Actua's AI Activities Series

Activity 9

Sentiment Analysis: Understanding the Emotion Behind Text

Sentiment Analysis: Understanding the Emotion Behind Text

If you're accessing this activity directly, did you know there are eight other activities in this series up on our website? If you find yourself unfamiliar with any of the AI concepts and terminology introduced in these activities, please refer to our AI Glossary. These activities also follow a space exploration narrative when done in order. It is recommended to complete the activities in order but they can also be done on their own.

You and your group mates are astronauts and scientists aboard the Actua Orbital Station. Unfortunately, your station just got bombarded by magnetic rays and your electronics have begun to shut down! The only one who can save you is the station's AI, DANN. DANN stands for Dedicated Actua Neural Network, and it's gone a little loopy. Brush up on your technical skills, learn about AI, and save yourself and your crewmates!

We've established DANN's Code of Ethics in "Ethics in AI: Don't Let DANN Turn Evil" Mission Control has been working diligently to get DANN's Language Processing Engine working. Now that it is, Mission Control has tasked you with restoring some of DANN's ability to understand the content of the language that's being used. DANN, needs to be able to identify and classify human emotions, so it knows when there's an emergency situation. Help train DANN to re-learn how to detect emotions in words. Once that's done, DANN will be fully back online! You're about to save the whole space station!

Activity Summary

In this activity, participants will learn about natural language understanding (NLU), a sub-area of natural language processing that focuses on how AI can comprehend the meaning of natural human language. Participants will learn about sentiment analysis, and how it can be used to detect emotion in text. They will then train and test their own sentiment analysis AI model.

Developed by Actua, 2022.

Delivery Environment	Activity Duration	Intended Audience
Classroom with computer access	60 minutes	Grades 9-12 (Ages 13-18)

Achievement Goals

Learning Goals

Learning goals are statements referring to the understanding, knowledge, skills or application participants acquire during the activity. **Following this activity, participants will:**

- **Define and describe** sentiment analysis in the context of natural language understanding.
- **Create** training data for, and train, a basic sentiment analysis AI model.

Logistics (Timing, Group Size, Materials)

Section Title	Time	Group Size	Materials
Opening Hook: Understanding language	10 min	*Entire group*	• Computer (with audio) • Internet access (to access video): The Turing test: Can...
Activity 1: What is "sentiment analysis"?	15 min	*Small groups or entire group*	Per small group • Pen/pencil • Paper
Activity 2: Training sentiment analysis models	30 min	*Individually or in small groups*	Per individual or small group • Computer • Internet access (to access Machine Learning For Kids): https://machinelearningforkids.co.uk/#!/login

Safety Considerations

Safety considerations have been provided below to support safety during this activity, however they are not necessarily comprehensive. It is important that you review the activity and your delivery environment to determine any additional safety considerations that you should be implementing for the delivery of these activities.

Online Safety

Some components of this activity require the use of devices connected to the internet.

- Facilitators should review the provided videos and read/explore provided websites and materials to determine if they are suitable for their participants.
- Where applicable, facilitators should remind participants to stay on task and only use links provided within this activity.

Activity Procedure

Opening Hook

The exact definition of the word "intelligence" is an important discussion in the field of artificial intelligence. Connections are often made between machine intelligence and human intelligence. **That is to say, a computer is often considered intelligent if it can do the same things a human does.**

Do you think that that would be a good definition of "intelligence" for computers and AI? What would some of the strengths and weaknesses of such a definition be? Think about these questions as you watch this short video, titled "The Turing test: Can a computer pass for human?" (running time: 4:42).

In a previous activity, you may have been introduced to "natural language processing" (NLP), a field that combines linguistics and computer science. **This field is concerned with having computers (and AI) make sense of, and sometimes respond to, everyday human language.** A common NLP application is a "chatbot": an AI that was trained to provide predetermined responses to a wide variety of input phrases.

As a class, discuss the following questions:

1. **What do you think intelligence is (for a computer or an AI)?** This question doesn't have a singular response, however responses could include:
    a. Sensing/perceiving and reacting to what is sensed/perceived
    b. Learning from own experience or the experiences of others
    c. Analysing problems or issues and suggesting rational/appropriate actions.
    d. Responses could also connect back to "human intelligence" and things that humans can do, for example:
        i. Compose music
        ii. Paint paintings
        iii. Write novels
2. Do you think the **capacity to carry on a conversation is an indicator of "intelligence"?**
    a. According to the video, not necessarily, since there are chatbots that can produce convincingly human responses without actually being truly intelligent.
    b. Since this is an opinion question, most answers that are supported with some amount of reasoning would be valid.
3. Does an AI need to **understand the meaning of language to be able to respond?**
    a. Not exactly. The chatbot worked by comparing input phrases to the phrases used in its training. It didn't try to extract meaning or ideas from an input sentence.
    b. As mentioned in the video, chatbots can create the illusion of understanding what they're being told by crafting responses using their input, without necessarily understanding the content of the input.
4. **What does it mean to understand a piece of text? How might that understanding be shown?**
    a. This is a philosophical question, so a variety of well-reasoned responses would be acceptable. This is expanded upon in the paragraph that follows.

In this activity, you will be learning about “natural language understanding” or NLU. NLU is a sub-area of NLP that is specifically concerned with the ability of an AI to, among other things, **comprehend the content of natural human language.** This allows the AI to extract and identify key ideas, embedded meanings, emotions, and/or intents from the text.

Section 1: What is “sentiment analysis”?

A key task in NLU is “sentiment analysis”. Before continuing, **{ in small groups / as a large group} consider the following questions:**

1. What is a sentiment (in the context of NLU)?
    a. “Sentiment”, in the context of NLU, refers to information like polarity (whether a statement is generally positive, negative, or neutral) and emotion (happy or angry). This information is embedded in the text of the author, and is very easy for humans to identify.
2. What do you think sentiment analysis is?
    a. “Sentiment analysis” is the process of analyzing language to understand the sentiments within it. This means determining, based on the words being used, whether a text is happy, or angry, or something else.
    b. Outside of NLU and analyzing text, there are also applications of sentiment analysis that try to guess an emotion based on facial expressions or tone of voice. We will not be addressing these applications in this activity.

Sentiment analysis is a difficult task because words can have multiple meanings, depending on their usage and context. A variety of sentence structures can be used to mean the same thing, or very similar sentences can mean very different things. It’s much more difficult for computers to process them. **Can you identify the subtext in the words below?**

For each of the following sentences, **{in small groups / as a class }**, **try to determine its polarity,** i.e. if it is positive, negative, or neutral. If there's a disagreement, note the arguments in support of each position and vote on the final result:

1. What an incredible experience!
2. She was an inspiration to us all.
3. I was told that it would be interesting to me.
4. You said that it would be fun, and it was.

Then, after going over each statement, **consider the following questions:**

1. Did everyone agree on how to classify each statement?
    a. If not, what were the disagreements?
    b. If so, do you think that there could be other ways to interpret any of the statements?
2. What clues (e.g., specific words or punctuation) did you use to determine if a statement was positive, negative, or neutral?

Now, for each of the following statements, **try to determine the possible emotion(s) being communicated:**

1. This is amazing!
2. I'm really upset...
3. I'm so angry at you!
4. It's my birthday today!
5. How could you do this to me?
6. How's it going?
7. I wish you could be here.
8. I don't know how to feel about this.

Likewise, after going over each statement, **consider the following questions:**

1. How many different emotions did you identify? Are any of the emotions that were identified similar to other ones?
2. Did everyone agree on how to classify each statement?
    a. If not, what were the disagreements?
    b. If so, do you think that there could be other ways to interpret any of the statements?
3. What clues (e.g. specific words or punctuation) did you use to determine the emotions in a statement?
4. The last few sentences were a bit trickier. Why were the last few sentences harder to define than the first ones?
5. How would a computer be able to tell the difference between a happy exclamation point and an angry one?

Finally, **reflect on the connections of sentiment analysis to artificial intelligence:**

1. How does a computer tell whether something is happy or sad (or neither)?
    a. An AI model can be trained to recognize “happy” or “sad” by analysing a large dataset that has been correctly labelled with those emotions. The training data would contain lots of words and sentences labelled as “happy”, and the AI would learn from those. The AI could then compare any new sentences against every other “happy” sentence it was given. If the new sentence seems similar to enough of those, then the AI declares it as happy!
2. Why is sentiment analysis important for the development of AI?
    a. Understanding emotions is an important step in developing AI programs that can think and act like humans. Sentiment analysis is one of the first pieces of that. It might be easy for us to tell whether a sentence is happy or sad, but it’s much more difficult for computers.

Section 2: Training sentiment analysis models

Now that you have an idea of what sentiment analysis is, you will be training a sentiment analysis model that can classify text as “happy” or “sad”. **{ Individually / in small groups }:**

1. Navigate to the Machine Learning for Kids website: https://machinelearningforkids.co.uk/#!/login
2. Click on the “Try it Now” button, and follow the instructions on screen to create your first project. The project can be given any (appropriate) name.
3. Select “text” for the “Recognising” field. Once that is done, you can click on your project and will be taken to this screen:

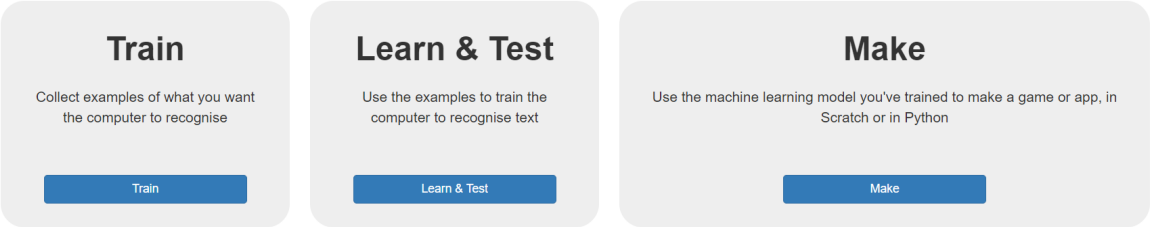

4. Click on “Train”, and follow the instructions to add a new label. This label will be the “Happy” label.
5. Repeat the same steps and create another label, with the title “Sad”.
6. Spend the next 5-10 minutes filling in both labels with any words, sentences, or punctuation that you think matches either “happy” or “sad”. These could be words and/or sentences from previous class discussions, or you can come up with your own.
    a. The program only asks for a minimum of ten examples for each label, but you should try to come up with more. The more data you have, the better your sentiment analysis model should work.
    b. **Make sure you have an equal amount of example data in each label, or the AI will automatically prefer the label with more data.**
7. When you’re done creating your data, click the “Back to Project” button and move to the “Learn and Test” page.
8. Click the “Train new machine learning model” button to start model training.

Model training may take a while depending on the amount of data provided. **While your model trains, compare your training data { with nearby groups / as a whole class }:**

1. What sort of text did you put in for the happy label? For sad?
2. Did everyone use the same examples?
3. Was there any input data that someone labeled that you would label differently?

If you're still waiting for your model to finish training after you're done sharing, there is also a quiz at the bottom of the page that you can do about AI and machine learning.

Once the training is complete, it's time to test your model! Put in sentences to see what emotion it gets labelled as. Make sure that you test with sentences that are different from the ones that you used as examples.

Reflection & Debrief

Having trained and tested your own sentiment analysis model, **{ reflect on / discuss as a class } the following questions:**

1. **Does everyone always agree on the polarity** (i.e. the general expression of a positive, negative, or neutral sentiment) or emotion(s) of a piece of text? What can affect each person's analysis of a statement?
    a. Many statements would probably have a general agreement on polarity and/or emotion, but there might also be sentences that can be interpreted in ways that people might understand differently.
    b. Experience, assumed context, or even a person's mood at a given moment might affect how a person analyses a sentiment
2. **Are statements always clear and defined in their meaning?** Can you think of any kinds of statements that might be difficult for an AI to understand?
    a. Many statements could be ambiguous and require you to make assumptions based on intuition for them to make sense. The video at the beginning of the activity uses the example, "I took the juice out of the fridge and gave it to him, but forgot to check the date."
    b. Some language concepts, such as sarcasm, qualified statements (e.g. "The food was surprisingly good, given the challenges."), and double negatives (e.g. "I don't dislike chocolate cake"), could make it hard to figure out the true meaning of a statement.
3. **How do you think AI sentiment analysis models can be used in a positive way?**
    a. A sentiment analysis model could potentially detect distress in a person's communication and suggest that they get help (or automatically send help).

Delivery Adaptations

How might you adapt the time, space, materials, group sizes, or instructions to make this activity more approachable or more challenging? **Modifications** are ways to make the activity more accessible, **extensions** are ways to make the activity last longer or more challenging.

Modifications

- If a participant is struggling to create data for the activity, have them read from a book or other writing source, and classify each sentence as “happy”, “sad”, etc. This provides them with a base to come up with their own examples as well.
- If you find that groups are having difficulty generating their training data, consider creating a Google Doc or similar collaborative document and having each group add their training data for each label. Groups can then choose to train a model using this data while still testing their models independently.

Extensions

- You can add more labels beyond “happy” and “sad” to your model. Consider adding labels such as “angry”, or “confused”, creating training data, and re-training your model.
- In a Google Doc or similar collaborative document, brainstorm testing data for your models. Remember that the testing data shouldn’t be used for training and that your training data shouldn’t be used for testing. Have all groups use the same testing data to evaluate their models.
- In a Google Doc or similar collaborative document, have each group add their training data for each label. Train a model using this data and compare its performance to the original models trained.
- The last question of the activity asks about the potential positives of sentiment analysis. As an extension, you can facilitate a discussion on the potential implications of sentiment analysis by watching this video, titled *Can Machines Read your Emotions?* (running time: 4:20,

https://www.youtube-nocookie.com/embed/QFk3e5PcK7s). After watching the video, discuss the positive and negative uses/implications that were identified. Do the positives outweigh the negatives?

References & Gratitude

TED-Ed (2016, April 26). The Turing test: Can a computer pass for a human? - Alex Gendler [Video file]. https://www.youtube-nocookie.com/embed/3wLqsRLvV-c.

TED-Ed (2016, November 29). Can machines read your emotions? - Kostas Karpouzis [Video file]. https://www.youtube-nocookie.com/embed/QFk3e5PcK7s.

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/.
- You may adapt the content for your program (remix, transform, and build upon the material), providing appropriate credit to Actua and indicating if changes were made. No sharing of content with third parties without written permission from Actua.

About Actua

Actua is Canada's leading science, technology, engineering and mathematics (STEM) youth outreach network, representing a growing network of over 40 universities and colleges across the country. Each year 350,000 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!

Appendices

Appendix A: Career & Mentor Connections

- Machine learning researcher
- Programmer
- Software Engineer
- Psychologist