Welcome to "Little Bits to Big Ideas"

Lab 7: Artificial Intelligence and ChatBots

TAs: Ashish & Shivani

CAs: Abhinav, Harry, & Sherry

ILLINOIS

What is Artificial Intelligence (AI)?

“Artificial intelligence” predates “computer science” by 5 years.

The concept predates computing by hundreds of years.

The term is an aspiration:

Computer systems that do things that humans think of as intelligent.

What is part of AI changes as perceptions of intelligence change.

Two broad types of AI

Classical AI: Humans model an aspect of intelligent behavior and program an algorithm that exhibits that behavior.

Machine learning (ML): Humans program a way for a computer to pick one of many possible algorithms by seeing which one comes closest to matching a bunch of examples, called **training data**.

Characteristics of Machine Learning

Supervised: The training data has desired outputs. We know what the system *should* do for each example input we provide.

Unsupervised: The training data is just inputs. We know characteristics of the output we want, but not exact outputs.

Categorical: The output is one of a discrete set of options.

Example: a **word** in a piece of text.

Continuous: The output is one or more real numbers.

Example: **color** of a pixel in an image (3 real numbers, RGB).

Common types of Machine Learning tasks

Classification: pick one of a known set of options for each input.

→ Cat

Clustering: identify similar groups and pick one group for each input.

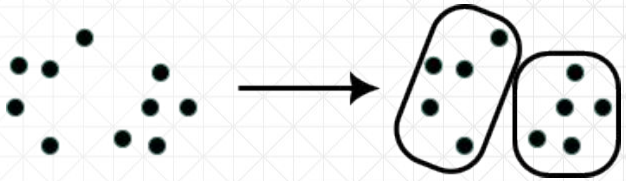

Generative: find a value that fits a gap in a provided pattern.

This is an|example

Regression: predict a number not in the input based on what is in the input.

2.4 hr/week study

⇒ 3.64

GPA

Overfitting

Parameters: Ways different functions may vary.

Training Data: Examples used to pick a good function.

Overfitting: Function selected that matches training data details, not the pattern the training data is supposed to exemplify.

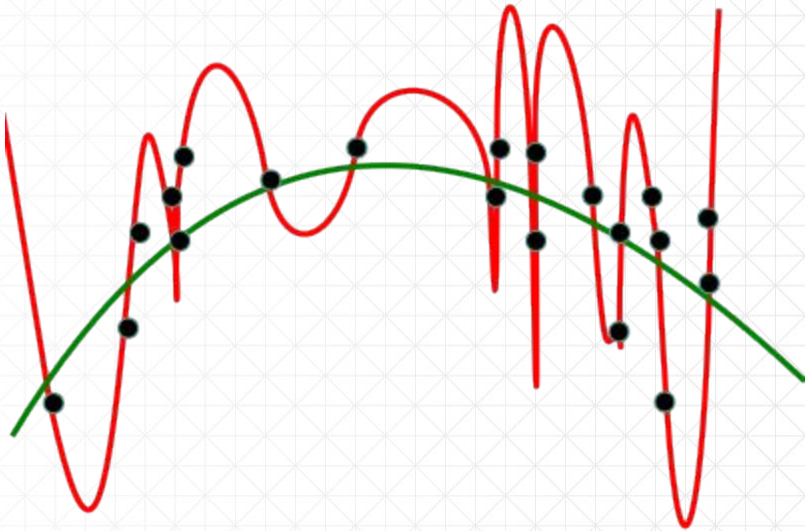

ELIZA, the first chatbot (1966)

ELIZA—A Computer Program For the Study of Natural Language Communication Between Man And Machine

JOSEPH WEIZENBAUM
Massachusetts Institute of Technology, Cambridge, Mass.*

ELIZA is a program operating within the MAC time-sharing

to that reserved for curios, fit to be discussed only with people less enlightened than he.

Work reported herein was supported (in part) by Project MAC, an MIT research program sponsored by the Advanced Research Projects Agency, Department of Defense, under Office of Naval Research Contract Number Nonr-4102(01).

* Department of Electrical Engineering.

The object of this paper is to cause just such a re-evaluation of the program about to be "explained". Few programs ever needed it more.

ELIZA Program

ELIZA is a program which makes natural language conversation with a computer possible. Its present implementation is on the MAC time-sharing system at MIT. It is written in MAD-SLIP [4] for the IBM 7094. Its name was chosen to emphasize that it may be incrementally improved by its users, since its language abilities may be continually improved by a "teacher". Like the Eliza of Pygmalion fame, it can be made to appear even more

Men are all alike.
IN WHAT WAY
They're always bugging us about something or other.
CAN YOU THINK OF A SPECIFIC EXAMPLE
Well, my boyfriend made me come here.
YOUR BOYFRIEND MADE YOU COME HERE
He says I'm depressed much of the time.
I AM SORRY TO HEAR YOU ARE DEPRESSED

ELIZA

- Classical AI (not ML)
- Pretends to be a therapist or counselor
- Only recognizes a few words in a few contexts
  - Has multiple rules for each context, so it doesn’t immediately repeat itself
  - Some rules store part of what you type, other rules repeat what was stored