Exploring Functions Using Machine Learning Concepts

Laurel E. Clifford, Mohave Community College

Introduction: The use of artificial intelligence (AI) is increasing exponentially in education, professional fields, security, and popular culture. As our society embraces this tool, we need to understand what AI really is: a tool making predictions from data, models, and algorithms and not actual human intelligence.

Understanding how machine learning trains computers to predict from data can help us make sense of AI and consider its limitations and ethical implications. While much of the mathematical computation behind machine learning may be beyond the precalculus level, the fundamental ideas are accessible to precalculus students and may motivate further study in STEM and AI fields. This module, written for use in an introductory unit early in the precalculus/college algebra course curriculum, explores function concepts and function classification through the lens of machine learning.

Target Audience: College Algebra/Precalculus freshman-entry level math course. This activity could also be used/adapted for Quantitative Literacy-related courses and Math for Elementary Teachers.

Prerequisites: Students should be familiar with the two-dimensional coordinate plane, relationships between inputs and outputs for relations, functions, and one-to-one functions, features of relations and functions, such as domain, range, increasing, decreasing, constant behavior. Students may have some familiarity with basic function types (linear, absolute value, quadratic, cubic), but this is not required. The prerequisite review revisits the concepts of functions, one-to-one functions, domain, and range, and informally previews inverse functions and logarithmic functions.

Topics: Unsupervised and supervised machine learning, relations and functions including one-to-one functions, related function types (linear, quadratic, cubic, absolute value), inverse functions, transformations of functions (reflections, translations, dilations). Test data used in the activity also includes square root, cube root, constant, piecewise functions, and a non-function (vertical line).

Technology and Preparation Needed: This module can be done with little or no technology other than internet access for further exploration. Part II of the activity is available as a copiable and editable Desmos activity at the website https://teacher.desmos.com/activitybuilder/custom/679a4a0bae1473f2b01600c6. For ground-based classes, the Training Dataset and Test Dataset cards for part II, attached at the end of the module, need to be printed and cut out so each group of students has the entire data set, randomly mixed for sorting into groups, and a copy of the Test Dataset for matching into groups. The image classifier extension requires access to the internet and a webcam. Screen-shot images can also be used.

Learning Outcomes:

Determine whether a relationship is function and whether it is a one-to-one function

Discern between function inputs and outputs

Identify from graphical representations the important features that distinguish different function types from each other and determine which function type an unknown function relates to using distinguishing features and function transformations

Discuss strengths and weaknesses of machine learning algorithms and ethical considerations of AI

Timeline: The Prerequisite review and parts I through III can be completed in a single 1 hour 50 minute session or divided into two 50 minute sessions. Extensions can be completed outside of class, or in an additional class session.

Prerequisite Review: The Concept of a Function

Relationships in mathematics can be viewed through a machine analogy. We put an input (x) value into a machine and it generates an output (y) value based on whatever rule the machine uses to do its job.

1. Consider the machine shown below.
   a. Determine the outputs the machine will produce for each of the inputs listed.

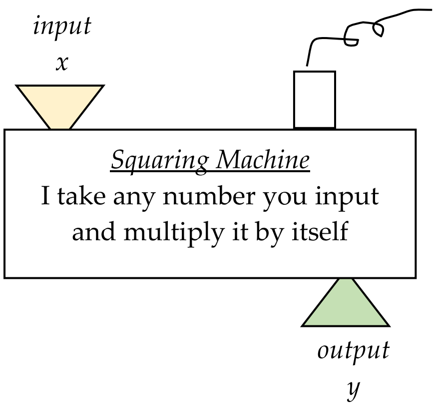

	Input: $x = 1$	Output: $y =$ 1
	Input: $x = -1$	Output: $y =$ 1
	Input: $x = 0$	Output: $y =$ 0
	Input: $x = -4.2$	Output: $y =$ 17.64
	Input: $x = 3/4$	Output: $y =$ 9/16
	Input: $x = 202$	Output: $y =$ 40804
You choose:	Input: $x =$ _____	Output: $y =$ __________

   b. Reflect on the inputs you used and the outputs that resulted.
   For each input you use, will you always get an output out of the machine? Why or why not?
   Yes—every real number can be multiplied by itself and get a real number result (closure)
   Will you always get just one output from the machine? Why or why not?
   Yes—each product of a number by itself produces just one output
   Are there any numerical inputs we cannot use in the machine? Why or why not?
   No—we can multiply any real number by itself
   Are there any numerical outputs we cannot get out of the machine? Why or why not?
   Yes—if we use only real number inputs, we can only get nonnegative outputs

The first two questions,

- Will you always get an output?
- Will you always get just one output?

address the definition of a function as a relationship between two variables where each input has one and only one output.
Function notation allows us to show that a relationship is a function as well as illustrate the dependence of one variable on the other. ***It has the advantage of showing the specific input value we're working with, listed inside parentheses.***

$$f\text{ (input variable value) = output variable value}$$

The next two questions,

- Are there any numerical inputs we cannot use in the machine?
- Are there any numerical outputs we cannot get out of the machine?

address the **domain**, the set of possible inputs (values of the independent variable), and the **codomain,** the set of possible outputs (values of the dependent variable). The **range** is the set of outputs that actually come out of the function, so a subset of the **codomain,** the values that could come out.

The rule this machine followed (multiplying the number by itself) to produce the output is an **algorithm**. In general, algorithms are processes used to make predictions. In machine learning, computers follow processes to analyze data and make predictions. In the squaring machine, we were given the algorithm to use. In other situations, we may need to determine the algorithm or rules to follow to produce the output.

2. Consider the machine below and its given inputs and outputs:
   a. Determine the algorithm the machine is using.

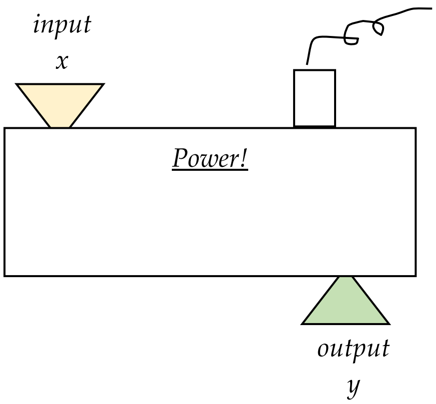

Power(1000, 10) = 3

Power(49, 7) = 2

Power(32, 2) = 5

Power(0.1, 10) = -1

Power(1/81, 3) = -4

Power(7, 49) = 1/2

This machine appears to be using the first input as a result from the second number used as the base of an exponential expression and the output of the machine is the exponent or power which produced the first input

Ask your instructor for one: *Power*(______________) = _____

   b. Reflect on the machine's inputs and outputs.
   What do you notice about the inputs of this machine?
   There are two inputs, the first a result from using the second as a base

   Does this machine represent a mathematical function? Why or why not?
   Yes: for each input pair, we get one output. It has a multivariable input.

3. Consider the machine below and its given inputs and outputs:
   a. Determine the algorithm the machine is using.

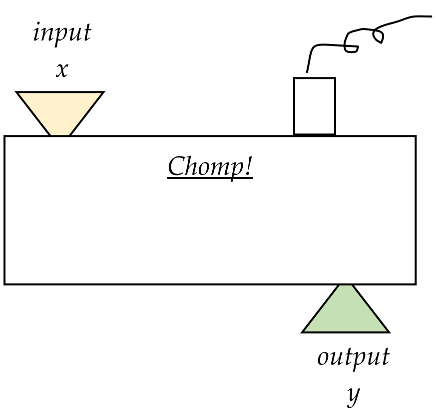

Chomp(Star Wars) = S

Chomp(Forrest Gump) = F

Chomp(Interstellar) = I

Chomp(Saving Private Ryan) = S

Chomp(Nosferatu) = N

Chomp(Wicked) = W

This machine appears to be "chomping' the title of a movie and returning the first letter of the movie title as the output.

Ask your instructor for one: *Chomp*(___________________) = _____

b. Reflect on the machine's algorithm.
How was this machine's process different from those in #1 and #2?
This machine is not a specifically defined math process like multiplication or exponentiation
What is the domain for this machine? The range?
The domain (inputs) are movie titles, the range (outputs) are letters of the alphabet
Does this machine represent a mathematical function? Why or why not?
Yes: each movie title has one and only one first letter
If we know the output for this machine, is it possible to determine what input it came from? Explain.
No—some movie titles have the same output (first letter), so if we have the output letter (like S in the examples) we may not be able to determine which movie it came from (this is not one-to-one)

4. Each of the machines we considered assigned an output for a specific input, and each input had one and only one output, making the results predictable. Based on your experience with the machines, you can determine the outputs for specific inputs:

Squaring Machine: $f(-100) = \underline{10000}$ *Power*(125, 5) = _3__ *Chomp*(Dune) = __D___

Reversing the processes, determining the unknown input which produced a known output, is less predictable. Consider what inputs could exist for each of the following outputs:

Squaring Machine: $f(x) = 9$	*Power*$(x, b) = 2$	*Chomp*(x) = H
x? *x = -3 or 3*	*x, b? x = 16, b = 4*	*x? x = Harry Potter…*
	x = 25, b = 5 etc.	*Happy Gilmore, etc.*

Each of these functions lacks the ability to predict a single input from a given output because they are not **one-to-one** (injective) functions. Functions which are not **one-to-one** are not **invertible**.

I. A Simple Introduction to Machine Learning

"Machine learning is Machine learning (ML) is a subset of artificial intelligence (AI) that focuses on developing algorithms and models that enable computers to learn and make predictions or decisions without being explicitly programmed for a specific task."[1] Machine learning aims to teach a computer to look at data and learn from it like a human being would, without being programmed explicitly for each task.

There are three main types of machine learning:

- **Supervised Learning:** the computer is given a training labeled dataset, where each input is associated with a corresponding output, then trained to classify new data. *This is like having your instructor give you a set of problems like questions 2 and 3 in the previous section and ask you to apply the rule to find new values.*
- **Unsupervised Learning:** the computer is given unlabeled data and looks for patterns and associations within the dataset without any prior knowledge, grouping the data into sets according to their similarities. *Your instructor isn't telling you what the association or rules are, you are looking for them.*
- **Reinforcement Learning:** the computer is trained to make a series of decisions based on trial and error, where the user gives positive or negative reinforcement to the model based on the output, and the computer's goal is to maximize the positive reinforcement. *This is like your instructor giving you a first step, telling you if you're correct or not, and then giving you the next step, and you repeat the process with the goal of getting as much of it correct as possible.*

In each case, the machine is performing a function, taking input data, and assigning an output to the data. In supervised learning, labeled training data teaches the machine to take a given input and assign it a label as an output. In unsupervised learning, the input is an unlabeled data set and the output is data sorted into specific groups.

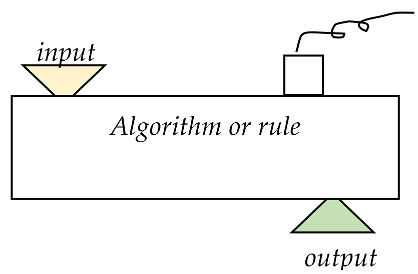

1. The graphic at the right, [*What is Supervised Learning*](www.educba.com/what-is-supervised-learning/)*?*[2] contrasts supervised and unsupervised machine learning.
   a. In the **supervised learning illustration**, what is the training data set and how is it labeled?
   A group of red apples is given as the input data and labeled "apples"

   b. What *input* is the model given?
   A red apple is given as input

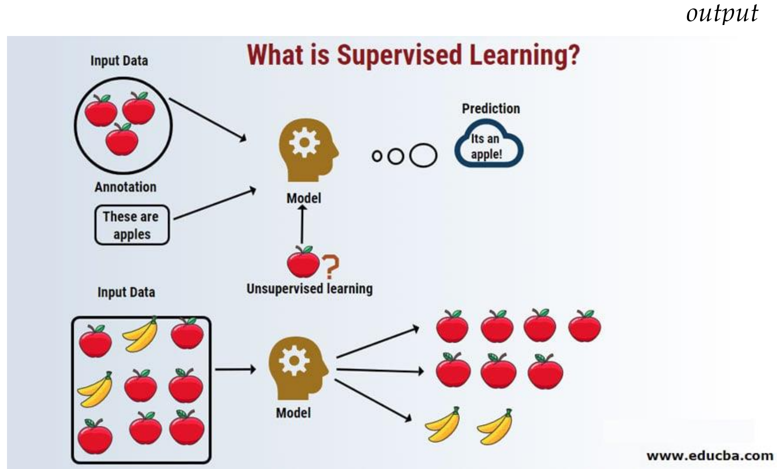

---

[1] From Nigatu, Shumet Tadesse. *Introduction to Machine Learning*, OER Commons, 8 Apr. 2024, Accessed 2025 oercommons.org/courseware/lesson/115010/overview. License: https://creativecommons.org/licenses/by/4.0/

[2] Pedamkar, Priya. "What Is Supervised Learning?" EDUCBA, www.educba.com/what-is-supervised-learning/. (accessed 2025)

c. What *output* does the model produce?
The machine assigns the apple the label “apple” (This is an apple)

d. Suppose we want the machine to identify any apples. How might the training data given to the machine limit the machine’s ability to identify an apple?
All of the apples in the training data are red, so the machine may not recognize green, yellow, or pink apples.

e. In the **unsupervised learning illustration**, what *input* is the model given?
The input is a group of fruit including both apples and bananas

f. What *output* does the model produce?
The output is the fruit sorted into two groups, apples and banans?

g. What do you think the model has learned in the process, and what might be limiting from its input data?
The machine recognizes two different shapes or possibly colors. Since the apples and bananas are different colors, the machine may be classifying only by color, or may be classifying by shape, or both. It may not identify green apples or green bananas correctly, or may identify a round fruit as an apple or an elongated fruit as a banana.

2. The picture on the right shows a herd of animals at a waterhole in Africa.

a. If we use unsupervised machine learning on this image as an input, what do you think the resulting output will be?

The output may be different sets of animals, such as giraffes, lions, gazelles (impala) and wildebeest 4 distinct groups)

b. Which of the images below do you think the model could classify correctly and which incorrectly? Explain your thinking.

There were no zebras in the training data, so this may be seen as one of the impala due to stripes or wildebeest probably not identified correctly

The giraffe matches the training data although facing a different direction.

The donkey may be classified with the wildebeest (close!)

II. Unsupervised vs. Supervised Machine Learning with Functions Activity

Next, we apply the ideas of machine learning with an investigation of function graphs.

1. In our first activity, we explore different function types using methods similar to **unsupervised machine learning** in AI.
   a. The Training Dataset is unlabeled and ungrouped. Your task is to group the graphs into 4 groups based on their features and patterns, looking for commonalities.
   Answers may vary depending on features noted. Sample student responses shown below.
   b. After you have grouped your functions, think about which commonalities are required to belong to each group, and if any commonalities are not needed.

Group A	Group B	Group C	Group D
Linear: Constant slope Different angles One-to one	V-shaped (absolute value) 90 degree angle Have a minimum Constant slope on each side Not one-to-one	U-shape (parabola curve), have a minimum Not one-to-one	(Cubing function) Have a curve that continues One-to-one Split parabola with reflected side Symmetric about the origin

   c. Looking back at your work, how do you think your groups would change if you could only group them into 2 groups stead of 4? *Be specific about the criteria for your two groups.*
   Answers vary; Students considered curved graphs vs. linear graphs
   One-to-one functions vs. not one-to-one functions
   Symmetrical vs. asymmetrical

   d. How would your groups change if you could have more than 4 groups? *Be specific about the criteria for your groups.*
   Answers vary. Students considered graphs located with a minimum at the origin or passing through the origin as another separation criteria

   e. **As a class,** decide on names for your groups of functions and what specific criteria are required for classification within the group.

Group A	Group B	Group C	Group D

2. In our second activity, we are going to look at new graphs in the Test Dataset and decide if we should assign them to a group or not, using the criteria we chose as a class in part (1e) above. Explain the reasoning behind each choice using the class criteria. *If you said a graph did not belong in a group, explain why you excluded it.*

Graph 5 Group? Reasoning: Students placed with other parabolas, despite the training set parabolas all opening up and this one opening down	Graph 6 Group? Reasoning: Students placed with other cubing functions due to its curve and one-to-one shape despite the training set passng through the origin and this one shifted.	Graph 7 Group? Reasoning: Students either didn't want to put in a group, or though closest the linear graph as it was one-to-one
Graph 8 Group? Reasoning: Students placed graph with the cubing function due to its "reflected parabola" curve and one-to-one nature	Graph 9 Group? Reasoning: Students commented on its linearity but didn't group it with lines as it wasn't at an angle and it was not one-to-one (felt it belonged in its own group)	Graph 10 Group? Reasoning: One student wanted to put it with the other lines, but most recognized it didn't match the training set as it was not a function (undefined slope)
Graph 11 Group? Reasoning: Students didn't think this belonged with a group although one commented on its symmetry and felt it was related to parabolas as a result (it is reciprocal squaring function)	Graph 12 Group? Reasoning: Most groups put with absolute value groups but one group commented it didn't belong as it did not have the same 90 degree angle and was opening downward	Graph 13 Group? Reasoning: Groups did not put it with any functions as it didn't have symmetry or constant slopes that the training groups had, despite being one-to-one
Graph 14 Group? Reasoning: Similar arguments to graph 13. One student commented on how all the training curves had some sort of change in direction and these two (13, 14) did not	Graph 15 Group? Reasoning: Groups were split. Half wanted to put this graph with the other absolute value functions as it had the constant sloped sections. Others did not want to put it in any group.	Graph 16 Group? Reasoning: One group put this with the cubing functions due to its increasing nature and being one-to-one. Several groups recognized it had pieces of the different functions.

Reflect:

3. How did you make your decision as to whether a graph belonged in a group? If you did not agree, what criteria did you choose to reconcile the disagreement? *What problematic issues arose in the process?*

Students discussed the criteria they listed and how they used this criteria to the best of their ability. They were bothered when graphs did not match what they had been given in the training set. They also discussed how the location of the graph in relation to the axes themselves influenced their decision. Some were hesitant to classify a translated graph as part of the group if the training group didn't include it.

4. One of the goals of unsupervised machine learning is **dimensionality reduction**, which reduces the number of features (criteria) while keeping the most essential information. Looking back at your function groups in problem 1, what features/criteria did you have identified that you didn't use to determine if a graph belonged within a group?

Things they felt they didn't need to use: domain, where the graph crossed the x-axis, if it crossed at all, if it stayed constant, and where it crossed the y-axis. There was some discussion of the direction graphs opened (up or down) and whether that was a characteristic which would exclude a graph even though the training data may not have included it. For example, a u-shaped parabola made sense to them no matter where it was located or which direction it opened, but they also acknowledged prior experience with the parabola graphs could bias their decision.

5. What other information would you like to have had for the Training Dataset or Test Dataset or both to help make the decision about sorting the graphs into a specific group?

One group would have preferred fewer groups to work with. Another group would have liked to know how much variation in how the patterns matched would be acceptable. One group added knowing whether a function had extrema as part of its defining features would help (more of a supervised approach).

6. How did we use unsupervised and supervised machine learning concepts? *Refer to Part I if needed.*

Students recognized the training set giving more freedom in how they viewed the functions and how prior learning/experience they brought with lines, parabolas may have been created a supervised effect as they were working with the training set.

III. Limitations and Ethical Concerns about Machine Learning and Artificial Intelligence

Suppose we want to teach a computer to identify images of a horse[3]. For a training data set, we give it the set of horses as shown in the lovely image shown below.

The artist states, "This is a chart on the possible colours of Eurasian wild horses or tarpans (*Equus ferus ferus*) - a surprisingly varied population, as shown by ancient dna and cave art. These animals only went extinct around year 1900."

1. What issues in the computer's predictions might arise from using this image as a training data set?

   All the horses in the same stance and direction as well as coming from extinct animals and cave artist interpretations

[3] https://www.deviantart.com/eurwentala/art/Wild-Horses-713222142, Creative Commons Attribution-NonCommercial-ShareAlike 3.0 License, Artist: Eurwenta at deviantart.com.

2. On March 20, 2018, Elaine Herzberg was killed by a self-driving car as she crossed a darkened road in Tempe, Arizona. The incident is described here: https://www.wired.com/story/ubers-self-driving-car-didnt-know-pedestrians-could-jaywalk/. What issues in machine learning for the self-driving car led to the fatal accident?
   The system didn't recognize the pedestrian as a pedestrian as it had not been trained to recognize pedestrians outside a sidewalk. It had to reclassify what it saw repeatedly as a car or other, which delayed its reaction.

3. A 2018 MIT study[4] found, "Examination of facial-analysis software shows error rate of 0.8 percent for light-skinned men, 34.7 percent for dark-skinned women."
   https://news.mit.edu/2018/study-finds-gender-skin-type-bias-artificial-intelligence-systems-0212
   Related sources for facial-recognition bias:
   https://www.scientificamerican.com/article/police-facial-recognition-technology-cant-tell-black-people-apart/, https://foundation.mozilla.org/en/blog/facial-recognition-bias/
   What accounts for this difference in the error rates, and what are the consequences?
   "the data set used to assess its performance was more than 77 percent male and more than 83 percent white" "Across all three, the error rates for gender classification were consistently higher for females than they were for males, and for darker-skinned subjects than for lighter-skinned subjects."
   Consequences: racial profiling and misidentification, psychological distancing between law enforcement and people
4. AI chatbots such as ChatGPT (OpenAI) are being used to a greater extent in the workplace and education to both write content as well as analyze information. These tools show biases. Consider the screen capture of a ChatGPT query performed on the author's phone from June 2024.
   a. There is one small difference between the two question prompts asked. What is this difference?
      Pronoun "she" in the first inquiry, "he" in the seond
   b. How are ChatGPT's replies different? Why do you think this happens?
      Chat GPT says "she" is ambiguous, but when "he" is used, the doctor is assigned the "he" pronoun (and not the nurse). Training data is biased as women have been historically underrepresented as doctors, and men historically underrepresented as nurses.
   c. What are the implications of this bias in the use of AI in our day-to-day activities?
      Gender bias, bias towards underrepresented groups, reinforced stereotypes and discrimination, etc. (link below discusses bias in AI with recommendation letters.

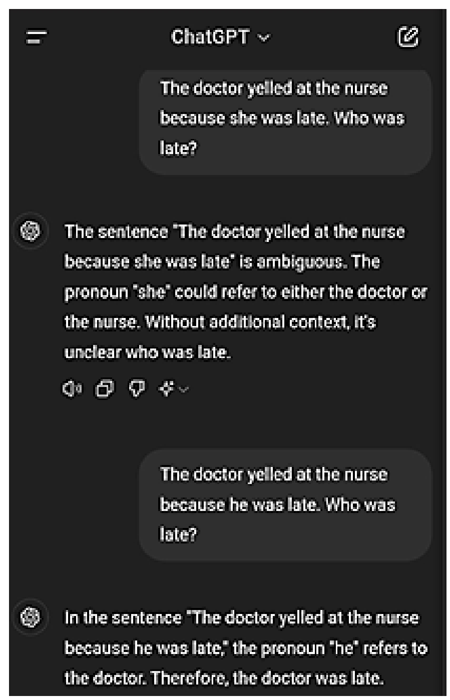

More resources about gender bias in AI:
https://www.scientificamerican.com/article/chatgpt-replicates-gender-bias-in-recommendation-letters/

[4] Hardesty, Larry. "Study Finds Gender and Skin-Type Bias in Commercial Artificial-Intelligence Systems." MIT News | Massachusetts Institute of Technology, 11 Feb. 2018, news.mit.edu/2018/study-finds-gender-skin-type-bias-artificial-intelligence-systems-0212.

IV. Explorations and Extensions

1. How well can AI match graphs? You can train MIT's Personal Image Classifier using a form of supervised machine learning by following the instructions at https://appinventor.mit.edu/explore/resources/ai/personal-image-classifier-part1.

   You can use a handheld graphing calculator to generate graphs and show the calculator window to the webcam, sketch graphs by hand and hold them up to the webcam, or upload images if you use https://classifier.appinventor.mit.edu/oldpic/

   Important: Although you are prompted to use a webcam, you can upload a picture file after assigning a label first, as shown in the image at the right.

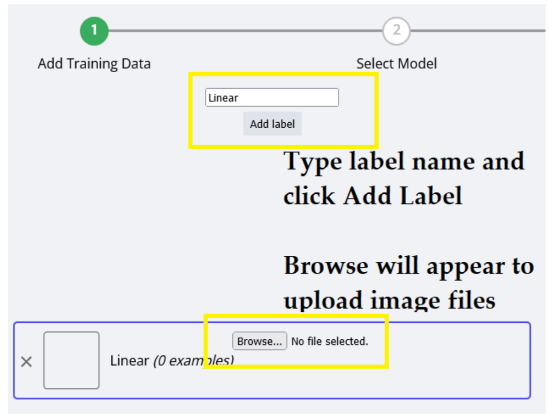

   The screenshot below shows the results from the image classifier after training and testing. The training data consisted of five absolute value functions, including some shifted and reflected over the x-axis, and five quadratic functions (parabolas), with all five parabolas opening upward.

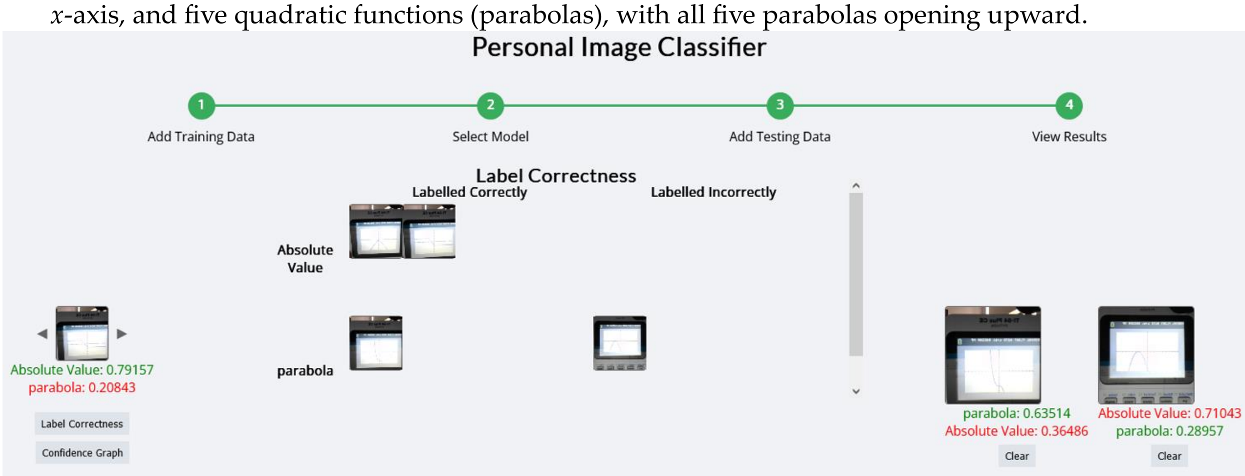

   For testing, the image classifier was given two absolute value graphs: one a linear combination of two absolute value functions, the other an absolute value function reflected over the x-axis. The classifier correctly identified the two absolute value functions.

   The classifier was given a cubing function and a quadratic function reflected over the x-axis. It classified the cubing function as more likely to be a parabola than an absolute value function, possibly due to its non-linearity. It classified the reflected quadratic as an absolute value function instead of a parabola since the training data for the parabolas did not include any reflected over the x-axis--all opened upward.

   Choose two function types and create and upload images of their graphs for a training dataset. Create test graphs using the same functions but different transformations from the training data. Use the Predict button to observe the results.

Discuss the accuracy of the predictions and why you think they labeled the test data correctly and/or correctly. If all your graphs are correctly identified, try creating one that could "trick" the model into misidentifying the function and discuss why. Finally, discuss whether it is possible to make the model predict with perfect (100%) accuracy.

2. Use a free artificial intelligence agent such as ChatGPT, Microsoft Copilot, Meta AI, Hugging Chat, Google Gemini, etc. to make queries to see if you can reveal AI's racial, gender, age, marital status, or other biases.

3. Explore other AI limitations, such as spelling (https://techcrunch.com/2024/03/21/why-is-ai-so-bad-at-spelling/) images issues, such as drawing six fingers on a hand, and AI hallucinations. What causes these issues and what are the ethical implications?

4. Discuss the unexpected results from using AI shared in stories at Machine Learning for Kids, such as "Machine learning doesn't always learn what we intend it to" https://machinelearningforkids.co.uk/#!/stories/learning-the-wrong-thing and "What happens when machine learning models are used for things they weren't trained to do" https://machinelearningforkids.co.uk/#!/stories/models-learn-to-do-specific-jobs

References:

Images without reference are public domain

GeeksforGeeks. (2025, January 24). *Machine Learning tutorial*. GeeksforGeeks. https://www.geeksforgeeks.org/machine-learning/

Hardesty, Larry. "Study Finds Gender and Skin-Type Bias in Commercial Artificial-Intelligence Systems." MIT News | Massachusetts Institute of Technology, 11 Feb. 2018, news.mit.edu/2018/study-finds-gender-skin-type-bias-artificial-intelligence-systems-0212.

Kwiatkowski, S. (2019, March 26). *Machine learning from scratch*. Medium. https://medium.com/towards-data-science/machine-learning-from-scratch-part-1-76603dececa6

Machine learning for kids. (n.d.). https://machinelearningforkids.co.uk/
An educational tool for teaching *kids* about *machine learning*, by letting them train a computer to recognize text, pictures, numbers, or sounds

Nigatu , Shumet Tadesse. *Introduction to Machine Learning*, OER Commons, 8 Apr. 2024, oercommons.org/courseware/lesson/115010/overview. https://creativecommons.org/licenses/by/4.0/

Pedamkar, P. (2024, May 13). *What is Supervised Learning?* EDUCBA. https://www.educba.com/what-is-supervised-learning/

Training Dataset (*Letters for reference in discussion; cards should be randomly sorted*)

A

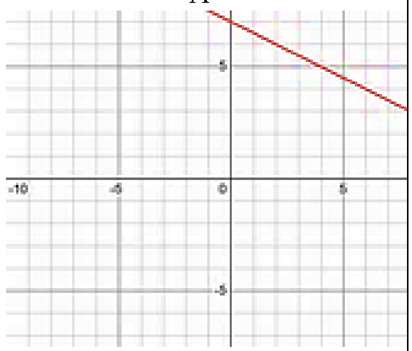

B

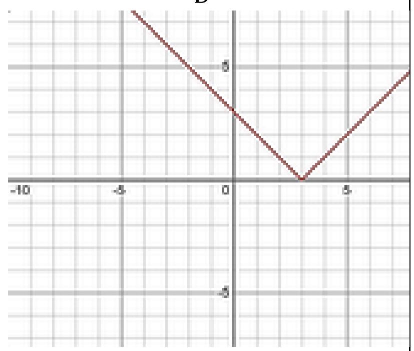

C

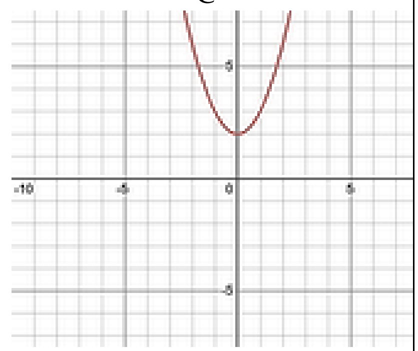

D

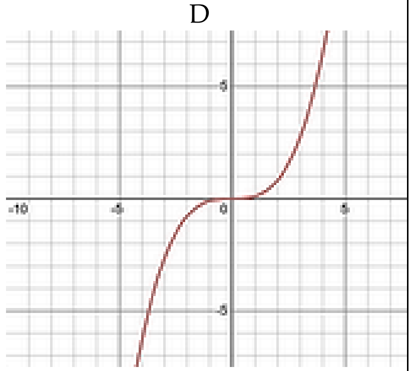

E

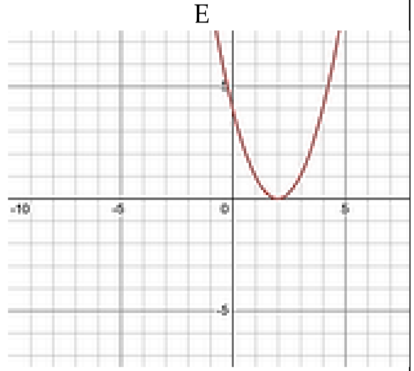

F

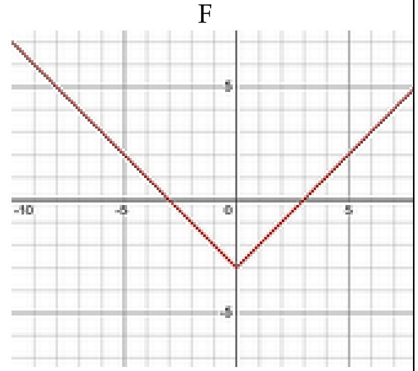

G

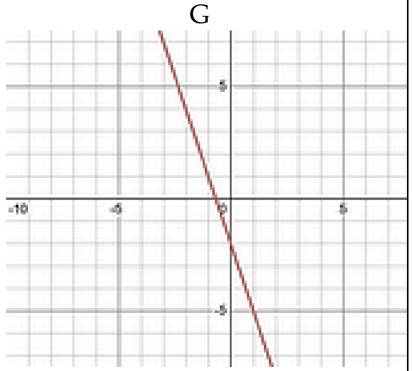

H

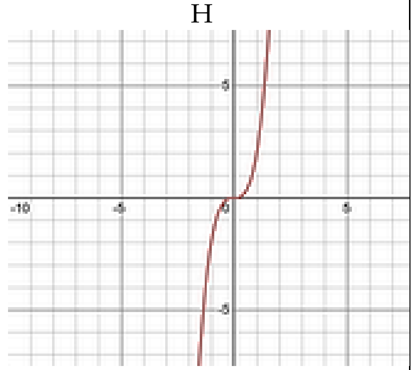

I

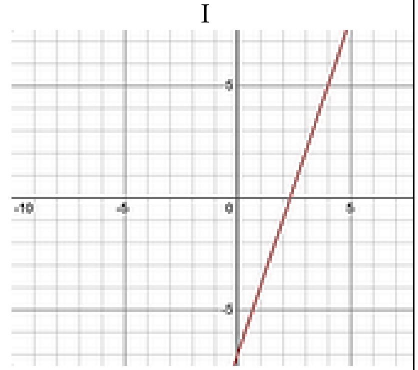

J

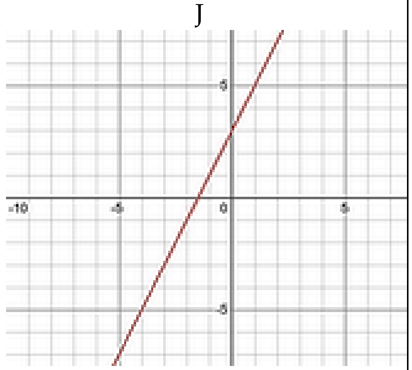

K

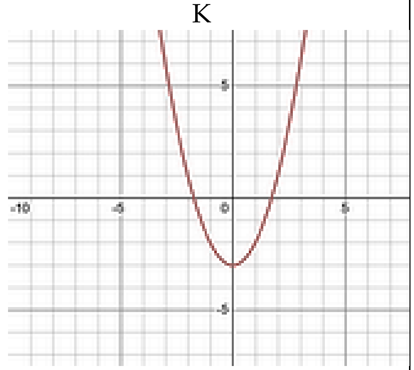

L

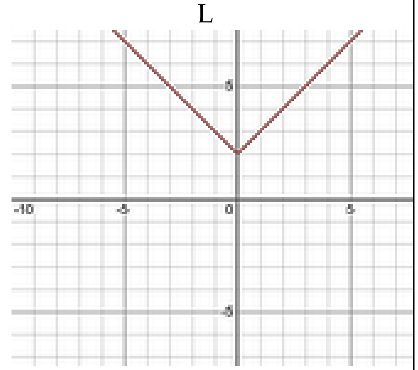

Training Dataset continued.

M

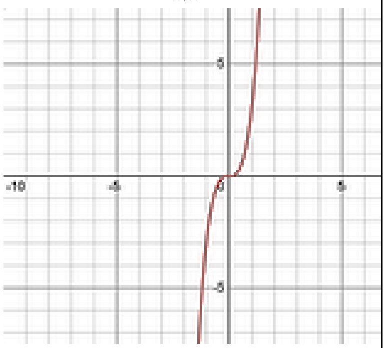

N

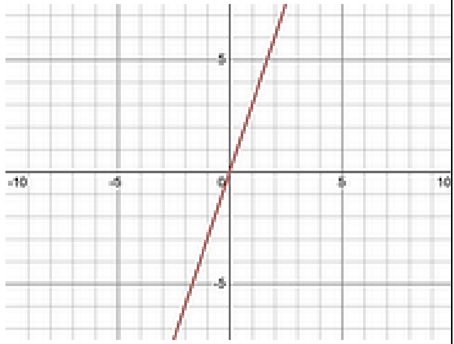

O

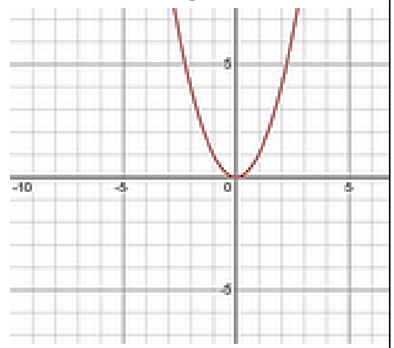

P

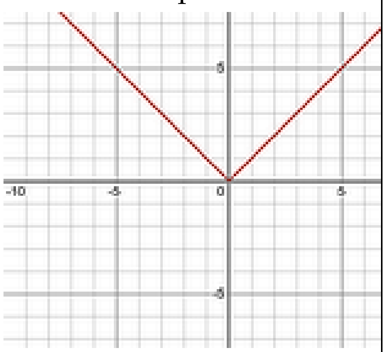

Q

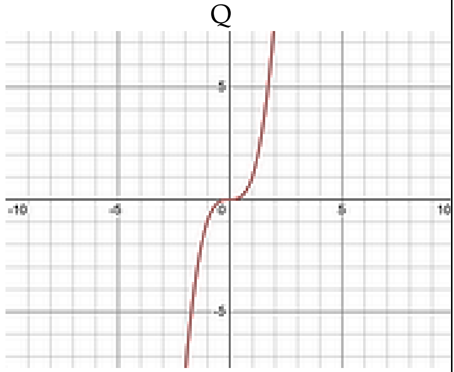

R

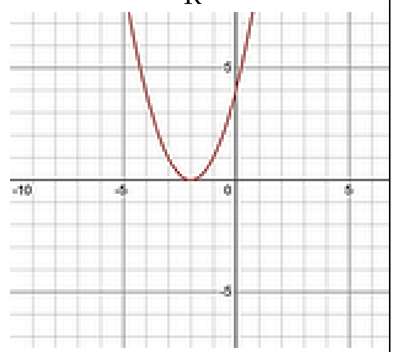

S

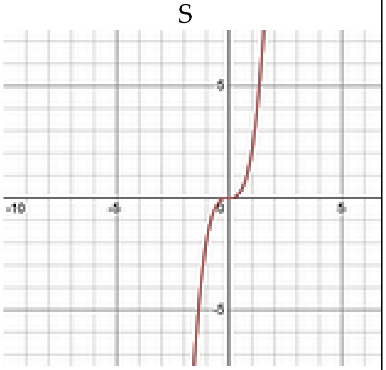

T

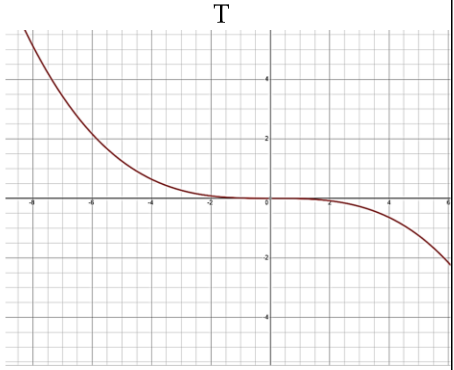

U

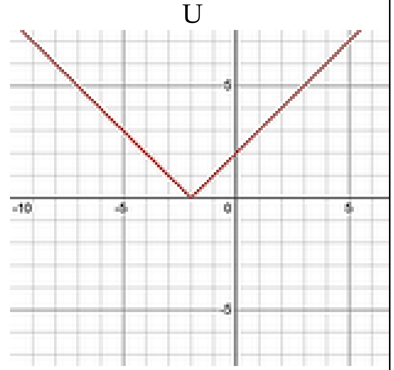

Testing Dataset *(Graph numbers match slides in Desmos activity)*

Graph 5

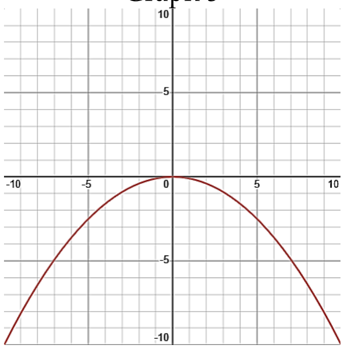

Graph 6

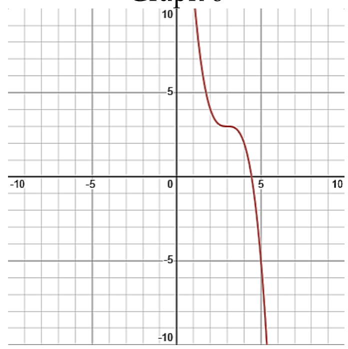

Graph 7

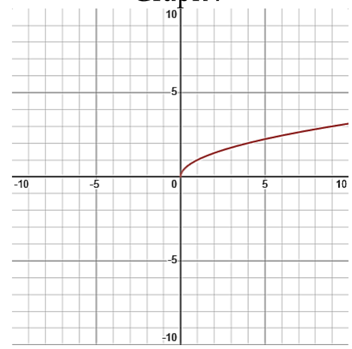

Graph 8

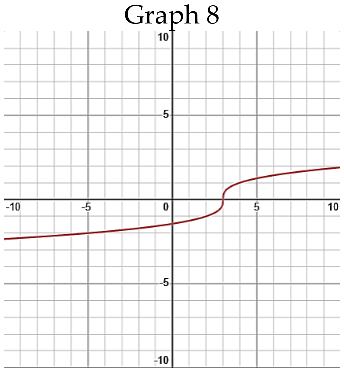

Graph 9

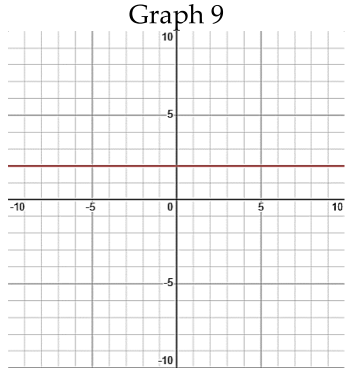

Graph 10

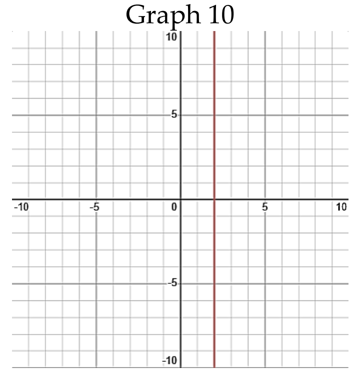

Graph 11

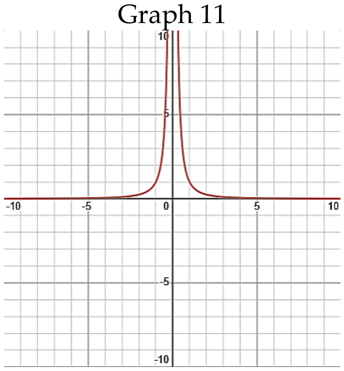

Graph 12

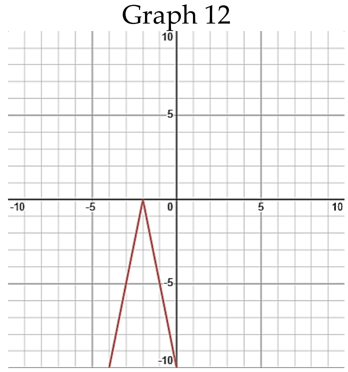

Graph 13

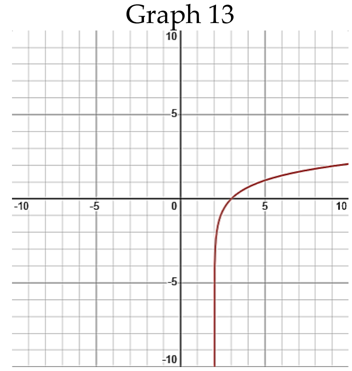

Graph 14

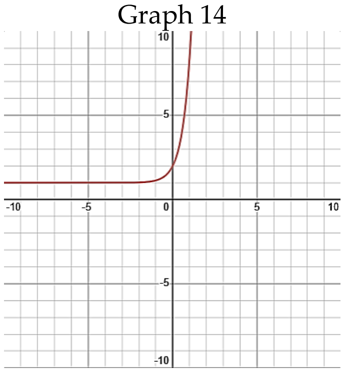

Graph 15

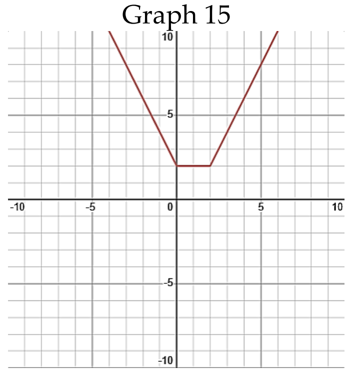

Graph 16