

I know my cats! Introduction to machine learning

Participants Perspective

Today I learned how machine learning networks are trained and operate **by** developing a neural network to identify cats.

Learning Goal

- Understand the basic architecture of a neural network
- Experience a machine learning process used to teach a neural network
- Understand how simple networks can become deep learning networks by adding layers of computational complexity
- Consider what it means to “Know” something

Logistics

- Suggested length: 1.5 hours
- Age range: Grade 8 or older

Part Number	Section Title	Group Size	Materials
Opening Hook	What is this?	Whole group	<ul style="list-style-type: none"> • Image of a cat projected • Whiteboard or chart paper with markers
Part 1	What’s a cat?	Small groups	<ul style="list-style-type: none"> • Digitally connected devices for participants • Paper & writing tool
Part 2	I know my cats!	Whole group	<ul style="list-style-type: none"> • Digitally connected device • Paper & writing tool • Cat images
Part 3	Deep into Cats	Whole group	<ul style="list-style-type: none"> • Paper & writing tool • Cat images
Part 4	Closing discussion	Whole group	

Safety Considerations

- **Laptops/Computers:** Electrocutation hazard from frayed or damaged cords and tripping over unsecured cords. These issues can be mitigated by securing all electrical cords, particularly extension cords, reaching over the floor (using tape) and immediately unplug them after use. Cords are plugged in/unplugged by adults, or by participants under adult supervision.

- **Internet:** Participants may come across materials that is not appropriate (including but not limited to violent or sexual images, racist/sexist commentary, and so on). Participants may use websites unrelated to the activity (social media, email, etc). When participants are using computers, activities should be monitored. When possible, each computer should be pre-set with the desired website already loaded. Review expectations with participants before you start the activity and develop a plan that they should follow if they come across inappropriate content.

Framework Connections

Toolsets	Skillsets	Mindsets
<p>Knowledge Consider what it means to “Know something”</p> <p>Basic structure and training of a neural network to achieve machine learning</p> <p>Introduction to deep learning and how increased complexity in a network can improve results.</p>	<p>Digital Skills Data analysis/analytics - understanding how data may be processed by a computer</p> <p>Machine Learning- Fundamental processes of machine learning</p>	<p>Digital Intelligence Awareness of the uses for machine learning, evaluating both positive and negative aspects</p>
<p>Resources Network training and test sets (image of cats) - Google images</p> <p>Other participants as nodes in your neural network</p>	<p>STEM Skills Computational thinking - systematic identification of features</p> <p>Problem Solving - Designing computer algorithms to solve complex problems</p>	<p>Computational Thinking Machine learning- chunking information into nodes for analysis. Using training data to teach neural network.</p> <p>Logical and systematic - evaluation of information</p>
<p>Experiences Work collaboratively to develop a neural network that can recognize cats.</p> <p>Computational thinking - Using logic, identification and classification to identify features of an image.</p>	<p>Essential Employability and Life Skills Collaboration - develop a network of nodes as a group</p> <p>Problem Solving/Efficiency - understanding that problems can be solved by building tools to solve them for you</p> <p>Creativity - identify unique characteristics that may be used in the categorization of images.</p>	<p>Digital Action Machine learning as an analysis tool to solve complex and time consuming problems.</p> <p>Data Science - analysis of (big) data to find valuable trends and categorize information.</p>

Nuts and Bolts

Hook: What is this? (10-15mins)

- Show the participants an image (preferably digital) of a cat. Ask the participants to identify what is in the image. After identifying the cat, ask the participants **how they know it is a cat?** Expand the questioning to cover the concept that the image of the cat is comprised of many things:
 - What aspects of the image show that it is a cat?
 - **How do YOU know it is a cat? Why do you know what a cat is?** Have you seen one before?
 - Would you be able to identify it as a cat if you had never seen a cat before?
 - Advanced - (if digital) Is this cat not a representation of 1s and 0s? **It is a digital image - Are there a series of 1s and 0s that represent what a cat should look like?**
 - Does your experiences identifying cats in the past help you to categorize this image as a cat now?

Participants should be able to identify that an **accumulated history of seeing cats** (in various forms - images, video, real, stories, etc.) **helps them to identify features of a cat and categorize them. The more experience with cats, the better they are at identifying cats.**

- Ask the participants to identify **What specific features define what a cat looks like?** Record the features that are identified as these will be used in the next step. Some common suggestions are:
 - Ears
 - Eyes
 - Nose
 - Tail
 - Fur
 - Whiskers
 - Activities
 - Environment
- Features should not be limited to this list, allow the participants to make a complete list.

Part 1 : What is a cat? Teaching the machine (20 mins)

This type of problem is a very common problem that computer scientists and data scientists want to solve. It is extremely **helpful if we can train a computer to identify, classify or cluster objects into categories. The process of setting up a computer to do this is called machine learning. Machine learning involves developing a network to analyze objects based on some parameters and letting it learn.** In many cases, the **goal is to have the machine learn how to categorize things itself (called Classification of Data),** based on some examples of objects that have been pre-categorized (called *Supervised learning process*). In simple terms - setup a network and show it what cats look like and what cats don't look like. After a while, it will be able to tell you if the object you show it is a cat.

- Explain to the participants that they will be setting up simple human powered machine learning network. **This network will be a simplified neural network (a network made up of nodes, kind of like the human brain).**
- Group the participants into small groups of 2-3 people. Allow the groups to take turns selecting an important feature (that they listed above) until each group has at least 1 feature and most of the features have been selected. Ask them to consider as they are selecting

what they think is the most important features are for identification. There can be no duplicate features at this time. The network will need at least 6 NODES to run well, but more will improve how well it works.

- The feature (or features) that the groups have selected will be the NODE or NODES that they are responsible for during this activity. Hand out 1 sheet of blank paper for each feature. This paper represents everything that the NODE knows about cats. Have the groups write the feature in the middle of the paper, with a circle around it. Plenty of space should be left around the label so that information can be recorded on the paper.
- It is now time for the participants to complete a [supervised learning process to train the NODE to recognize the feature that has been selected](#). Participants are taking the [role of computer scientists](#), feeding information to the NODE of what to look for in that feature and writing it down on the NODE sheet. Participants should use google image to search for the evidence to feed into the NODE. Don't go off memory - [find images that provide evidence of the information](#) you write down. For example, let's say we were training the "Nose" NODE:
 - Document what a cat nose looks like. Draw some examples, describe its shape, size & placement (relative to something else, like the cats eyes), what colour is it, etc. [Include things that are obvious! Remember that the NODE has never seen a cat before, so you need to teach it everything it needs to know.](#)
 - Document what it doesn't look like too. Build some parameters around what it can not be. It shouldn't have more than two nostrils, it shouldn't be much bigger than the cat's eye, it shouldn't be lime green, etc.
 - Give the participants 15-20 minutes to complete the supervised training of the NODE, doing image searches and writing down everything they can about the feature.
 - [Encourage collaboration between the groups - what strategies are each team using to train their NODE?](#)
 - Note that groups should take their time to write clearly, as others will be reading the information.

Part 2: I know my cats!

- Now that we have trained NODES, it is time to see if our neural network is working as expected. Have the groups trade their NODE sheets to other groups so that every group has a new NODE. Give the groups a couple of minutes to familiarize themselves with their new NODES.
- Show images of different things (cats and not cats) and see how well the NODES work. Participants should take on the [role of the NODE](#) and only use the information on the sheet to determine if the NODE finds the feature it is looking for. [NODES have no other context, so they base it only on what they have on their NODE papers \(what they have been trained for\).](#)
 - Example: If the picture is of a cat but its nose can't be seen, the nose NODE should determine that the image is not a cat.
- Show an image and have the NODES call out one at a time if the image is of a cat. Tally the counts of "Cat" and "Not a Cat" on the whiteboard or chart paper. Whichever vote has the majority determines if the image is a cat or not.

- After categorizing each image, ask the participants if a human would have categorized this image as a cat or not. Place a ✓ or ✗ depending on if the network made the correct prediction. This is the real test of if a neural network is working well - how well does it mimic human decision making?
- Continue to show additional images from the collection as desired.
- Discuss as a group the outcome of the test.
 - How well did your machine learn what a cat is? Was your neural network effectively trained?
 - What types of changes would you make to how you trained the NODES?
 - What types of changes would you make to the network itself - other NODES, etc?
 - Were there certain aspects that worked well or didn't work well?
 - As a Computer Scientist, what did you learn from this test?
 - What about other items using the term cat (CAT 5 cable, cat treats, etc), should these be identified as cats? Why or why not?

Part 3: Deep into Cats (15 mins)

Computer Scientists use the term “Deep learning” to describe when additional layers of complexity are added to a neural network. This can be done in a number of ways:

- Having layers of NODES, where NODES will take information from other NODES and uses that to help categorize. For example, the network could have the nose, eyes, mouth and ear NODES feed into a face NODE. This NODE could use information like where in the picture the parts were seen to better determine if a cat's face is present.
- Weighting NODES can make results from one NODE more important than another. If the legs NODE sees 2 legs, that might be less important data than seeing eyes, ears and a nose that look like a cat.
- Allowing NODES to cluster results. This means that the NODES select a whole bunch of images that have similar features and group them together. These images can then be analysed using a different network to look for specific things.
- Commonly, undefined NODES can be used where the network itself is allowed to define what the NODES are looking for. Instead of training the neural network to look for a cat's nose, you let the network learn to pick out its own important aspects. As the training images are processed, the NODES will be defined by the network.
- Ask the participants to consider how they could develop their neural network into a network that thinks deeper?
 - What changes might the group make to how the network operates?
 - Additional training to the NODES?
- Give the participants 10 minutes to improve their neural network.
- Once ready, select a few images from the previous test to re-test again. Include at least one image that was properly identified and a number of images that were not. Again keep a tally of the results on the whiteboard.

Part 4: Closing Discussion

- Discuss the final results from the test of the neural network:

- Was the network more effective at identifying cats?
- Were the improvements effective?
- What other changes might want to be made?
- In this activity, there has been a lot of human involvement with the training of the NODES. As described above, many machine learning processes actually take undefined NODES and train them to look for features themselves. This eliminates the time consuming process of teaching the NODES.
- Other discussion points: Allow the participants some time to personally reflect on the question before discussing as a group:
 - Consider how Google images knows to look for cats when you type 'cats'. Do you think that Google uses machine learning to categorize information?
 - How is the machine learning process similar to human learning? How are human brains setup like a neural network? Are these networks trained in the same way?
 - How might machine learning benefit your life? Could machine learning be harmful?

Modifications and Extensions

Ways to make this activity more approachable:

- Create a few sample NODE pages for the exercise.
- Brainstorm information to include on the NODE pages as a whole group.
- Where internet access is limited - Print a few pages of google image searches for "cats"
- Prompt groups with questions to improve their NODE construction:
 - What if the cat is facing the away from the camera?
 - What if the cat is a cartoon?
 - What is a dog is in the picture instead?

Ways to make this activity more challenging:

- Create 3 layers in the deep learning network. How might a 3rd layer be beneficial?
- What other information might a NODE want to output, besides cat/no cat? How can this be used to improve the deep learning of the network?
- Consider how a computer might process these images - The network does not "See" the image, rather it processes the image pixel by pixel examining the colour of the pixel relative to the pixels around it. What are some examples of information the NODES might be looking for in this case?

Assessment and Evaluation

- Whole group evaluation of the neural network based on ✓ or ✗.
- Seeing improvements in performance when adding deep learning strategies to the process.
- Discussion of the results and considering how these networks may be similar to the human brain. Personal reflection.

Credits, Kudos, Shoutouts

- Background information on Machine learning in neural networks
 - <https://skymind.ai/wiki/neural-network>

- https://en.wikipedia.org/wiki/Machine_learning
- https://en.wikipedia.org/wiki/Deep_learning

Terms of Use

Prior to using this activity or parts thereof, you agree and understand that:

- It is your responsibility to review all aspects of this activity and ensure safety measures are in place for the protection of all involved parties.
- Any safety precautions contained in the “Safety Considerations” section of this write-up are not intended as a complete list or to replace your own safety review process.
- Actua shall not be responsible or liable for any damage that may occur due to your use of this content.
- This work is licensed under a Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License. For more information, please see <https://creativecommons.org/licenses/by-nc-sa/4.0/>.

About Actua

Actua is Canada’s leading science, technology, engineering and mathematics (STEM) youth outreach network representing a growing network of university and college based members. Each year 250,000 young Canadians in over 500 communities nationwide are inspired through hands-on educational workshops, camps and community outreach initiatives. Actua focuses on the engagement of underrepresented youth through specialized programs for Indigenous youth, girls and young women, at-risk youth and youth living in Northern and remote communities. For more information, please visit us online at www.actua.ca and on social media: [Twitter](#), [Facebook](#), [Instagram](#) and [YouTube](#)!

Cat Images Slide:

<https://docs.google.com/presentation/d/1zfUBNrgwchQa4vTwEX6GcqpbqgcD3hsczUzR6OEzsIE/edit#slide=id.p>