

Anti-Cyberbullying Assistant

By: Bilal Qadar and Nima Boscarino

Duration: 60-90 minutes

LEVEL	SUBJECTS	PROVINCES / TERRITORIES	TOOL
Grades 4-6	Science and Technology, Social Sciences, Applied Design, Skills, and Technologies	Across Canada	Scratch, Machine Learning for Kids

Overview

In this lesson, learners will create an AI-powered assistant that will help them identify cyberbullying.

Prep Work

Technology (per learner):

- Laptop

Before the lesson...

- Review lesson plan
- Create a free teacher account on machinelearningforkids.co.uk, which allows educators to have managed class accounts (optional)
- Work through the project to get familiar with the software, and ensure everything is working correctly
- Review key coding concepts and ensure you are confident describing them to your group

Key Coding Concepts

- ✓ Artificial Intelligence
- ✓ Algorithms
- ✓ Conditional Statements

Terminology

Artificial Intelligence

Ability of a machine or computer program to think and make decisions.

Machine Learning

Machine learning algorithms are algorithms with the ability to automatically learn and improve from experience without being explicitly programmed.

- ❑ Print the **solution sheet**: bit.ly/MLsolutionkey
- ❑ Print “The Social Butterfly Effect”:
bit.ly/clc-cyberbullying-comic

Tutorial

Watch *A Detailed Introduction to Machine Learning for Kids* (video): bit.ly/MLforkids-video (optional)

Lesson

Introduction

Programming usually involves writing out clear and unambiguous instructions for a computer to follow, in order to solve a problem or complete a task. For some complicated tasks, however, computer scientists have developed “machine learning” **algorithms** that can discover patterns by analyzing huge amounts of **data**.

Machine learning is present in many technologies in today’s world! Many people may be familiar with voice assistants in our devices, like Siri and Alexa. Machine learning is often used for applications where developers want to replicate human actions, such as playing a video game, finding routes on a map, understanding text, and more!

Like humans, these kinds of systems learn through experience and practice. They train on data to improve and we can retrain them as necessary. For example, we might retrain the system once we get access to new data. A trained machine learning algorithm is called a model.

(OPTIONAL) Let’s Go on a Picnic

To introduce your learners to the way the machine learning algorithms work, have them simulate the data-collection and model training process by repeatedly guessing which items are allowed on a picnic. The instructions for the activity can be found in the **solution sheet**: bit.ly/MLsolutionkey.

Data

Pieces of information usually formatted in a specific way (e.g. numbers, pictures, and words)

Algorithms

A step-by-step set of operations to be performed to help solve a problem.

Conditional Statements

Making decisions based on conditions i.e. if some condition is met do something, else do nothing or something else.

Curricular Connections

To see how this lesson maps to provincial curricula, visit bit.ly/CLClessons.

References

Machine Learning for Kids
Machinelearningforkids.co.uk

The “Social Butterfly” Effect
bit.ly/clc-cyberbullying-comic

Activity

Say: "In this lesson, we'll be building a machine learning model that will help us identify cyberbullying."

Do: Read *The "Social Butterfly" Effect* (bit.ly/clc-cyberbullying-comic) together and follow the prompts at the end of the story."

Do: Use the **solution sheet** (bit.ly/MLsolutionkey) to guide learners through the following steps:

- Adding Data and Training
 - Creating a new project
 - Adding training labels
 - Adding the comic book data
 - Training and testing the Machine Learning model

- Creating the assistant in Scratch 3.0
 - Setting up the project
 - Deleting and adding a sprite
 - Resizing characters
 - Asking a question
 - Recognizing Nice and Mean phrases
 - Turn that frown upside down!

Reflection

Do: Have learners Think-Pair-Share (bit.ly/think-pair) the following questions:

What are/is:

- 3 new things that you learned today
- 2 areas that challenged you in this lesson
- 1 way you can use something from this lesson in your own life

Assessment

Learning Outcomes

I can find patterns in data

I can train an algorithm using supervised learning

I can use conditional statements to control what happens in my project

I can use machine learning to tackle cyberbullying

Assessment Ideas

Encourage learners to tinker with their models to verify that mean and nice phrases are appropriately categorized. If some phrases are miscategorized, learners may need to add to their training data and retrain their models.

Have learners demonstrate an understanding of the parts of a machine learning project, including:

- Data collection
- Training and testing a model
- Using the model in an appropriate application (e.g. a scratch Scratch project)

Extensions

Complete the add-ons in the solution sheet (e.g. happy spin and bounce):

bit.ly/MLsolutionkey

For advanced learners, you can challenge them to make the Scratch logic a bit more complicated, especially if the learners have prior Scratch experience. As a couple of examples:

- The Anti-Cyberbullying assistant could be extended to suggest nice phrases if someone enters a mean phrase.
- The Anti-Cyberbullying assistant could display the *confidence* level of its classification. This is a measure of how sure it is that a phrase is mean or nice. (*Hint: There's a special block for this!*)
- If the assistant isn't very confident in the classification of a phrase (e.g. less than 20%), it could ask the user to select which label it should be classified as, add that new phrase to the training data, and retrain the model. (*Hint: There are blocks for this, too!*)

TACKLING CYBERBULLYING WITH MACHINE LEARNING

SOLUTION KEY

LET'S GO ON A PICNIC ACTIVITY

(In breakout rooms of maximum 5 learners)

- 1) Arrange students into a "circle", about 5 people each.
- 2) Choose 1 mentor to be the picnic organizer/trainer.
- 3) Have this mentor choose a rule that all objects being brought on the picnic must follow (See below for rules). An example rule could be that the object must be yellow. Make sure no one else in the circle knows what the rule is!
- 4) The mentor should start by saying 'I am bringing ----- on my picnic.' The blank should be filled by an object that follows the rule they created.
- 5) Go around the circle, one learner at a time. Each student says "Can I bring _ on your picnic?"
- 6) The picnic organizer should respond with "yes" or "no".
- 7) The game should continue until most students have discovered the rule. Remind students who have discovered the rule that they shouldn't reveal it.
- 8) Once everyone has figured out the rule and time permits, perform another round.

Round 1 Rule: Anything Orange

Round 2 Rule: Anything Spherical/Round

Round 3 Rule: Anything made of paper

Round 4 Rule: Anything which first letter starts with "O"

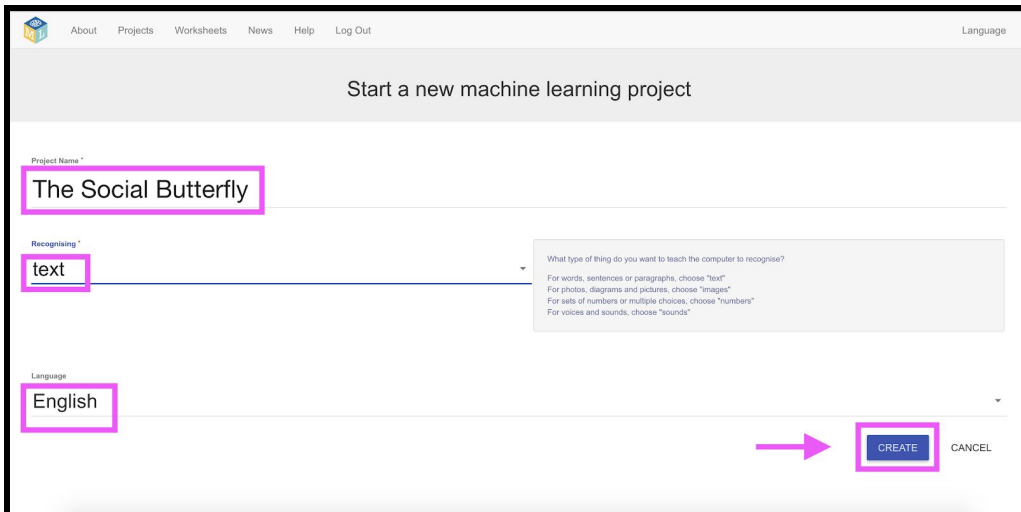
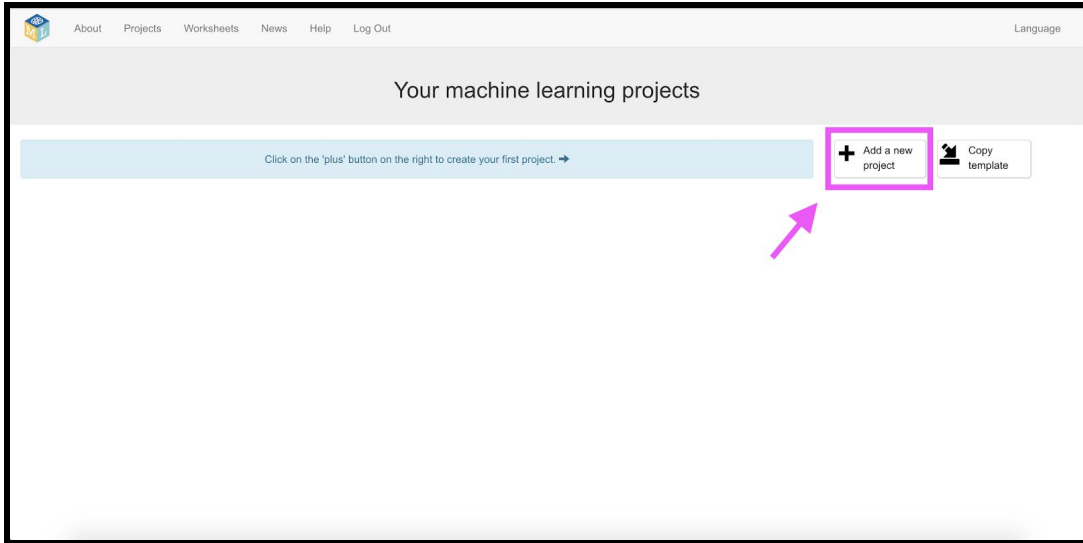
ADDING DATA & TRAINING

STEP 1: Creating a New Project

Start with this link: <https://machinelearningforkids.co.uk>

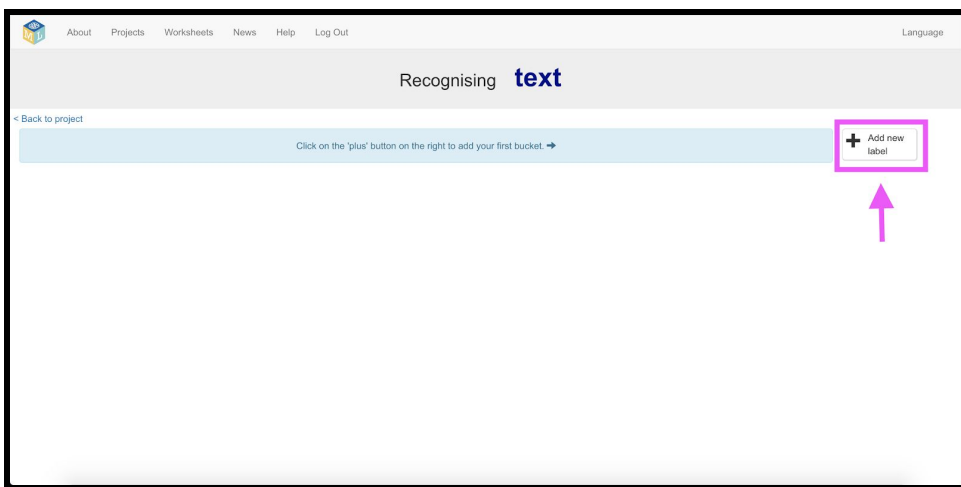
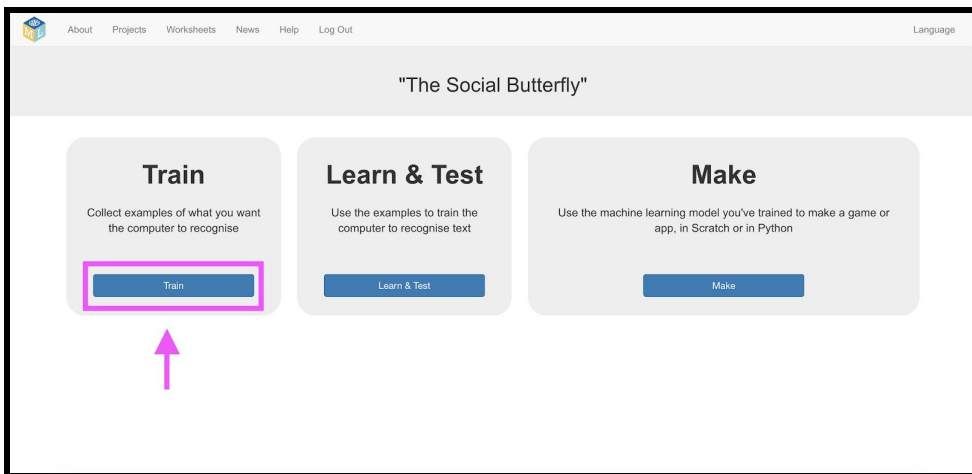
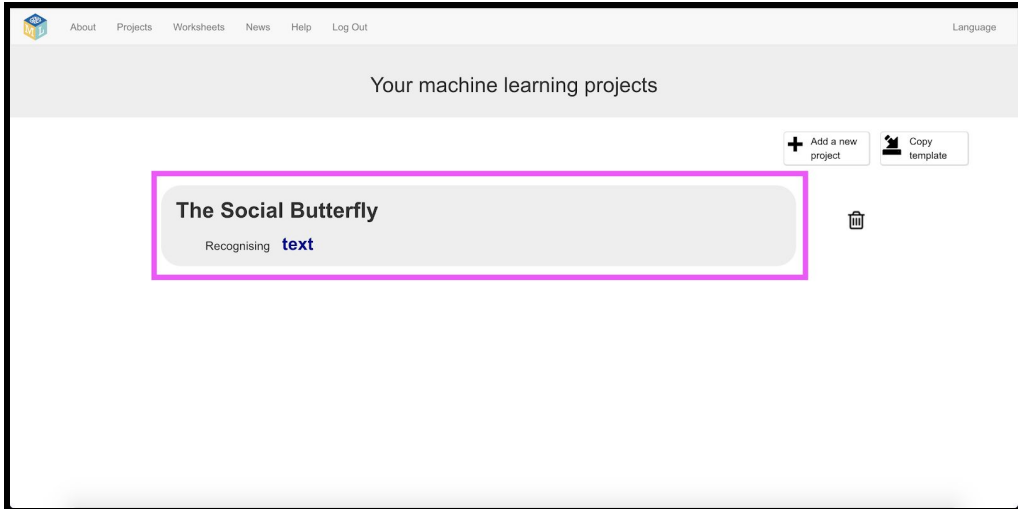
→ **Go to your Projects** if you've already been on the website OR **Get started** → **Try it Now**

1. Click on the 'add new project' button
2. Ensure the project is named and is recognizing text and either English or French depending on region the workshop is being conducted, click create



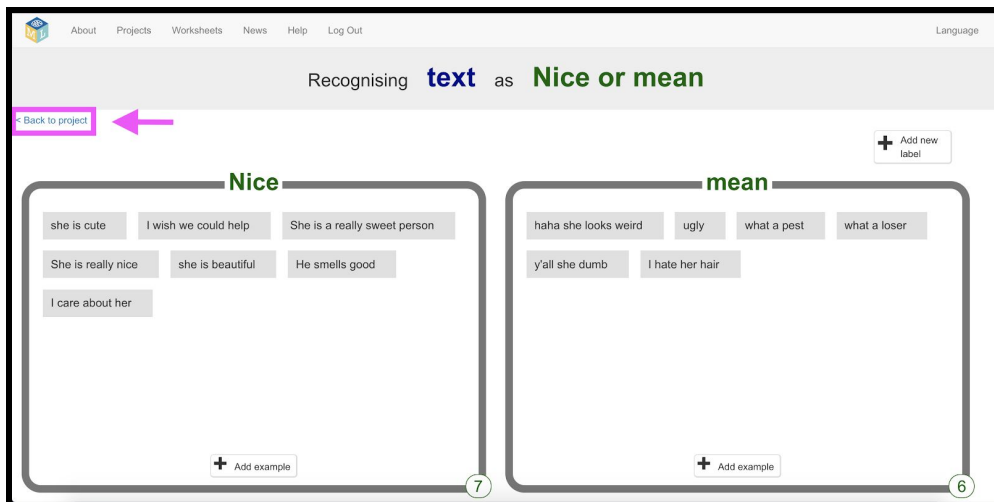
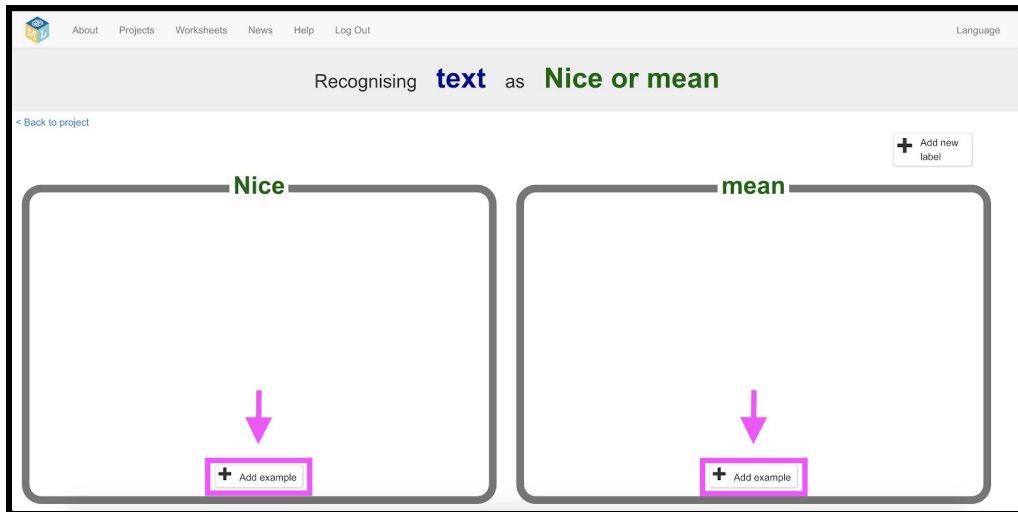
STEP 2: Adding Training Labels

1. Click on the newly created project and then click on the train button
2. Add two labels called 'mean' and 'nice'. This is where we will be adding our comic book data
3. Make use of smaller breakout rooms. Time: No more than 5 minutes.



STEP 3: ADDING THE COMIC BOOK DATA

1. Add data from the comic book using the 'Add Example' button
2. Encourage learners to think of their own 'nice' phrases but **to only add mean sentences from the comic book**
3. Once at least **6** mean and nice examples have been added click on the 'back to project' button
4. Example nice data: [she is cute, I wish we could help, she is a really sweet person, she is really nice, she is beautiful, he smells good, I care about her]
5. Example mean data: [haha she looks weird, ugly, what a pest, what a loser, y'all she dumb, I hate her hair]

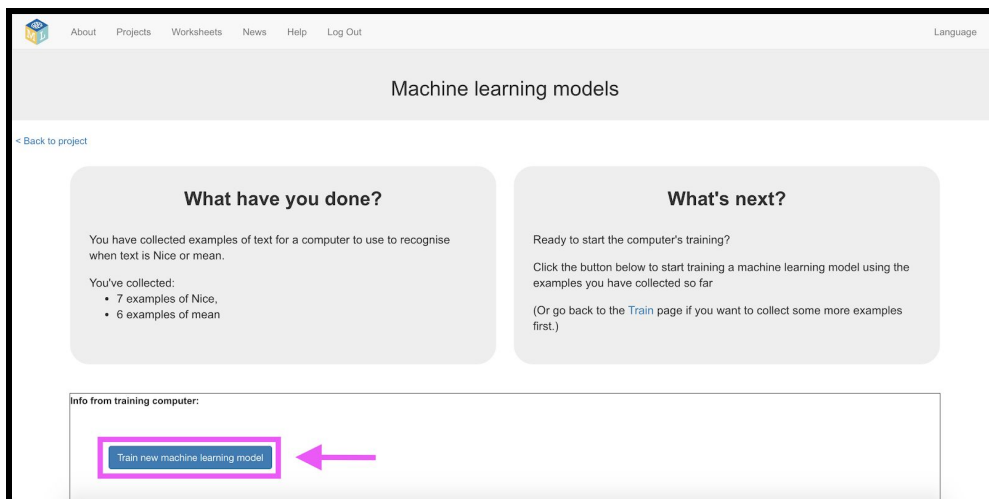
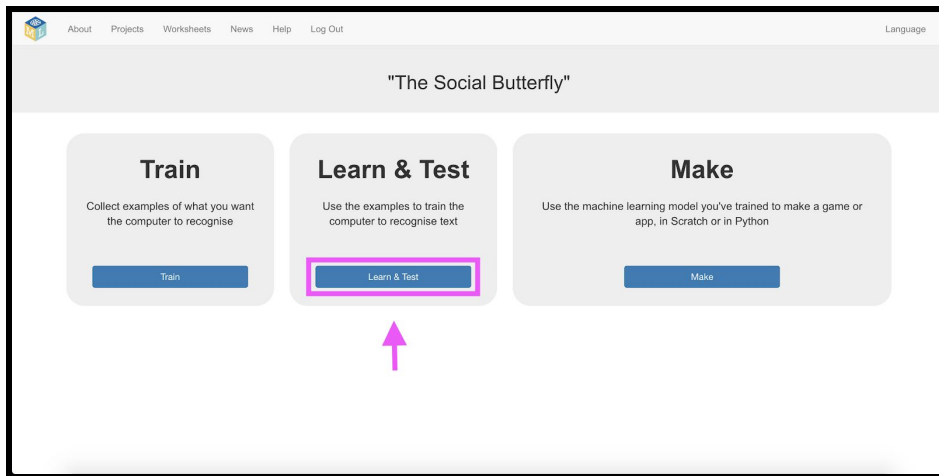


STEP 4: Training & Testing the Machine Learning Model

1. Click on the 'Learn & Test' button
2. Select 'train new machine learning model'
3. Keep in mind this may take up to 5 minutes. The more examples added the longer the process will take

- Once the model has been trained a dialogue box for testing will appear. Learners can try out their model to see if the computer is guessing correctly
- Make use of smaller breakout rooms. Time: No more than 7 minutes.

NOTE (Optional): Overfitting data is when you select examples that are too specific to a given case. If you just use "Amira" when training the computer will not understand "he or she"



What have you done?

You have trained a machine learning model to recognise when text is Nice or mean.

You created the model on Friday, November 29, 2019 12:31 PM.

You have collected:

- 7 examples of Nice,
- 6 examples of mean

What's next?

Try testing the machine learning model below. Enter an example of text below, that you didn't include in the examples you used to train it. It will tell you what it recognises it as, and how confident it is in that.

If the computer seems to have learned to recognise things correctly, then you can go to Scratch and use what the computer has learned to make a game!

If the computer is getting too many things wrong, you might want to go back to the [Train](#) page and collect some more examples

Once you've done that, click on the button below to train a new machine learning model and see what difference the extra examples will make!

Try putting in some text to see how it is recognised based on your training.

Recognised as **mean**
with 82% confidence

Info from training computer:

CREATING AN ASSISTANT WITH SCRATCH 3.0

STEP 1: Setting up the project

1. Click on the 'Make' button
2. Select Scratch 3 button

About Projects Worksheets News Help Log Out Language

"The Social Butterfly"

Train

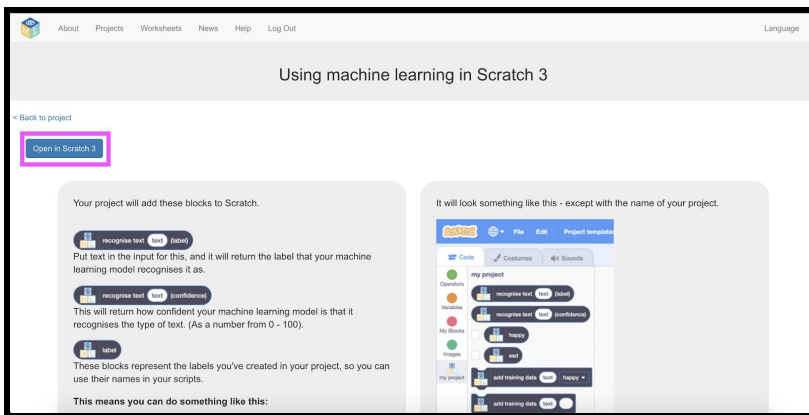
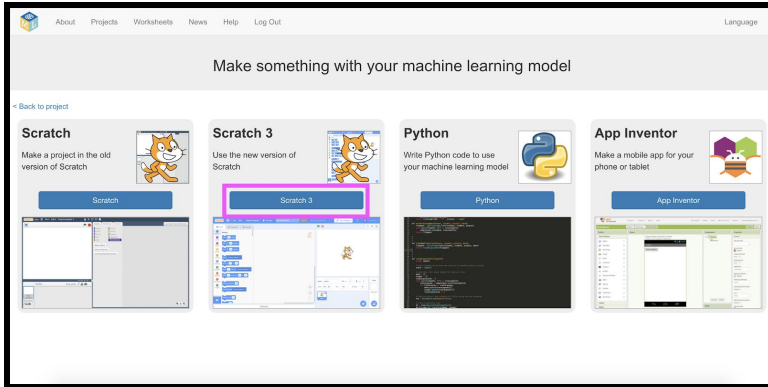
Collect examples of what you want the computer to recognise

Learn & Test

Use the examples to train the computer to recognise text

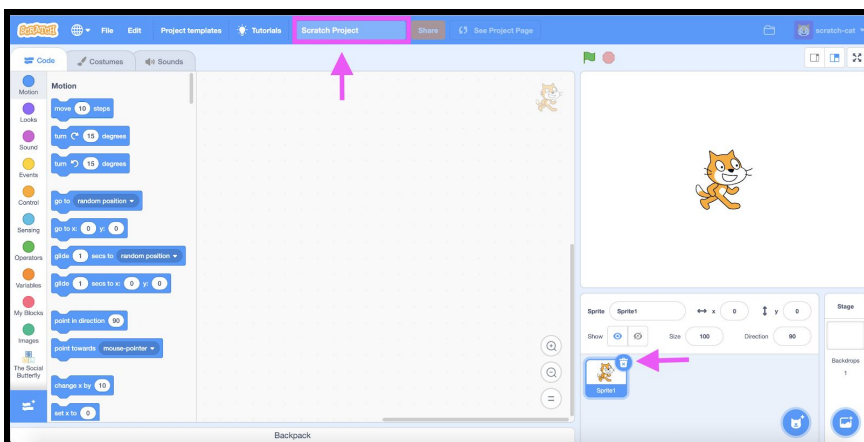
Make

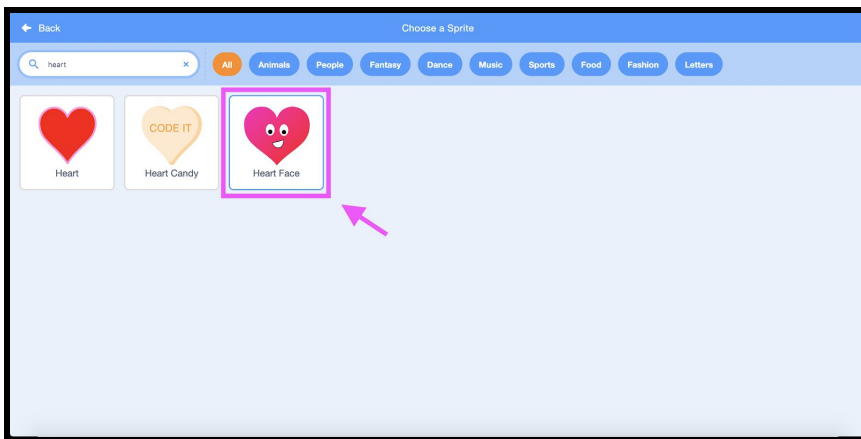
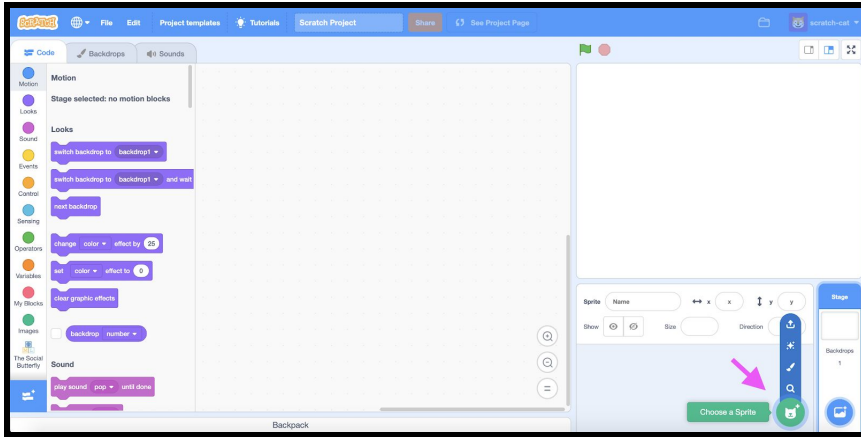
Use the machine learning model you've trained to make a game or app, in Scratch or in Python



STEP 2: Deleting & Adding a Sprite

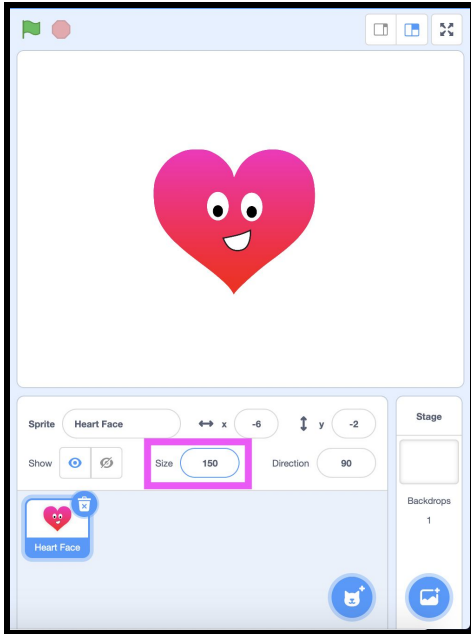
1. Rename the project
2. Delete scratch the cat by clicking on the trash can icon
3. Click on the 'choose a sprite' icon and select a new character





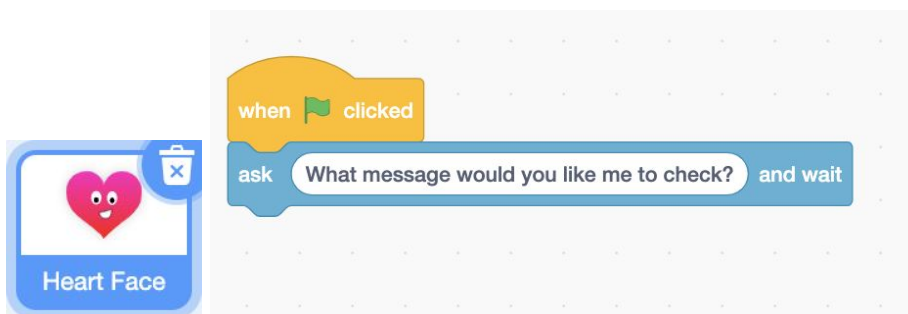
STEP 3: Resizing Characters

1. Click on the 'Size' button under the sprite
2. Type a larger number to increase the sprites size or a smaller number to decrease the size



STEP 4: Asking a Question

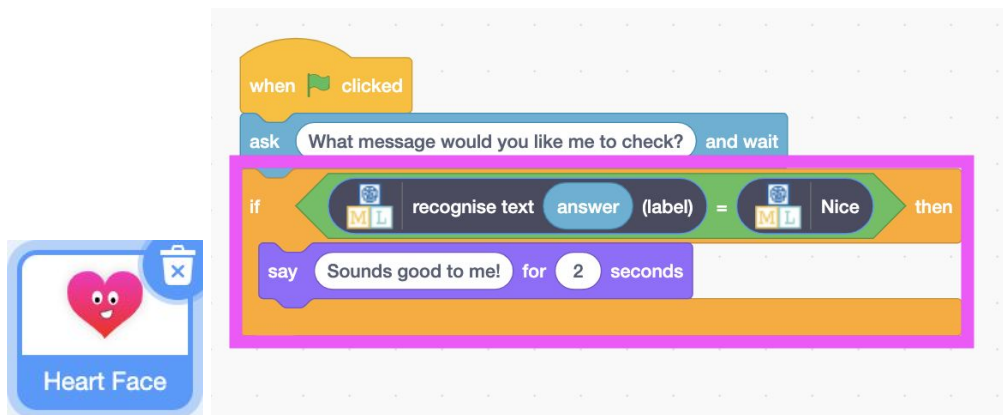
1. Ask: What is the first thing we need our intelligent assistant to do? (A: We need our assistant to ask a question)
2. Ask: What blocks do we need to do this? (A: Ask and wait block and a when green flag is pressed)
3. From the 'Sensing' category, add a 'ask and wait' block
4. Type in a guiding message such as 'What message would you like me to check?'
5. From the 'Events' category, add a 'when green flag clicked' block at the beginning of the program



STEP 5: Recognizing Nice Phrases:

1. From the 'Control' category, add an 'if then block'
2. Ask: What should our assistant say if a phrase is nice? (A: Something positive)

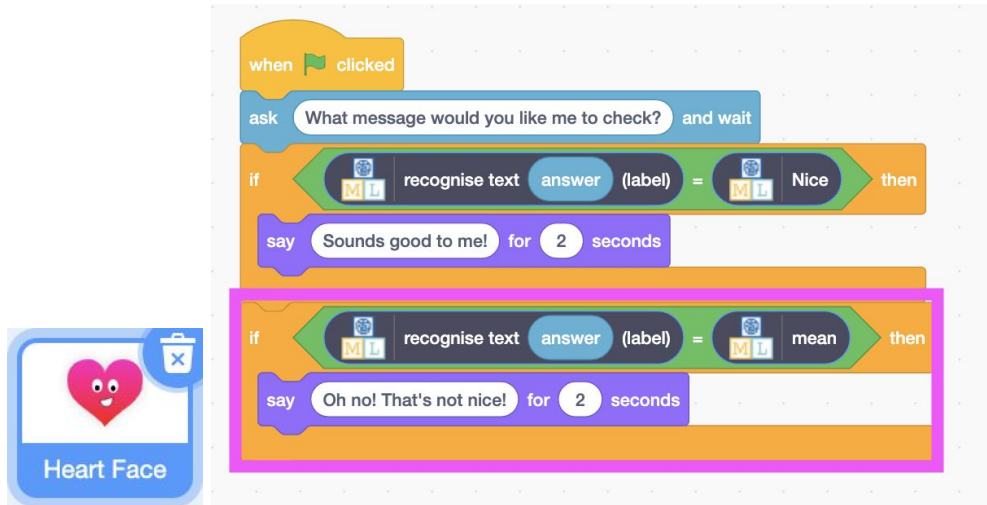
3. From the 'Looks' category, add 'say' block inside of the 'if then block'. Type something positive such as "Sounds good to me!"
4. Ask: Look in the Machine Learning menu. How can we check if a phrase is nice? (A: Recognise text block and Nice block)
5. Drag out the recognize text block and nice block
6. Add answer inside the 'recognize text label'
7. From the 'Operators' category, drag out an equal comparison block
8. Add 'recognize text' and 'Nice' blocks in the equal comparison block
9. Place the equal comparison block in the "if then block"



STEP 6: Recognize Mean Phrases



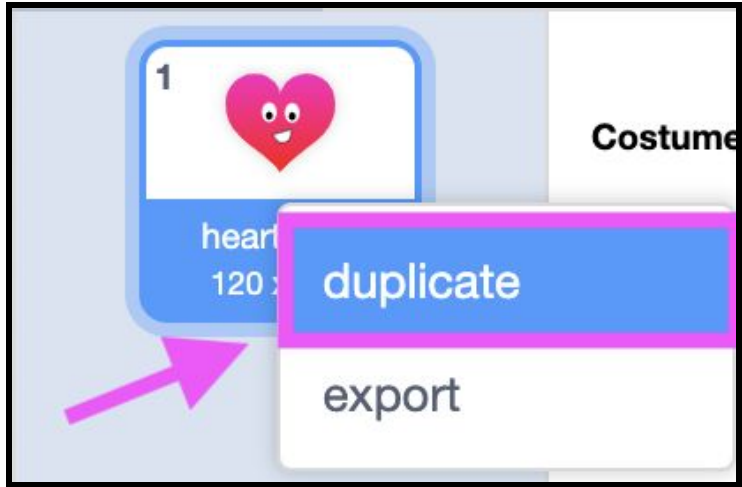
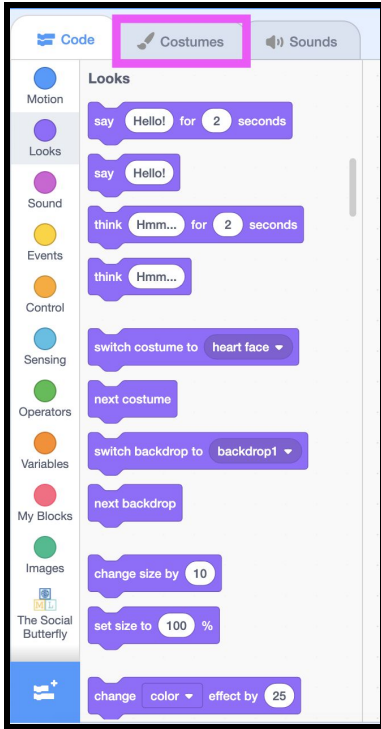
**Recommended to let learners write this code themselves and take it up after 5 minutes

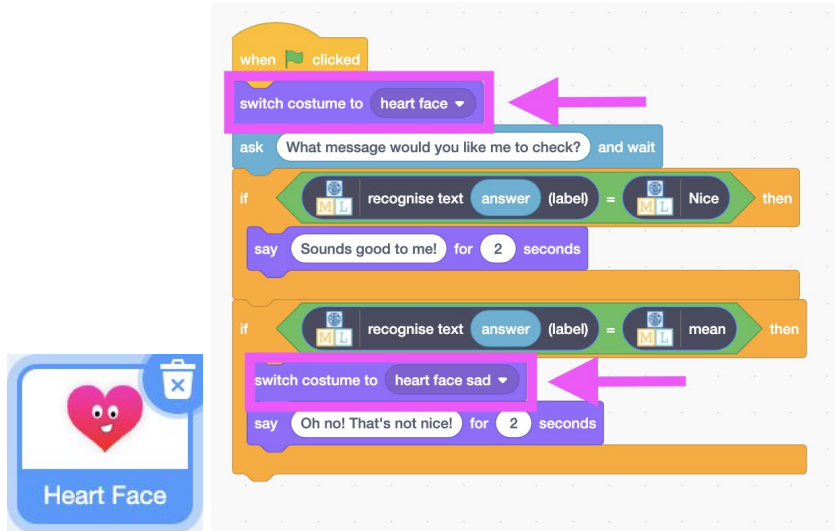
1. From the 'Control' category, add an 'if then block'
2. Ask: What should our assistant say if a phrase is mean? (A: Something helpful for the user to re compose their message)
3. From the 'Looks' category, add 'say' block inside of the 'if then block'. Type something positive such as "Oh no that's not nice! Remember to THINK!"
4. Ask: Look in the Machine Learning menu. How can we check if a phrase is mean? (A: Recognise text block and mean block)
5. Drag out the recognize text block and mean block
6. Add answer inside the 'recognize text label'
7. From the 'Operators' category, drag out an equal comparison block
8. Add 'recognize text' and 'mean' blocks in the equal comparison block
9. Place the equal comparison block in the "if then block"



STEP 7: Turn that Frown Upside Down

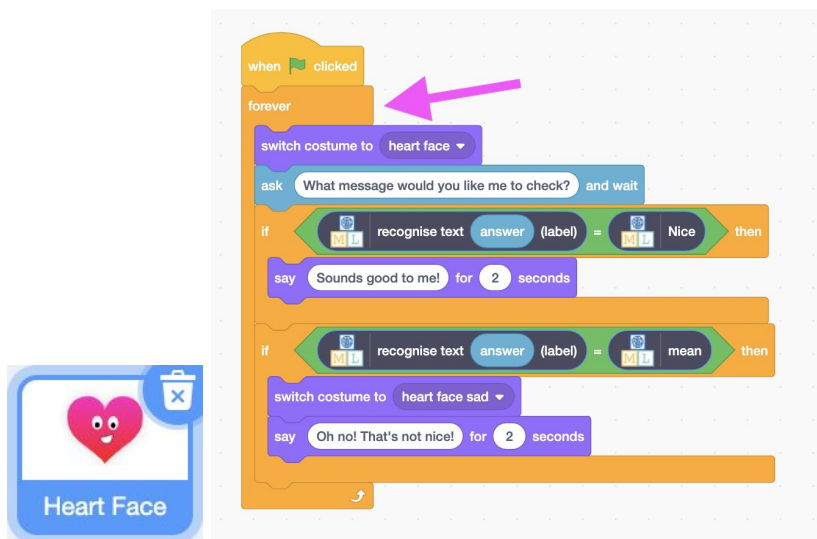
1. Click on the the costumes tab
2. Right click the smiling costume and select duplicate
3. Rename the costume
4. Click on the smile and rotate it around
5. Go back to the 'code' menu
6. Ask: When do we want our character to be happy? When do we want it to frown? (A: Frown when a mean phrase is said and happy the rest of the time)
7. Add a 'switch costume' to sad in the mean if statement
8. Add a 'switch costume' to happy underneath the when green flag is pressed





STEP 7: Making the Assistant Always Checking

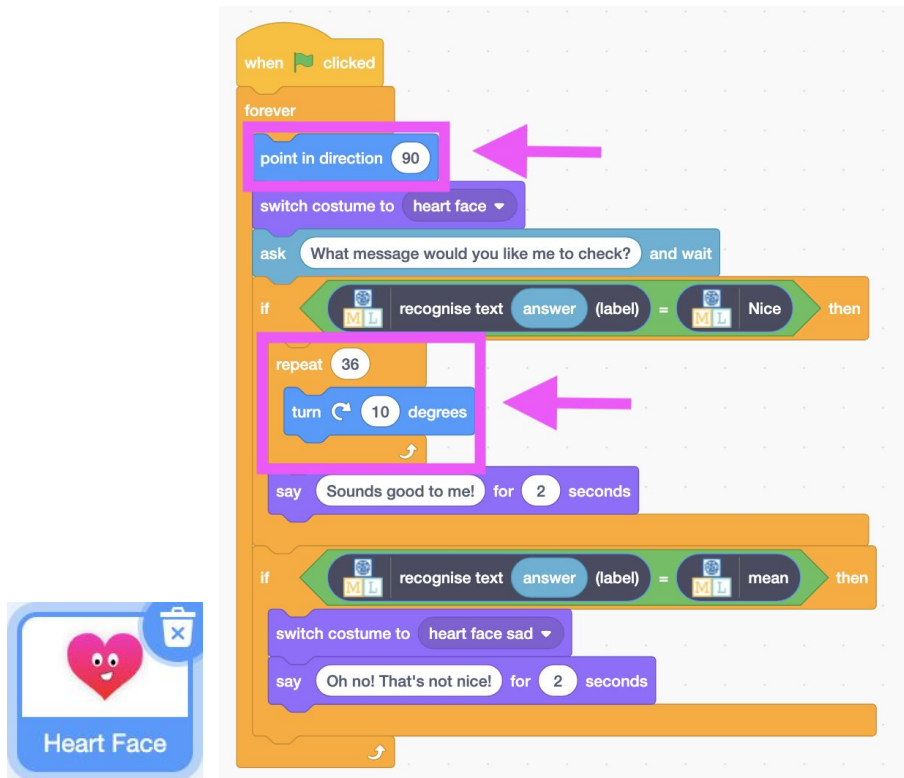
1. Ask: How can we make our assistant check messages forever? (A: Forever block)
2. Add a forever block around all the blocks to make them repeat



ADD ON: Happy Spin

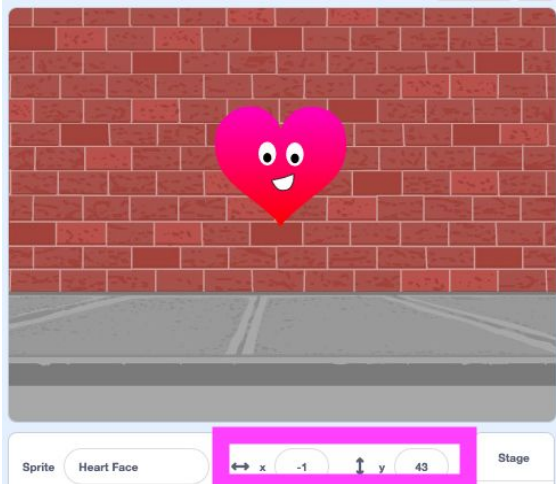
1. Give this as a stretch goal to learners if time permits. Have them try on their own to add this functionality and take it up afterwards
2. Add a 'repeat' block inside the nice if statement. Set it to repeat 36 times.
3. Add a turn 10 degrees block inside the repeat

4. Add a point in direction 90 at the beginning of the program inside the forever block



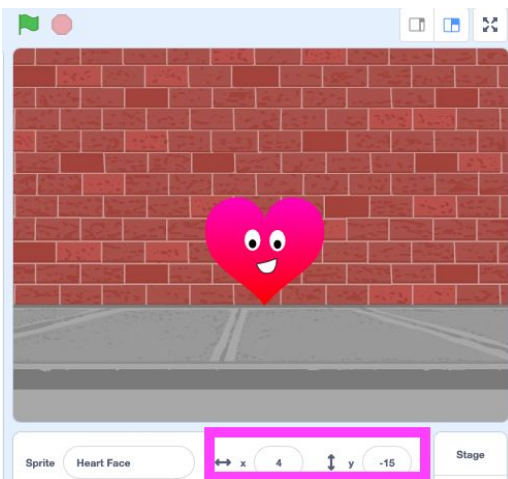
ADD ON: BOUNCE

1. Give this as a stretch goal to learners if time permits. Have them try on their own to add this functionality and take it up afterwards
2. Drag the sprite to however high you would like it to bounce
3. Drag the 'Glide to x,y' block and place it under the 'point in direction 90' block. The x and y position in the glide block should be the same as where your sprite currently is
4. Drag the sprite back to where you want it to end the bounce
5. Drag the 'Glide to x,y' block and place it under the block placed in step 3. The x and y position in the glide block should be the same as where your sprite currently is.
6. Set the glide blocks to 0.5 seconds to make them smoother
7. Duplicate the two glide blocks if you would like the sprite to bounce multiple times
8. Duplicate the glide blocks and place them under the 'switch costume to' block in the 'if then' block which recognizes mean phrases if you want the sprite to bounce when a mean phrase is said



```

when green flag clicked
  forever loop
    point in direction 90
    glide 0.5 secs to x: -1 y: 43
  
```



```

when green flag clicked
  forever loop
    point in direction 90
    glide 0.5 secs to x: -1 y: 43
    glide 0.5 secs to x: 4 y: -15
    glide 0.5 secs to x: -1 y: 43
    glide 0.5 secs to x: 4 y: -15
    switch costume to heart face
    ask What message would you like me to check? and wait
    if recognise text answer (label) = nice then
      repeat 36
        turn 10 degrees
      say Sounds good to me! for 2 seconds
    if recognise text answer (label) = mean then
      switch costume to heart face sad
      say Oh no! That's not nice! for 2 seconds
  
```



```
when clicked
  forever
    glide 0.5 secs to x: 1 y: 34
    glide 0.5 secs to x: 1 y: -14
    glide 0.5 secs to x: 1 y: 34
    glide 0.5 secs to x: 1 y: -14
    point in direction 90
    switch costume to heart face
    ask What message would you like me to check? and wait
    if recognise text answer (label) = nice then
      repeat 36
        turn 10 degrees
      say Sounds good to me! for 2 seconds
    if recognise text answer (label) = mean then
      glide 0.5 secs to x: 1 y: 34
      glide 0.5 secs to x: 1 y: -14
      glide 0.5 secs to x: 1 y: 34
      glide 0.5 secs to x: 1 y: -14
      switch costume to heart face sad
      say Oh no! That's not nice! for 2 seconds
```

