

Exploring Artificial Intelligence: Concepts, Applications, and Interactions

Ziyue Luo

Postdoc

The Ohio State University

Artificial Intelligence (AI)


- What is **Artificial Intelligence (AI)**?
 - AI is a branch of Computer Science, Engineering, and Mathematics that aims to create machines and software capable of performing tasks that usually require human intelligence
 - Not one thing. It's hundreds of different things
 - Good idea to talk about a particular function of AI
 - Ex. Deep Neural Networks, Reinforcement Learning, Active Learning, Meta Learning...
 - Technically Machine Learning (ML) is a sub-field of AI, but today encompasses most of AI activities
- Many different **applications** of AI/ML

Reinforcement learning
for Self Driving



vecteezy.com

Deep neural networks for
image recognition/classification




Large language models for
natural language processing



Brief History of Artificial Intelligence

- **Alan Turing:** One of the Founders of Modern CS
 - Universal (Turing) Machine (1937):
 - Turing Test or Imitation Game (1950)
 - To determine if a machine can exhibit intelligent behavior
- **Claude Shannon:** Creator of Information Theory (1940s)
 - Notion of the bit, digital circuits, Boolean algebra
 - Father of the Digital Age
 - No computers/communication systems would be possible without him
- **Dartmouth Conference (1956):**
 - Often called the “Birthplace of AI”
 - A program called **Logic Theorist** was presented
 - Mimic problem solving skills of humans




1956 Dartmouth Conference:
The Founding Fathers of AI




Brief History of Artificial Intelligence

- **1960s-2000s:** Several AI Summers and Winters
 - Great promise followed by disappointment
 - AI never appeared to meet the sky-high expectations
 - Steady progress nonetheless
 - Foundational Advances in ML, expert systems, etc.
 - **Deep Blue (IBM) beats Gary Kasprov...**
- **2010 onwards:** AI starts to go mainstream (**Why?**)
 - Computers become much faster and cheaper
 - Data rich environment → explosion of ML
 - Recommender systems, apps, search engines, all use AI
- **2016-present:** Major Milestones for DNN and RL
 - Alpha-Go beats world champion in the Board Game Go (2016)
 - Alpha-zero uses beats Alpha-Go without human input (2017)
 - AlphaFold “solves” the protein folding problem (2020)
 - Development and release of LLMs like ChatGPT




The Pillar of Modern AI: Deep Learning


- What is **Deep Learning**?
 - A *subfield of AI (ML)* that uses neural networks with multiple layers “deep” structures to extract high-level features from data
- Typical Components in Deep Learning
 - Large data sets, high-performance computing, and effective training algorithms
 - **Training Phase:** Train the NN on extensive amounts of data to find patterns or features so that it can make accurate predictions on the training data itself
 - **Testing Phase:** Test the NN on new data to see how well it performs on new datasets (generalizability of the NN).
 - If it hasn't done a good job, may need retraining or collect better training data



Neural network



Convolutional neural network



Transformer

Example: Image Classification for Cats

- We want the **computer** to **recognize a cat!**
- **Training phase:** We need several hundred training images with labels
 - Some images with cats
 - Other images without cats




Cat



Not cat


- We feed these images and labels into a neural network



Broken down into a
vector of pixels





Feed into NN



Example: Image Classification for Cats (cont'd)

- **Deeper Look** at the Training Phase

- NN has lots of **adjustable parts** called **weights**
- For each image, NN makes a guess of whether it is a cat or not, based on the current weights
- If the guess is wrong, then NN will update its weights to correct the guess in the next time
- After training over hundreds of images, NN will develop the ability to recognize specific features
- Training phase typically ends after reaching a **certain error threshold**




Feature: Pointy ear

- **Testing phase**


- Ask NN to identify a cat from a new set of images not used in the training phase
- If NN can identify most of them correctly, then our training is successful
 - Otherwise, may need to redo the training with new NN models

Game 1: Image Classification by Teachable Machine

- **What is Teachable Machine?**
 - Lets you teach a computer to recognize **images, sounds, or poses**
 - No coding needed!
 - Works in your web browser
- **Website:** teachablemachine.withgoogle.com
- **Demo:** We'll teach the computer to recognize hand gestures for rock, paper, and scissors



Demo: Gesture Classification



The interface is divided into three main sections: Training, Preview, and Output.

Training Section:

- Rock:** 56 Image Samples. Includes a Webcam icon, an Upload icon, and a row of 8 hand gesture images.
- Paper:** 53 Image Samples. Includes a Webcam icon, an Upload icon, and a row of 8 hand gesture images.
- Scissors:** 58 Image Samples. Includes a Webcam icon, an Upload icon, and a row of 8 hand gesture images.
- Add a class:** A dashed box with a plus icon and the text "Add a class".

Training Status: A central box labeled "Training" shows "Model Trained" and an "Advanced" dropdown menu.

Preview Section:


- Preview:** A button to view the model's performance.
- Export Model:** A button to export the trained model.
- Input:** A toggle switch set to "ON" and a "Webcam" dropdown menu.
- Switch Webcam:** A button to switch the input source.
- Webcam Feed:** A live video feed showing a hand making a Rock gesture.

Output Section:

- Rock:** 97% confidence.
- Paper:** 0% confidence.
- Scissors:** 0% confidence.


Demo: Prepare the Training Data

- Prepare the data using webcam or collected images




Demo: Train the Model

- Train the model using the prepared dataset



Demo: Inference

- Use the trained model to recognize hand gestures for rock, paper, and scissors



Hand-on Project

- **Dataset: Tiny-Imagenet** <https://tinyurl.com/yrz5zht7>
- 100,000 images of 200 classes (500 images for each class) downsized to 64×64 colored images. (Original Imagenet: at least 224 pixels width/height)
- Each image around 2 KB (dataset total size: ~200 MB)
- You will find 5 folders in the downloaded folder: car, orange, teddy, cat, and dog (with folder names specified). These are called the **training data**
- You will find a test folder, which contains the **testing data**
- Note: complete dataset is [on Huggingface](#)

Dataset: Tiny-Imagenet



Object Localization




Image Segmentation



(Video) Object Tracking



From YOLOv8

Where can all these be used?

- **Medical/Health care:** analyze medical images (X-rays, MRIs, CT scans) to assist in the early detection and diagnosis of diseases
- **Traffic/Autonomous Vehicle:** identify traffic signs, pedestrians, other vehicles...
- **Agriculture:** monitor crop health, identify diseases, pests...
- **Security:** identify potential threats or suspicious activities

Chatting with AI

- **What is ChatGPT?**
- Developed by OpenAI; GPT-4 released in 2023
- Considered as one of the most advanced AIs right now
- Trained on tons of text from the internet
- Can write stories, answer questions, explain stuff, and more



How do ChatGPT work?

- Generate text by predicting the next word in a sentence
 - Use the context of the conversation - the preceding words and sentences - to make these predictions by using a transformer-like learning architecture
 - Do not possess consciousness or beliefs
 - Instead, LLMs mimic the patterns they learned during training

What Can It Do?

- Answer questions
- Help write essays or poems
- Translate languages
- Write code
- Solve math problems (sometimes)
- Be a chatbot buddy

Turing Test Game

- Two answers for the same question: One by human, the other by ChatGPT
- Which one do you think is written by human?

Question: Describe your perfect day.

Answer 1: Sitting at home doing nothing while knowing I have nothing important to do so I don't feel bad for doing nothing.

Human

Answer 2: A perfect day would be spent enjoying nature, having meaningful conversations with loved ones, achieving a small goal, and ending with a good book or movie to unwind.

ChatGPT

Turing Test Game

- Two answers for the same question: One by human, the other by ChatGPT
- Which one do you think is written by human?

Question: What makes a good friend?

Answer 1: A good friend is someone who is loyal, supportive, honest, and always there to listen and uplift you.

ChatGPT

Answer 2: A good friend is willing to help you as much as you are willing to help your friend, so there is balance.

Human

Turing Test Game

- Two answers for the same question: One by human, the other by ChatGPT
- Which one do you think is written by human?

Question: You have six horses and want to race them to see which is fastest. How many rounds do you need? Just give me a number.

Answer 1: 3.

ChatGPT

Answer 2: 1.

Human


Timeline for ChatGPT

- Jun. 2018, GPT-1 released
 - Large language model with 117 million parameters, the groundwork for GPT series
- Feb. 2019, GPT-2 released
 - 1.5 billion parameters
- Jun. 2020, GPT-3 released
 - 175 billion parameters, very strong in many NLP applications
- Nov. 2022, ChatGPT (GPT-3.5) released, attracted great attention
- Mar. 2023, GPT-4 released
 - Unknown number of parameters, stronger than ChatGPT


Exam	GPT-4	GPT-4 (no vision)	GPT-3.5
Uniform Bar Exam (MBE+MEE+MPT)	298 / 400 (~90th)	298 / 400 (~90th)	213 / 400 (~10th)
LSAT	163 (~88th)	161 (~83rd)	149 (~40th)
SAT Evidence-Based Reading & Writing	710 / 800 (~93rd)	710 / 800 (~93rd)	670 / 800 (~87th)
SAT Math	700 / 800 (~89th)	690 / 800 (~89th)	590 / 800 (~70th)
Graduate Record Examination (GRE) Quantitative	163 / 170 (~80th)	157 / 170 (~62nd)	147 / 170 (~25th)
Graduate Record Examination (GRE) Verbal	169 / 170 (~99th)	165 / 170 (~96th)	154 / 170 (~63rd)
Graduate Record Examination (GRE) Writing	4 / 6 (~54th)	4 / 6 (~54th)	4 / 6 (~54th)
USABO Semifinal Exam 2020	87 / 150 (99th - 100th)	87 / 150 (99th - 100th)	43 / 150 (31st - 33rd)
USNCO Local Section Exam 2022	36 / 60	38 / 60	24 / 60
Medical Knowledge Self-Assessment Program	75 %	75 %	53 %
Codeforces Rating	392 (below 5th)	392 (below 5th)	260 (below 5th)
AP Art History	5 (86th - 100th)	5 (86th - 100th)	5 (86th - 100th)
AP Biology	5 (85th - 100th)	5 (85th - 100th)	4 (62nd - 85th)
AP Calculus BC	4 (43rd - 59th)	4 (43rd - 59th)	1 (0th - 7th)
AP Chemistry	4 (71st - 88th)	4 (71st - 88th)	2 (22nd - 46th)
AP English Language and Composition	2 (14th - 44th)	2 (14th - 44th)	2 (14th - 44th)
AP English Literature and Composition	2 (8th - 22nd)	2 (8th - 22nd)	2 (8th - 22nd)
AP Environmental Science	5 (91st - 100th)	5 (91st - 100th)	5 (91st - 100th)
AP Macroeconomics	5 (84th - 100th)	5 (84th - 100th)	2 (33rd - 48th)
AP Microeconomics	5 (82nd - 100th)	4 (60th - 82nd)	4 (60th - 82nd)
AP Physics 2	4 (66th - 84th)	4 (66th - 84th)	3 (30th - 66th)
AP Psychology	5 (83rd - 100th)	5 (83rd - 100th)	5 (83rd - 100th)
AP Statistics	5 (85th - 100th)	5 (85th - 100th)	3 (40th - 63rd)
AP US Government	5 (88th - 100th)	5 (88th - 100th)	4 (77th - 88th)
AP US History	5 (89th - 100th)	4 (74th - 89th)	4 (74th - 89th)
AP World History	4 (65th - 87th)	4 (65th - 87th)	4 (65th - 87th)
AMC 10 ³	30 / 150 (6th - 12th)	36 / 150 (10th - 19th)	36 / 150 (10th - 19th)
AMC 12 ³	60 / 150 (45th - 66th)	48 / 150 (19th - 40th)	30 / 150 (4th - 8th)
Introductory Sommelier (theory knowledge)	92 %	92 %	80 %
Certified Sommelier (theory knowledge)	86 %	86 %	58 %
Advanced Sommelier (theory knowledge)	77 %	77 %	46 %

GPT-4 Technical Report, OpenAI, Mar. 2023

An Internet-Scale LLM Training System



An Internet-Scale LLM Training System



Demo