3g



GPU accelerated weather forecast model

Predict extreme weather events such as hurricanes and atmospheric rivers

The model was trained on 10TB of Earth system data

It is five orders of magnitude faster than traditional weather forecasting models

The State of the Art: Character Animation

- AI-powered character animated using reinforcement learning.
- <https://www.youtube.com/watch?v=8oIQy6fxfCA>

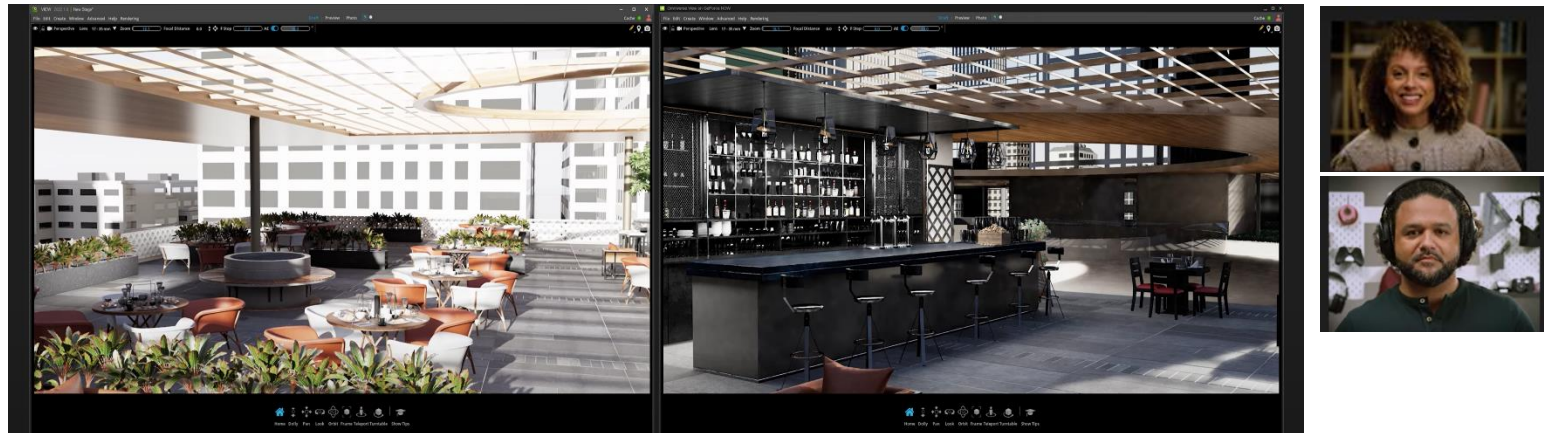


The character learns to perform life-like motions by imitating human motion data, such as walking, running, and sword swings (we can view it as a virtual robot who learns to act humanly)

Trained by massively parallel GPU simulator which only takes 3 days

The State of the Art: Omniverse

- Omniverse: a virtual platform for collaborative 3D design
- https://www.youtube.com/watch?v=PUC_Dgo8U6I



An architecture design team reviews a project using Omniverse in a web conference.

Discussion

- What tasks are humans better at? What tasks are machines better at?
- Strong AI and Weak AI, which one do you support?

	Human	Machine
Thinking rationally		
Acting rationally		
Thinking humanly		
Acting humanly		