

Demystifying AI


Image Classification

Instructor: Pat Virtue

Image Classification

Input: Image


Output: Category that object in image belongs to (label)



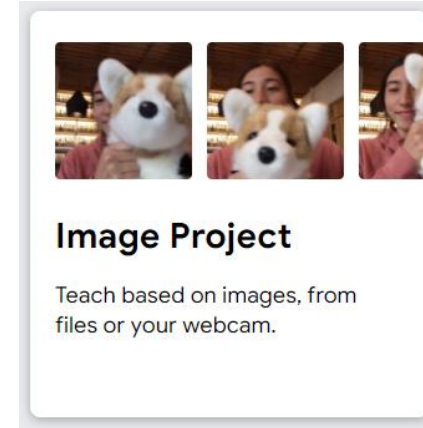
Teachable Machine

1. <https://teachablemachine.withgoogle.com/>

2. Click Get Started

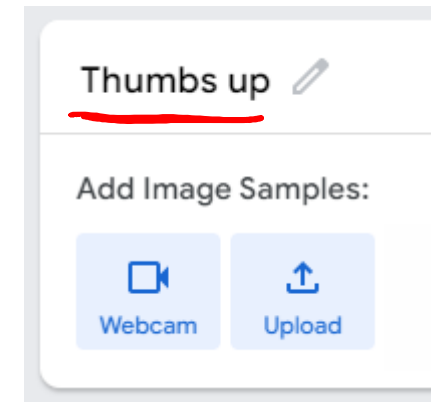


3. Select Image Project




4. Select Standard Image Model


5. Rename Class 1 and 2: Thumbs up and thumbs down



6. Use webcam to collect images



7. Train your model



Reminder: Machine Learning


Using (training) data to learn a model that we'll later use for prediction

Training Data
Input and
Measured Output



Model
Structure and
Parameters

Input



Predicted
Output

ML Image Classification Training



Training Data:

Images and corresponding labels



Model:

Neural network

- Humans design the neural network *structure*
- Network *parameters* (think slope and intercept, but a whole lot of them) are determined during training to minimize training error

↓
 $y = mx + b$


Very Loose Inspiration: Human Neurons



Neural Networks

Inspired by actual human brain

Input
Signal



Output
Signal

DOG



CAT



TREE



CAR

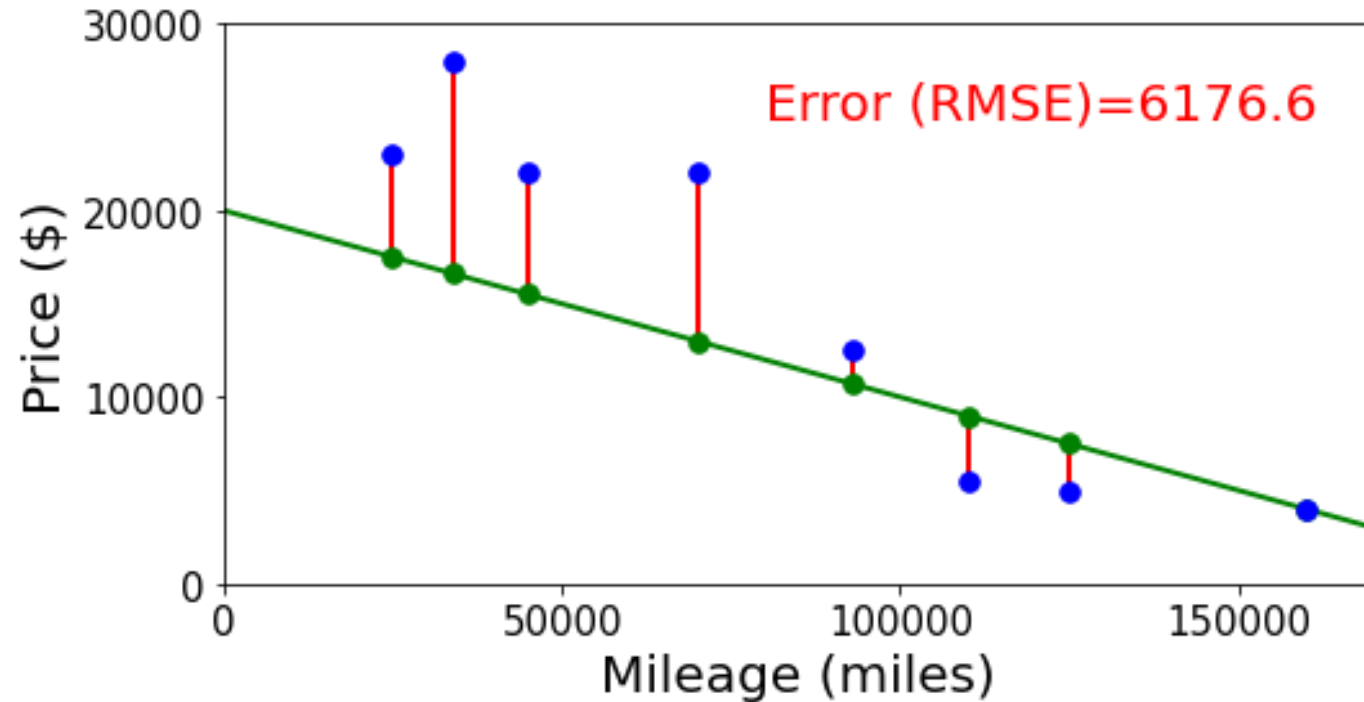



SKY

Neural Networks

Building on optimization for linear and logistic regression

- Selling my car



y_{meas}

y_{pred}

Neural Networks


Many layers of neurons, millions of parameters

Output

Signal

Input
Signal

y_{pred}



DOG



CAT




TREE



CAR



SKY





Neural Networks

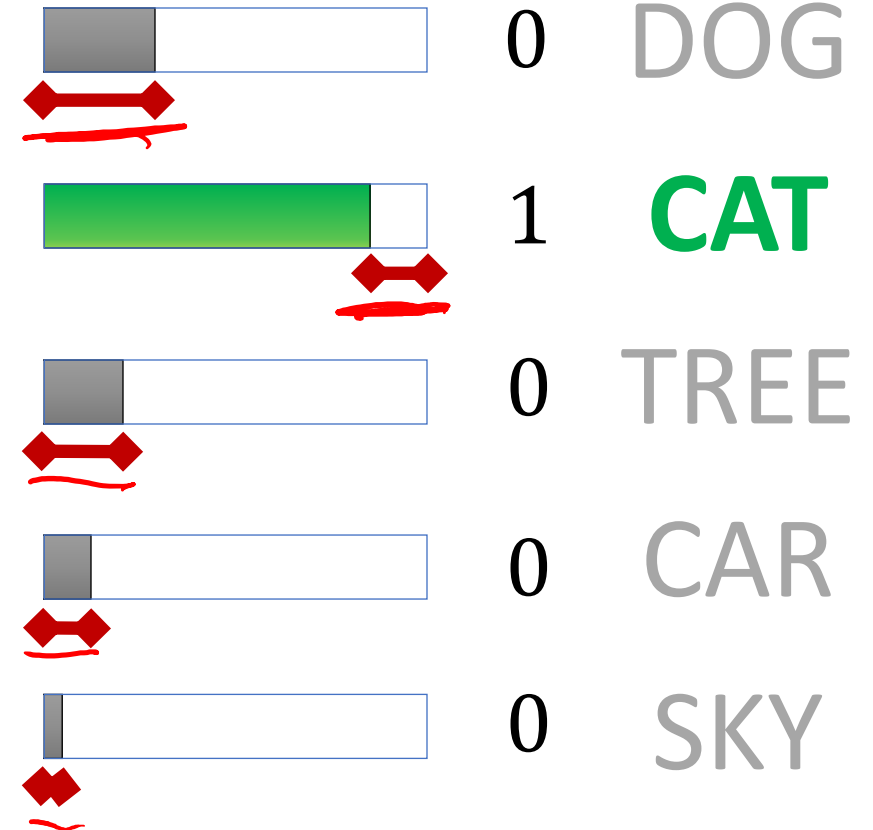
Many layers of neurons, millions of parameters

Output
Signal

Input
Signal



y_{pred} y_{meas}






Neural Networks

Many layers of neurons, millions of parameters

Output
Signal

Input
Signal




y_{pred}




LEFT




RIGHT



UP



DOWN



BUTTON

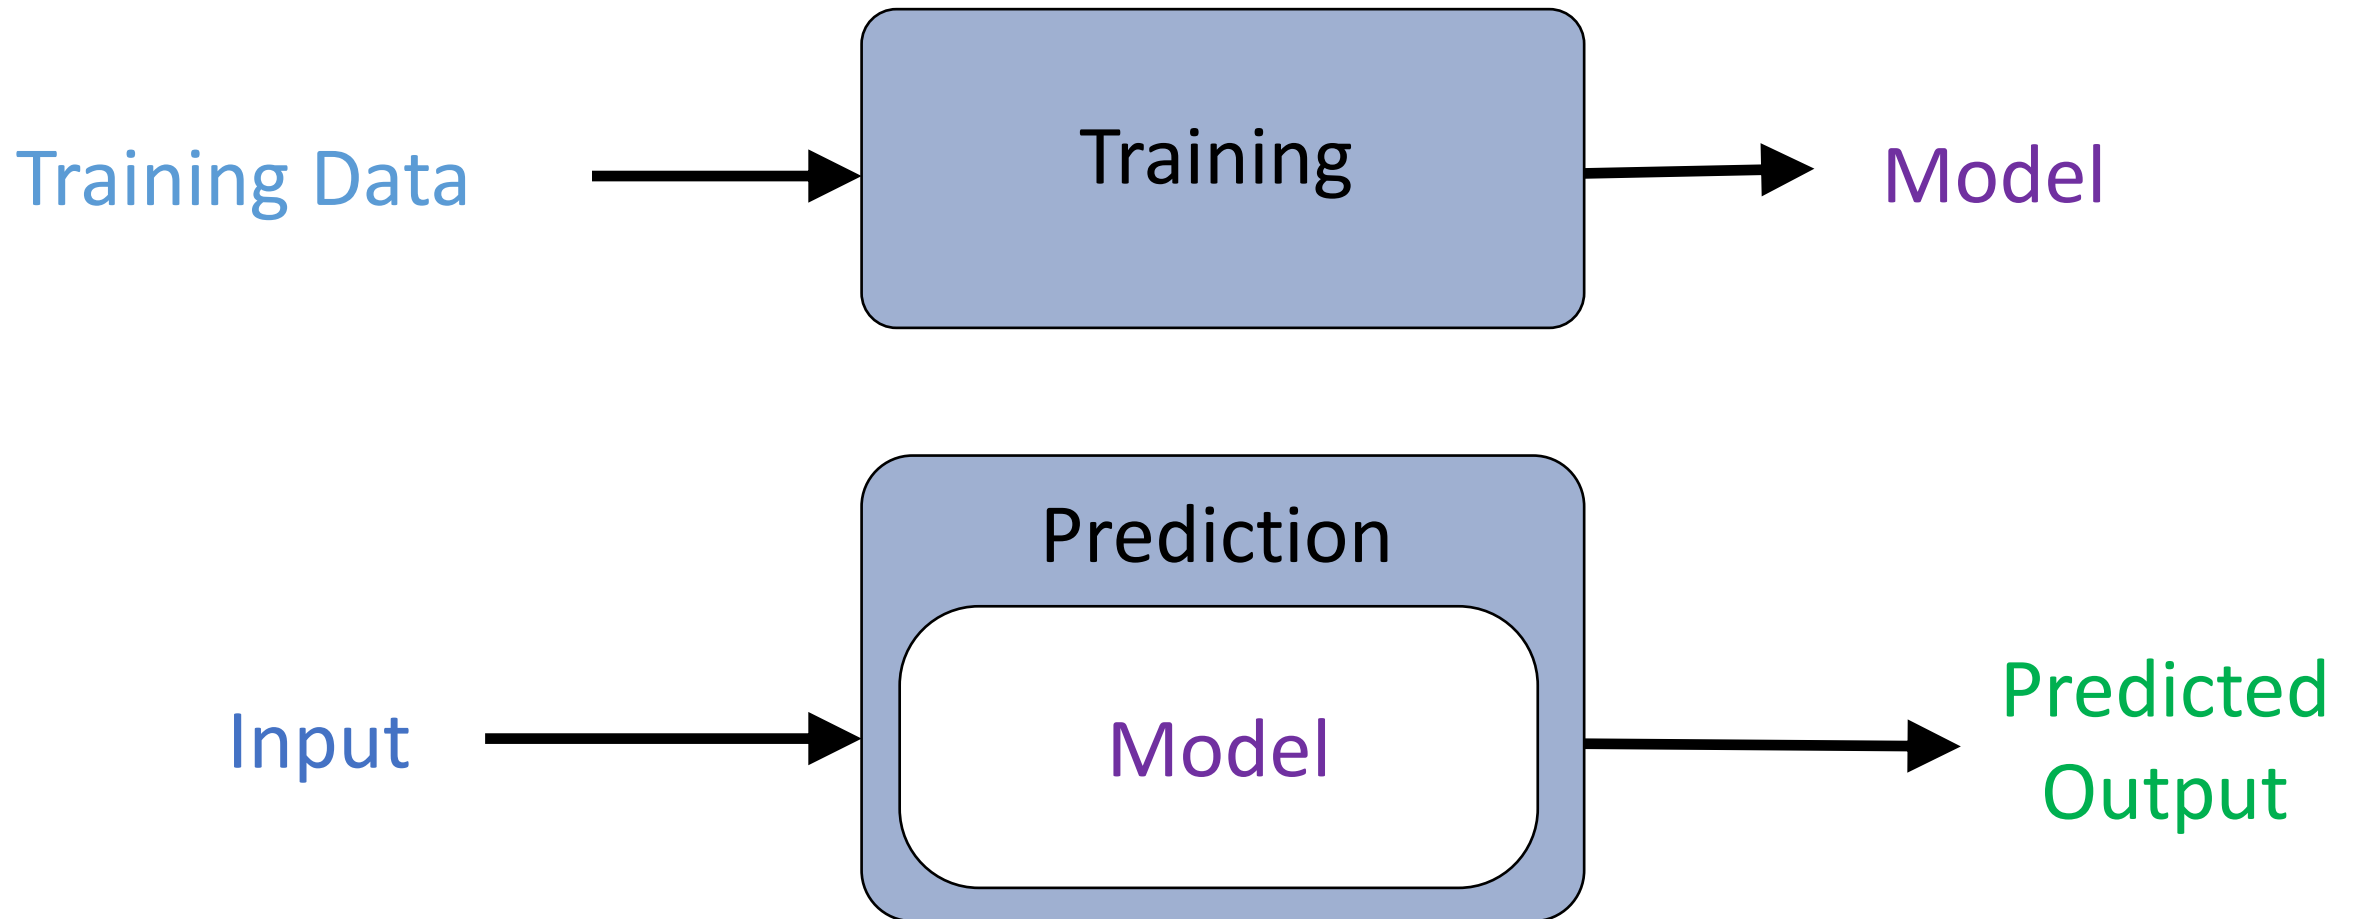
Teachable Machine

<https://teachablemachine.withgoogle.com/>

1. Task 1: Create a new image classification project with at least three classes
2. Task 2: Try training a project with very few images (one image per class maybe). Is this enough? Why or why not?
3. Task 3: Try training a project with training data that looks decent but is actually really biased.

Reminder: Machine Learning

Using (training) data to learn a model that we'll later use for prediction



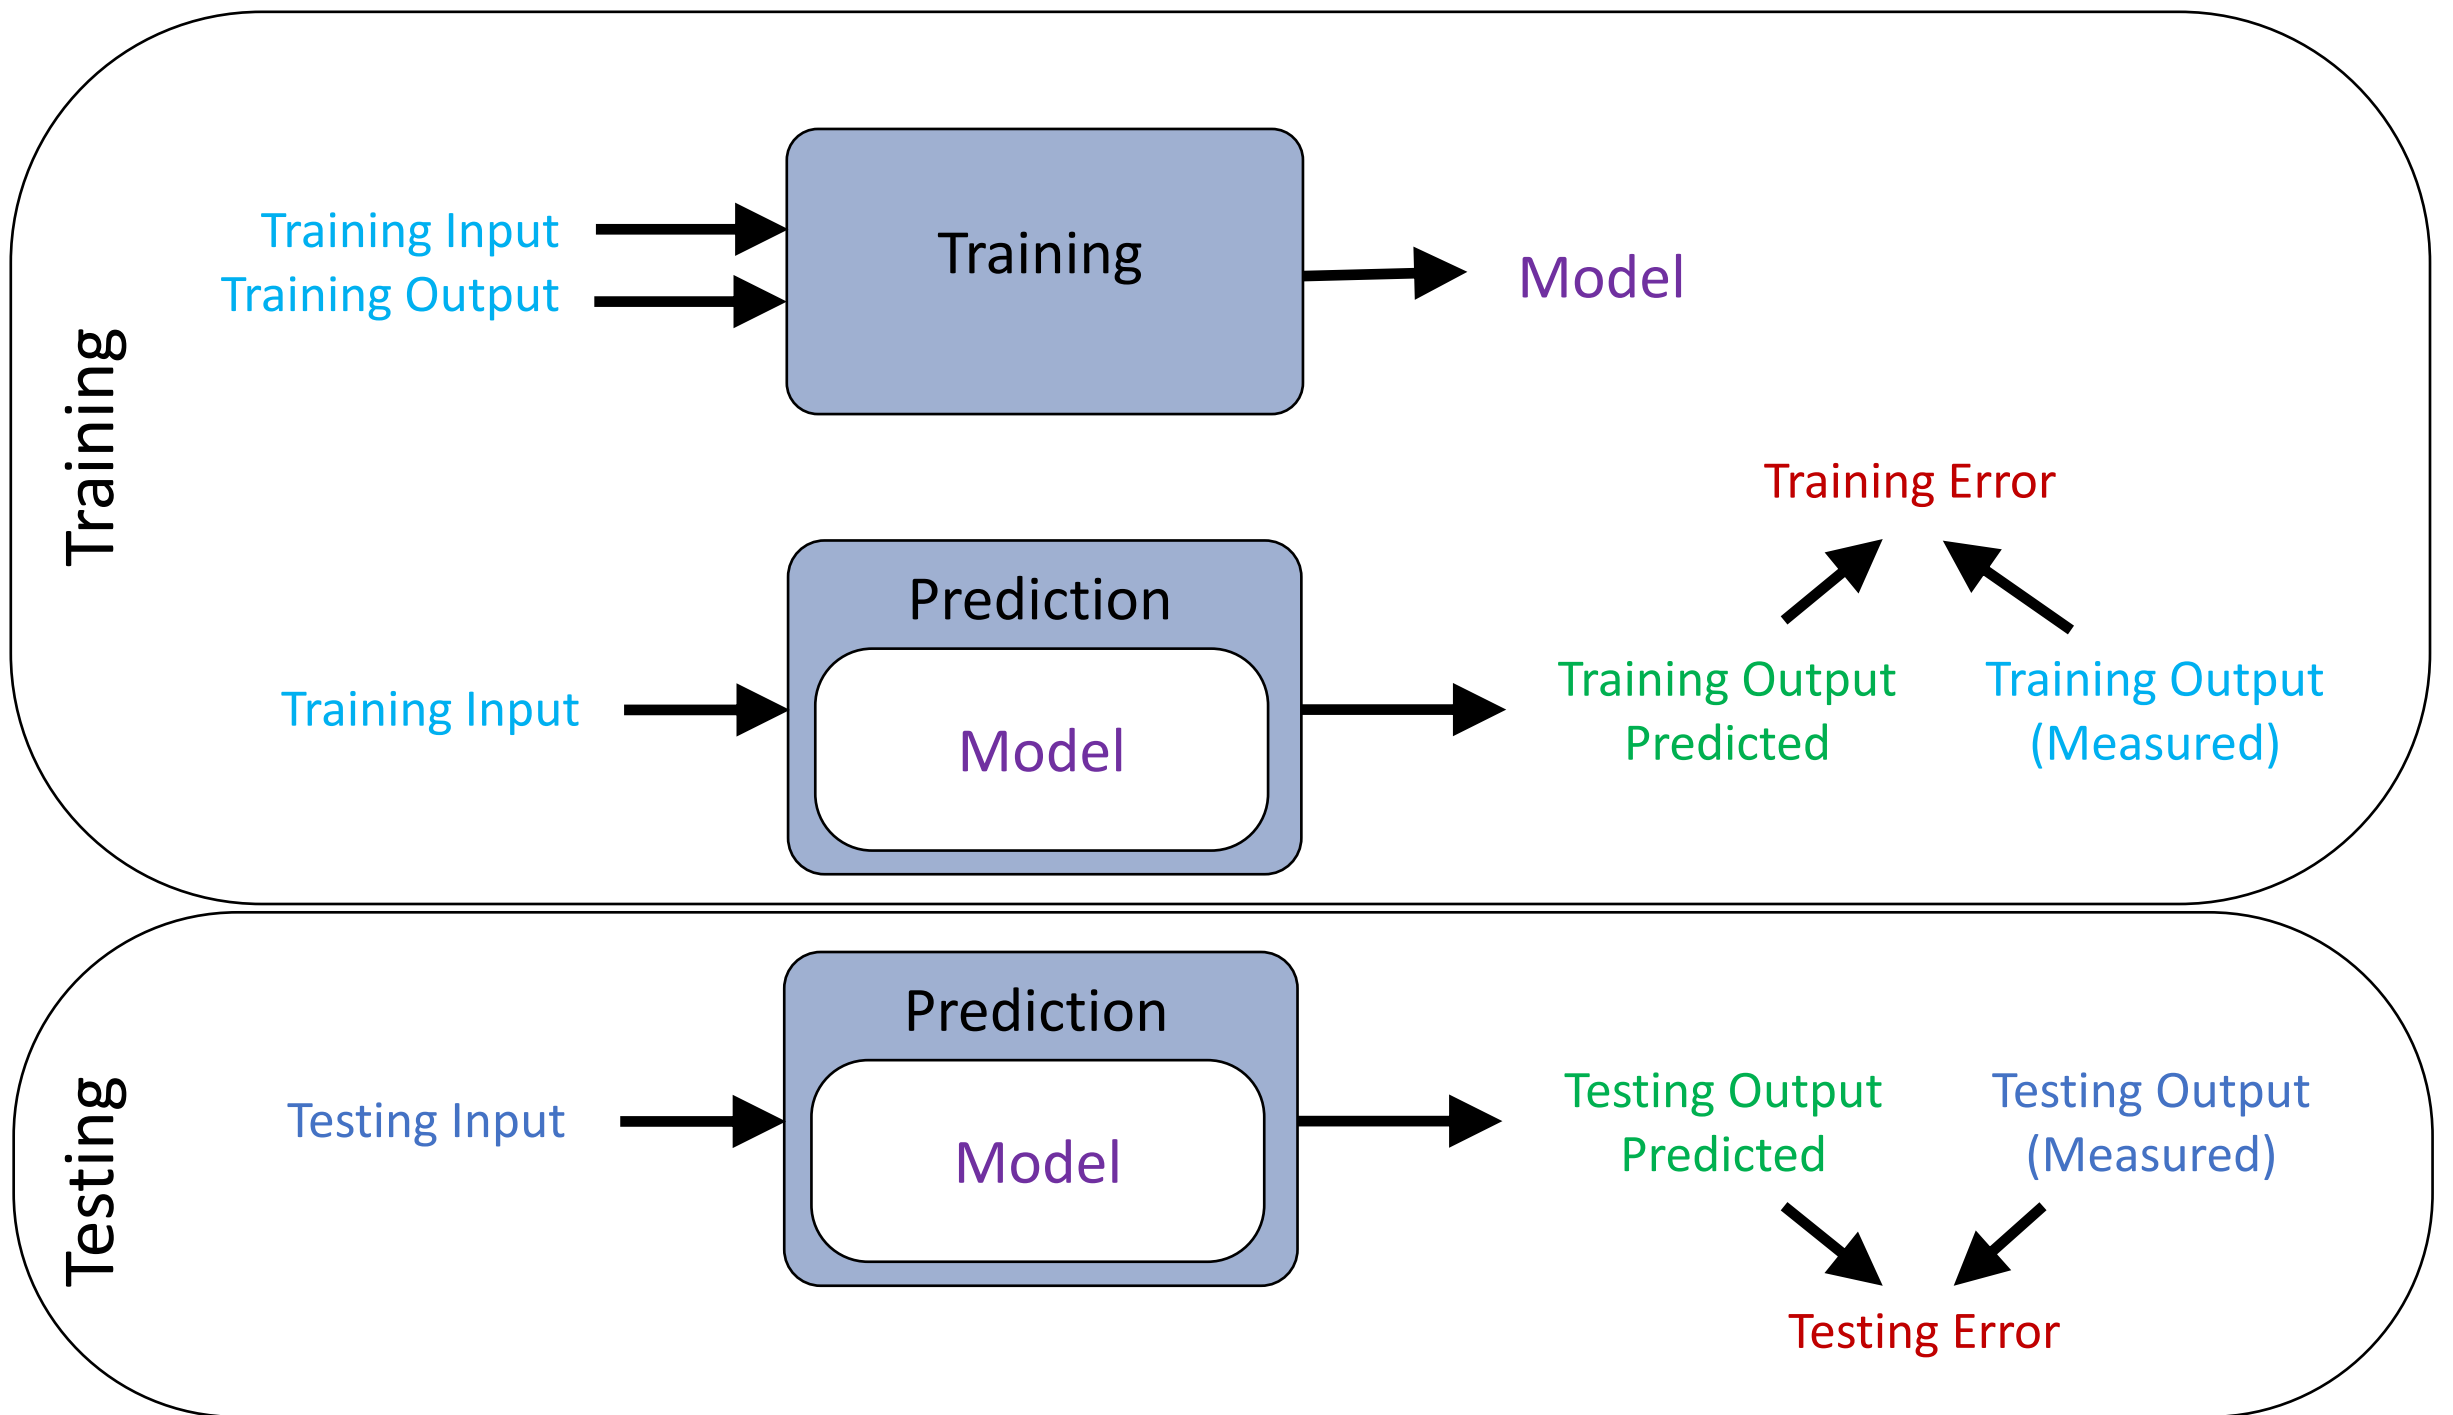


Image Classification Data

Where is the data on the JetBot?