Supervised Learning

Session 3

Slides Link:
https://teachla.uclaacm.com/classes/ml

ACM AI + ACM TeachLA

What are you going to be for Halloween?

Recap

Key Terms

- **Artificial Intelligence** - The theory and development of computer systems able to perform tasks that normally require human intelligence.
- **Machine Learning** - A type of AI that provides computers with the ability to learn without being explicitly programmed.
- **Deep Learning** - An ML learning method based on artificial neural networks.

AI

- if-else statements
- Decision Trees + Bayesian decision-making
- data mining
- natural language processing

ML

- supervised learning
- unsupervised learning
- reinforcement learning
- deep learning
  - neural networks

Supervised Learning?

It's one of the most widely used forms of machine learning!

guided

directed

Supervised

helped

Let’s imagine a basketball game...

...with 3 seconds left, down by 1!

Our captain (**Mr. Supervisor**) has the ball (**input**)!

but is trapped!

Mr. Supervisor passes to **Mr. Model** who can score!

However, **Mr. Model** actually doesn't know how to play basketball 😳

. . .he was given the ball (**input data**) but. . .

. . .doesn't know what to do (**output**)

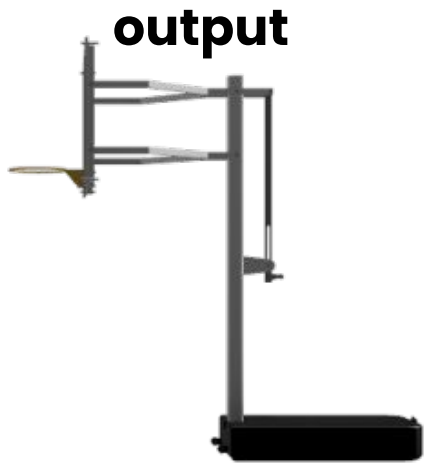

Mr. Model needs a supervisor to point him in the right direction. . .

Mr. Supervisor is here to tell him the **output** (what he should do)

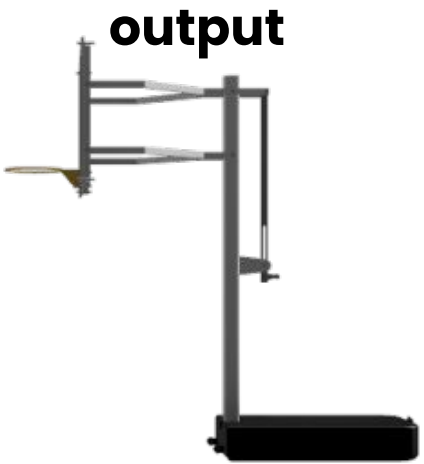

Mr. Supervisor: I DON'T CARE HOW you do it, just make sure the ball goes in the basket!

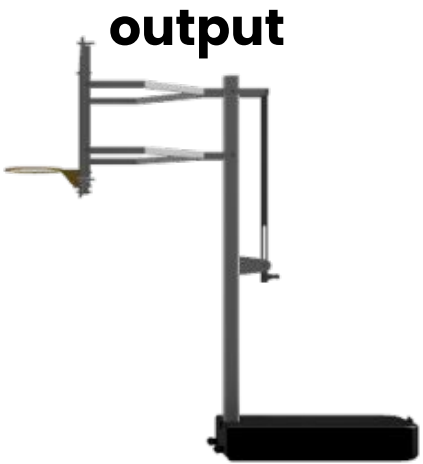

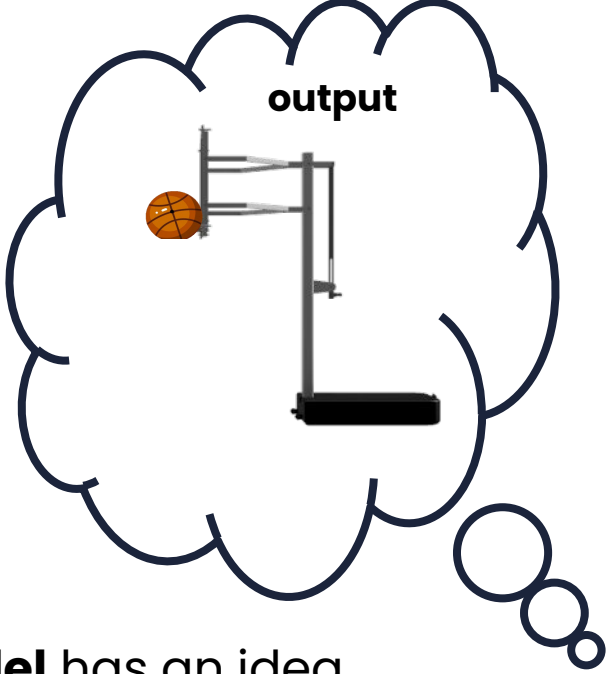

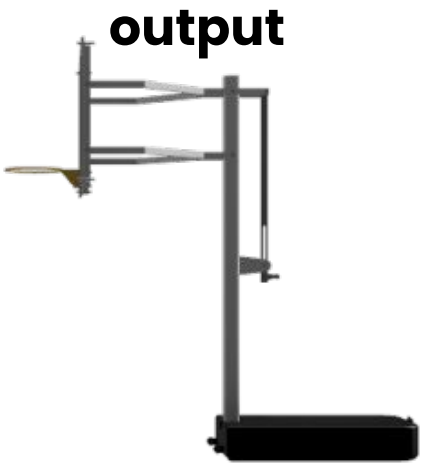

Now that **Mr. Model** has an idea of the <u>end result</u> (**output**). . .

Mr. Model can use the ball (**input**) to score!

In other words. . .

The **model** uses input *and* labeled output data to map the **input** to the **output**

labeled output

Now that **Mr. Model** has an idea of the end result (**output**). . .

Mr. Model can use the ball (**input**) to score!

Supervised Learning:

The **model** (function) uses input *and* labeled output data to map the **input** to the **output**

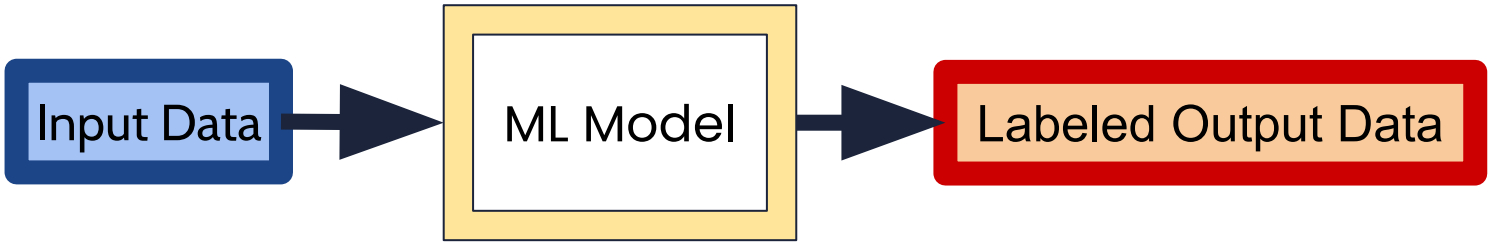

$$f(x)$$

it's like a big math function!

*Except instead of figuring out the output, we're figuring out the *relationship* of input->output

For reference...

$$f(x) = y$$

Normal/AI outside ML

- **Given:** input and defined relationship / function
- **Find:** the output

$$f(x) = y$$

Machine Learning

- **Given:** input and output
- **Find:** the relationship / function

Supervised Learning:

The **model** (function) uses input *and* labeled output data to map the **input** to the **output**

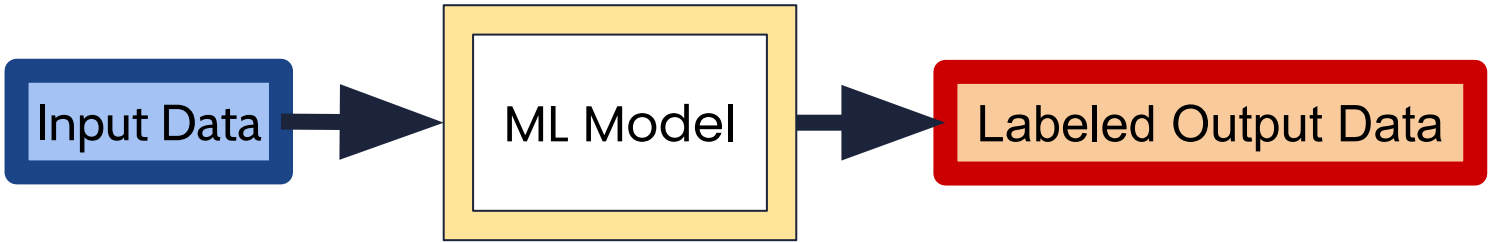

It's "supervised" because we *tell* the model the labeled output *beforehand*

take a moment to let that sink in!

Recap So Far:

- A Machine Learning *model* is a big function
- A Supervised Machine Learning model uses labeled output data, hence the name “supervised”
  - Labeled: output data we help *give* the model
- Supervised Learning **maps the relationship of:**

input → labeled output


LABELED DATA
Hello
Data
UNLABELED DATA

Training Models

Not these types of models!

Mapping? Learning?

- Supervised machine learning is all about **“mapping”** the relationship of input to output
- But how exactly does it **“map?”** (given input and output)

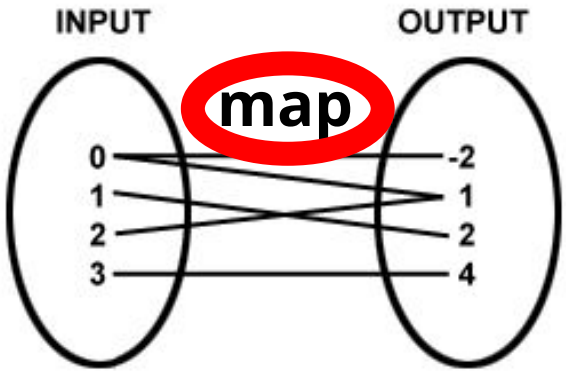

Mapping? Learning?

- Supervised machine learning is all about **“mapping”** the relationship of input to output
- But how exactly does it **“map?”** (given input and output)

Let's take a look at our basketball example again...

Recall that **Mr. Model** doesn't know how to play basketball

He doesn't know how to:

- shoot
- dribble
- dunk!

He has to **<u>train</u>** to learn those skills!

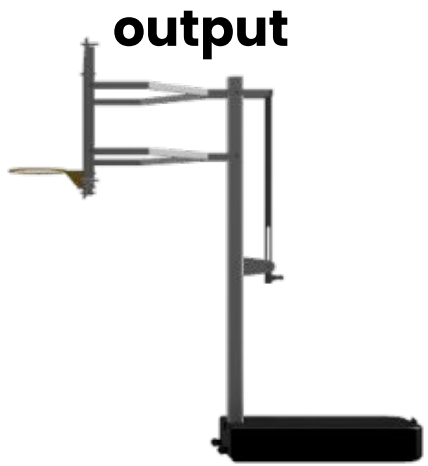

Training

So in order to **map** input to output consistently...

Training

So in order to **map** input to output consistently...

Mr. Model must **train** so he gets good at scoring baskets

(input → output)

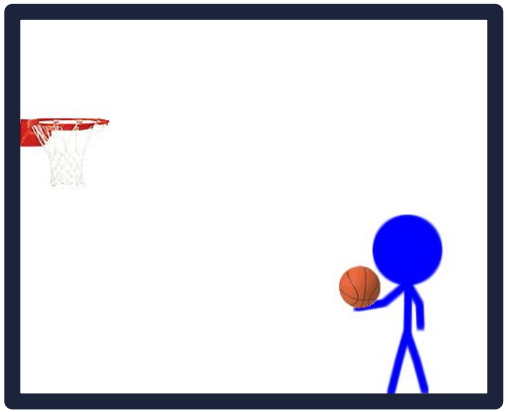

A Model in Training

- A model “training” means it’s ***learning*** ways to correctly *map* input data to output data
- Hence “machine *learning*”
- “Training” and “Learning” are interchangeable terms

A “Trained” Model

So after hours of training… **Mr. Model** is now trained and able to score hoops!

Training a machine learning model is a long, hefty mathematical process

We’ll learn how it specifically works in the coming weeks!

Me: *uses machine learning*

Machine: *learns*

Me:

How to Effectively Train Models

Effective Training

During preseason, in preparation for the NBA season **Mr. Model** trains *a lot*

He trains by playing against the B team, his **training data**

A **training dataset** is any dataset that allows our model to *train* (map input -> output)

Let's take a closer look at our B Team data!

B Team

In our case, the B Team is our training dataset: **Mr. Model** trains on them in order to learn how to play bball

Training on the B Team

4 feet

The average height of the B Team is 4 feet. . .

So **Mr. Model** just dunks on them <u>everytime</u>

It's so effective, it becomes the <u>only</u> thing **Mr. Model *<u>learns</u>* how to do

Testing our training abilities

Now that preseason is over, it's time to *test* **Mr. Model's** new bball skills...

...on the NBA!

Trained Model

Test Dataset

When we want to test to see if our *model* is well trained, we *test* our model on a **test dataset**

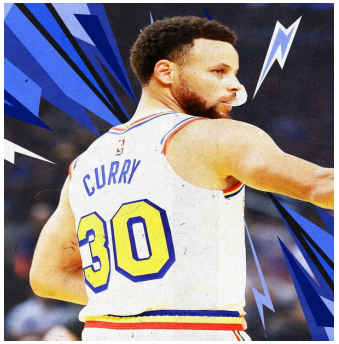

Test Data (NBA)

The **test dataset** is an example of what the model will see in the real world... it's a true test of its abilities!

Effective Training??

As we test **Mr. Model** plays against the tall, skilled players of the NBA we almost forgot...

...that he only knows how to dunk!

He's only able to dunk on people that are 4 feet tall! (B Team)

(In)effective Training

Did **Mr. Model** have good training? Especially if he plans to play against the tall, skilled players of the NBA?

(In)effective Training (cont'd)

Did **Mr. Model** have good training? Especially if he plans to play against the tall, skilled players of the NBA?

NO

(In)effective Training (cont'd)

Mr. Model has become so accustomed to dunking on the short players of the B Team that. . .

. . .he sucks at playing against anyone taller than 6 feet because dunking won't work on them!

Mr. Model had bad training!

Overfitting

Whenever a model has become too accustomed to a specific dataset (playing against B team too much), it's called ***overfitting***.

- The model is too mode-specific
- Not "general" enough to handle real-world randomness
  - Randomness of players: tall AND short, fast AND slow, good AND bad, clueless AND intelligent

Overfitting (cont.)

- **Overfitting** occurs when model becomes too sample-specific (only knows how the B Team plays)
- Generally, a model *should* be trained to perform well on the entire population of data (well-rounded players)
  - Better for Tesla cars to drive all streets decently well rather than one specific street perfectly

Combatting Overfitting

What are some ways **Mr. Model** could've avoided overfitting?

- Playing against diverse set of teams/players
  - Use a diverse input/output dataset for your model
- Train more! Train skills *other* than dunking too
  - Using *more* diverse data for your model <u>always</u> helps

Combatting Overfitting (cont.)

What does this tell us about the **training dataset**?

- Our training dataset was NOT diverse enough!
- Ideally, we want the training dataset to fairly represent the *majority* of the population

Avoid training against a small subset of your population, or else you'll get creamed like **Mr. Model in the NBA!

Underfitting

Conversely, if **Mr. Model** is lazy (doesn't have much data to train on) he can be susceptible to ***underfitting***

Underfitting: when the model is too *basic*, and doesn't work effectively in most situations

<u>How do we combat this?</u>

Gather more data to train on!

Training (graphically)

- Predictor variable: any input variable used to measure output

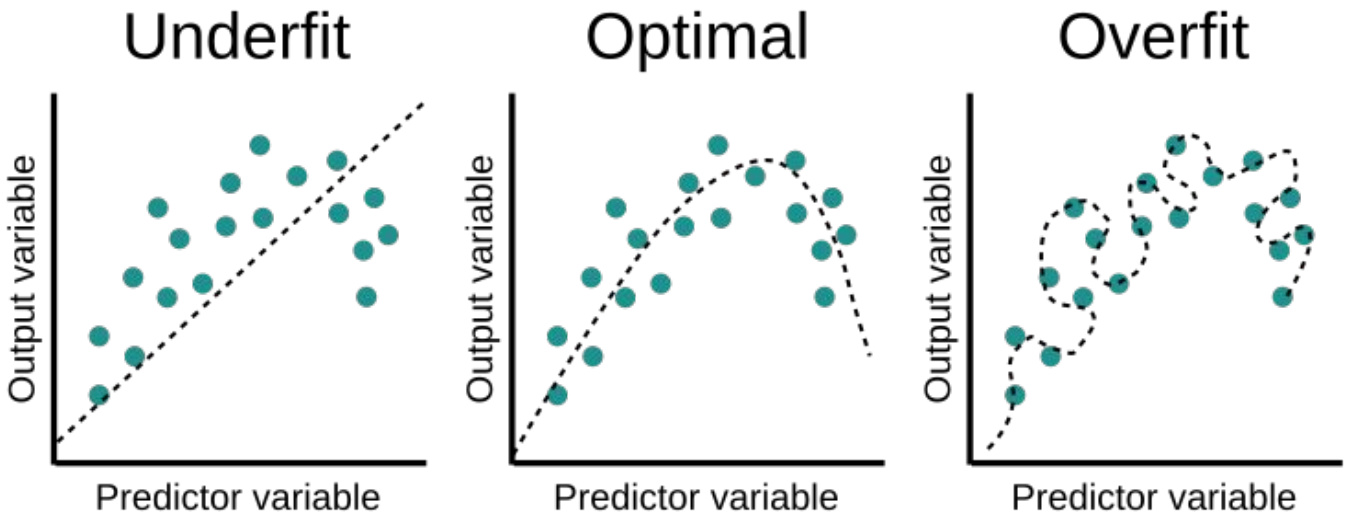

Training (graphically)

- Predictor variable: any input variable used to measure output

The black dotted line represents our model.

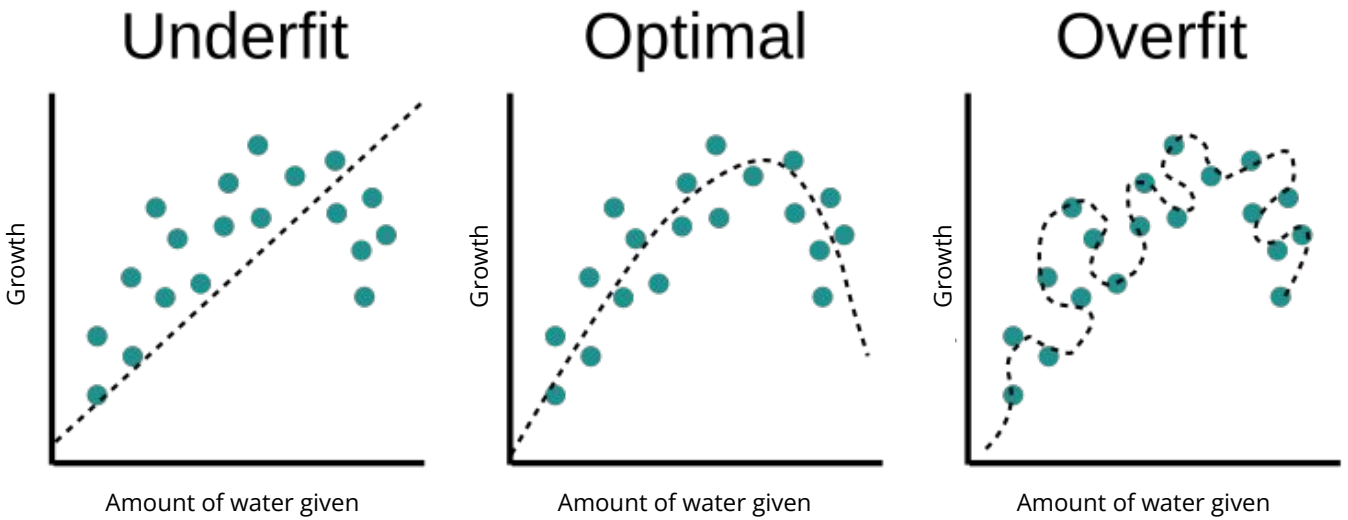

¡Kahoot!

Homework (optional)

Closing Comments

Key Takeaways

1. Supervised ML Models are just big functions!
2. ML Models “*train*”: finding the relation between input → output
3. Train on A LOT of diverse data to avoid overfitting/underfitting

ACM AI

Thanks!

Fill out our Feedback Form:
https://tinyurl.com/aimlworld

\- - -

Check out our Github
Read our Blog