

CS50 for MBAs

Artificial Intelligence

Artificial Intelligence

| | | |
|---|---|---|
| | O | X |
| O | X | X |
| O | | X |

**Game
Playing**




Deep Blue versus Garry Kasparov, 1997

Game Playing






AlphaGo versus Lee Sedol, 2016

Game Playing



Classification





Classification


Hello, world!



Hello, world!

**Handwriting
Recognition**

Watch History




Recommendations



**Recommendation
Systems**



Question Answering




ChatGPT: Optimizing Language Models for Dialogue

We've trained a model called ChatGPT which interacts in a conversational way. The dialogue format makes it possible for ChatGPT to answer followup questions, admit its mistakes, challenge incorrect premises, and reject inappropriate requests. ChatGPT is a sibling model to InstructGPT, which is trained to follow an instruction in a prompt and provide a detailed response.

[TRY CHATGPT ↗](#)

Text Generation



Larry Buchanan/The New York Times

TheUpshot

Did a Fourth Grader Write This? Or the New Chatbot?

Don't be surprised if you can't always tell. Neither
could a fourth-grade teacher — or Judy Blume.

Text Generation

Essay 1

... I like to bring a yummy sandwich and a cold juice box for lunch, and sometimes I'll even pack a tasty piece of fruit or a bag of crunchy chips. As we eat, we chat and laugh and catch up on each other's day. ...

Essay 2

... My mother packs me a sandwich, a drink, fruit, and a treat. When I get in the lunchroom I find an empty table and sit there, and eat my lunch. My friends come and sit down with me. ...

Text Generation

Image 1




Image 2



**Image
Generation**

Artificial Intelligence



Decision Trees

Is ball left of paddle?

Yes

No

Move paddle left.


Is ball right of paddle?

Yes

No

Move paddle right.

Don't move paddle.



```
while game is ongoing:  
    if ball left of paddle:  
        move paddle left  
    else if ball right of paddle:  
        move paddle right  
    else:  
        don't move paddle
```

| | | |
|---|---|---|
| O | | X |
| | X | |
| O | | |

Can I get 3 in a row
this turn?

Yes

No

**Play in square to
get 3 in a row.**

Can my opponent get 3 in a
row on their next turn?

Yes

No

**Play in square to block
opponent's 3 in a row.**

?

Minimax

MAX (X) aims to maximize score.

MIN (O) aims to minimize score.

| | | |
|---|---|---|
| O | X | X |
| O | O | |
| O | X | X |

-1

| | | |
|---|---|---|
| X | O | X |
| O | O | X |
| X | X | O |

0

| | | |
|---|---|---|
| O | | X |
| | X | O |
| X | O | X |

1

| | | |
|---|---|---|
| O | X | O |
| O | X | X |
| X | X | O |

VALUE: 1

TURN = O

VALUE:
 \emptyset

| | | |
|---|---|---|
| | X | O |
| O | X | X |
| X | | O |

VALUE:
1

| | | |
|---|---|---|
| O | X | O |
| O | X | X |
| X | | O |

VALUE:
 \emptyset

| | | |
|---|---|---|
| | X | O |
| O | X | X |
| X | O | O |

VALUE:
1

| | | |
|---|---|---|
| O | X | O |
| O | X | X |
| X | X | O |

VALUE:
 \emptyset

| | | |
|---|---|---|
| X | X | O |
| O | X | X |
| X | O | O |

TURN = 0

VALUE:
 \emptyset

| | | |
|---|---|---|
| | X | O |
| O | X | X |
| X | | O |

VALUE:
1

| | | |
|---|---|---|
| O | X | O |
| O | X | X |
| X | | O |

VALUE:
 \emptyset

| | | |
|---|---|---|
| | X | O |
| O | X | X |
| X | O | O |

VALUE:
1

| | | |
|---|---|---|
| O | X | O |
| O | X | X |
| X | X | O |

VALUE:
 \emptyset

| | | |
|---|---|---|
| X | X | O |
| O | X | X |
| X | O | O |

TURN = X

VALUE:
1

| | | |
|---|---|---|
| | X | O |
| O | X | |
| X | | O |

VALUE:
0

| | | |
|---|---|---|
| | X | O |
| O | X | X |
| X | | O |

VALUE:
-1

| | | |
|---|---|---|
| X | X | O |
| O | X | |
| X | | O |

VALUE:
1

| | | |
|---|---|---|
| X | X | O |
| O | X | |
| X | X | O |

VALUE:
1

| | | |
|---|---|---|
| O | X | O |
| O | X | X |
| X | | O |

VALUE:
0

| | | |
|---|---|---|
| | X | O |
| O | X | X |
| X | O | O |

VALUE:
-1

| | | |
|---|---|---|
| X | X | O |
| O | X | O |
| X | | O |

VALUE:
0

| | | |
|---|---|---|
| X | X | O |
| O | X | |
| X | O | O |

VALUE:
1

| | | |
|---|---|---|
| O | X | O |
| O | X | X |
| X | X | O |

VALUE:
0

| | | |
|---|---|---|
| X | X | O |
| O | X | X |
| X | O | O |

VALUE:
0

| | | |
|---|---|---|
| X | X | O |
| O | X | X |
| X | O | O |

Minimax

```
if player is X:  
    for all possible moves:  
        calculate score for board  
    choose move with highest score  
  
else:  
    for all possible moves:  
        calculate score for board  
    choose move with lowest score
```

Possibilities

Tic-Tac-Toe

All possible games

255,168

Chess

First 4 moves only

288 million

Go

First 4 moves only

266 quintillion

Depth-Limited Minimax

evaluation function

function that estimates the expected utility of the game from a given state



Chess


Deep Blue versus Garry Kasparov, 1997


Why AI is Good at Games


Inputs and outputs that can be represented mathematically.


A **goal** in the form of a **function to maximize/minimize**.


Reinforcement Learning




















Explore vs. Exploit

Explore vs. Exploit Strategy

```
epsilon = 0.10
```

```
if random() < epsilon:  
    make a random move
```

```
else:
```

```
    make the move with the highest value
```


**Robot Motor Skill
Coordination with EM-based
Reinforcement Learning**

**Petar Kormushev, Sylvain Calinon,
and Darwin G. Caldwell**

Italian Institute of Technology

Source: YouTube, Italian Institute of Technology

Watch History



Recommendations




**Recommendation
Systems**




AlphaGo versus Lee Sedol, 2016

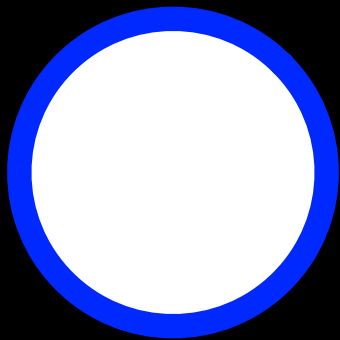
Go









Neural Networks







































A diagram illustrating a process flow. It consists of two square boxes, one on the left and one on the right, both with white outlines. The left box contains the word "input" and the right box contains the word "output". A white arrow points from the right side of the "input" box to the left side of the "output" box, indicating a directional flow from left to right.


input

output

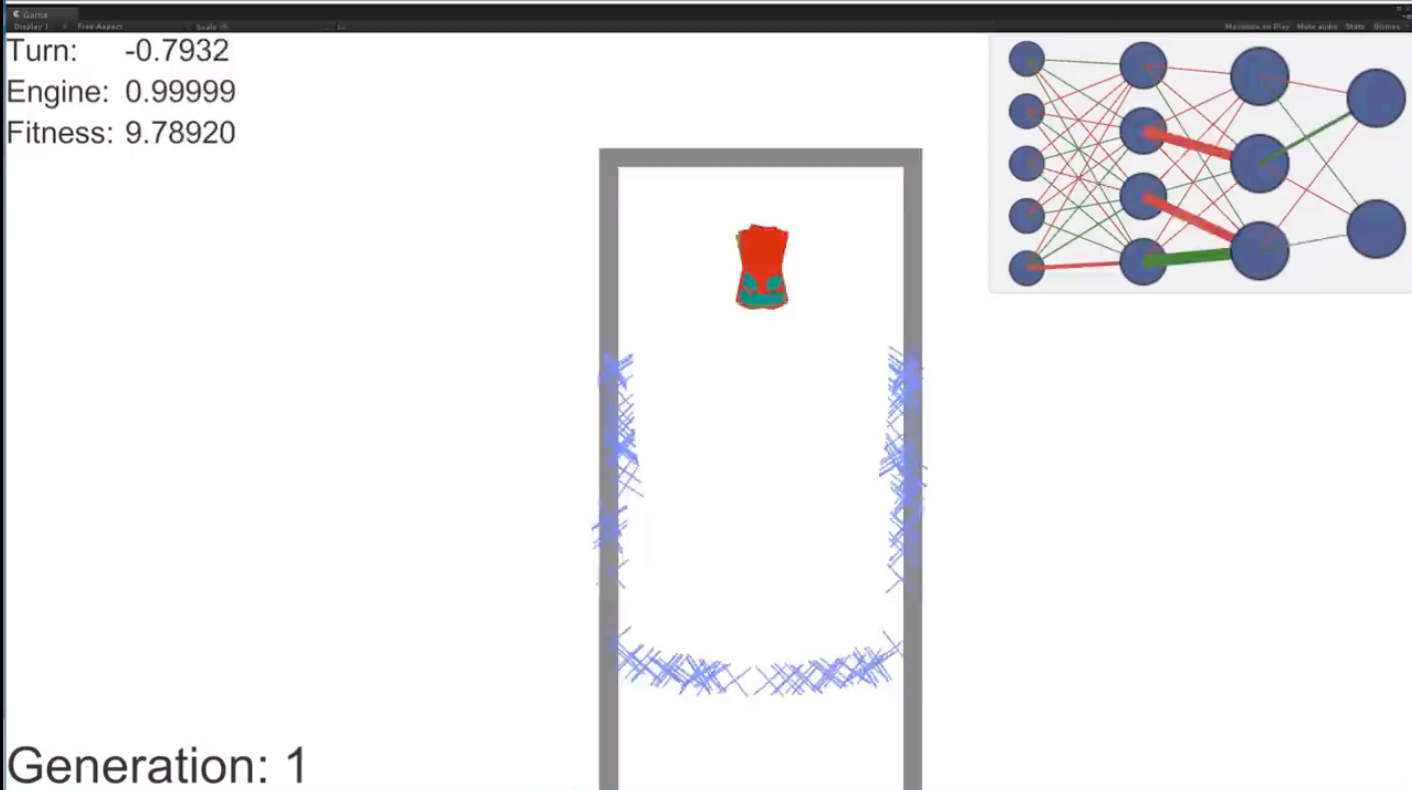








Game
Turn: -0.7932
Engine: 0.99999
Fitness: 9.78920




Generation: 1

The image shows a game window with a white background. In the center, a grey rectangular frame contains a small red and green creature. The frame is surrounded by blue scribbles. In the top right corner, there is a neural network diagram with blue nodes and red and green connections. The top left corner displays game statistics: Turn: -0.7932, Engine: 0.99999, and Fitness: 9.78920. The bottom left corner displays 'Generation: 1'.


Source: *YouTube, Samuel Arzt*



Dataset of Real Images




Generator





Discriminator

Dataset of Real Images



Generator




Discriminator


Real



Dataset of Real Images



Generator





Discriminator


Generated



Dataset of Real Images



Generator




Real





Discriminator

Dataset of Real Images



Generator




Real




Discriminator




Dataset of Real Images



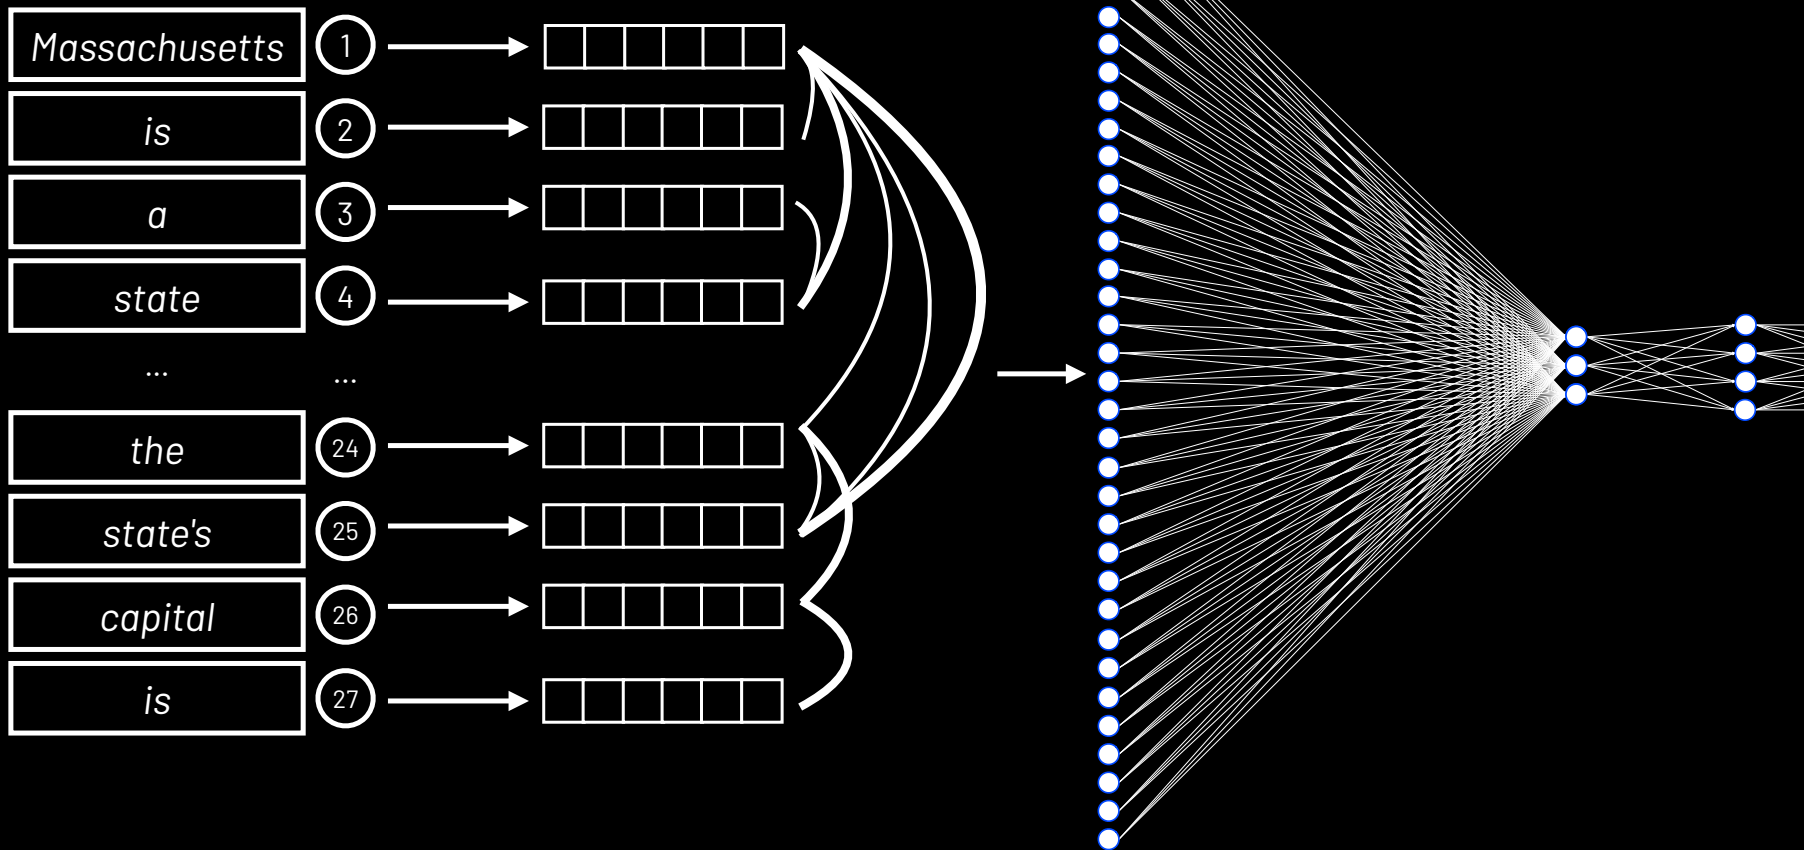
Generator



Discriminator



"Massachusetts is a state in the New England region of the Northeastern United States. It borders on the Atlantic Ocean to the east. The state's capital is..."



Attention Is All You Need

Ashish Vaswani*
Google Brain
avaswani@google.com

Noam Shazeer*
Google Brain
noam@google.com

Niki Parmar*
Google Research
nikip@google.com

Jakob Uszkoreit*
Google Research
usz@google.com

Llion Jones*
Google Research
llion@google.com

Aidan N. Gomez* †
University of Toronto
aidan@cs.toronto.edu

Lukasz Kaiser*
Google Brain
lukasz.kaiser@google.com

Illia Polosukhin* †
illia.polosukhin@gmail.com

Abstract

Transformer Model
2017

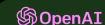
Language Models are Few-Shot Learners

Tom B. Brown* Benjamin Mann* Nick Ryder* Melanie Subbiah*
Jared Kaplan† Prafulla Dhariwal Arvind Neelakantan Pranav Shyam Girish Sastry
Amanda Askell Sandhini Agarwal Ariel Herbert-Voss Gretchen Krueger Tom Henighan
Rewon Child Aditya Ramesh Daniel M. Ziegler Jeffrey Wu Clemens Winter
Christopher Hesse Mark Chen Eric Sigler Mateusz Litwin Scott Gray
Benjamin Chess Jack Clark Christopher Berner
Sam McCandlish Alec Radford Ilya Sutskever Dario Amodei

OpenAI

Abstract

GPT-3 Language Model
2020



ChatGPT: Optimizing Language Models for Dialogue

We've trained a model called ChatGPT which interacts in a conversational way. The dialogue format makes it possible for ChatGPT to answer followup questions, admit its mistakes, challenge incorrect premises, and reject inappropriate requests. ChatGPT is a sibling model to InstructGPT, which is trained to follow an instruction in a prompt and provide a detailed response.

[TRY CHATGPT ↗](#)

ChatGPT Model
2022

Artificial Intelligence

Assignment 4

Office Hours

Lab 2

CS50 for MBAs

Artificial Intelligence