CAIS++
CENTER FOR AI IN SOCIETY'S STUDENT BRANCH

Improving the Accuracy of Facial Expression Recognition Using Computer Vision

Rida Faraz, Arjun Bedi, Jonathin Aydin

BACKGROUND:

- **Facial emotion recognition** (FER) is the process of using a medium such as photos to recognize human emotion
  - It comprises of three steps: **face detection**, **facial expression detection**, & **emotion classification**
- FER enables us to study the **nonverbal communication and behaviors** of people, which is crucial in many facets of life
- **Education**: FER allows teachers and mentors to determine whether or not students are understanding concepts
- **Mental Health**: Expressions are an indicator to an individual's true emotions and wellbeing, communicates emotional state without requiring the person to explicitly express their thoughts
- **Marketing**: Allows vendors/companies to determine consumer satisfaction with their product
- Overall, FER provides a window into people's emotional states without the need for verbal communication

OBJECTIVE:

- Facial emotion recognition is important in many different facets of life, from therapy to education
- Our aim is to build a CNN to **accurately identify** and read facial expressions from an image and **determine the emotion** the face is conveying
- Ensuring that the CNN has high accuracy would make FER a **more reliable** process and therefore, make it more **effective** in its various applications

METHOD:

- To create our classification model, we decided to use a standard CNN utilizing the Keras sequential function
- Other features:
  - Optimizer: Adam
  - Learning Rate: 0.0001
  - Epochs: 60

DATA:

FER-2013 Data Set

- 28,709 training samples
- 3,589 test samples
- Seven Classes: Anger, Disgust, Fear, Happiness, Neutral, Sadness, & Surprise
- Contains 48x48 pixel greyscale images of faces with various expressions
- Images were already standardized (similar centering and placement), leaving us with the task of applying the convolutions to the images
- **Potential Problem**: Variation in size of each class

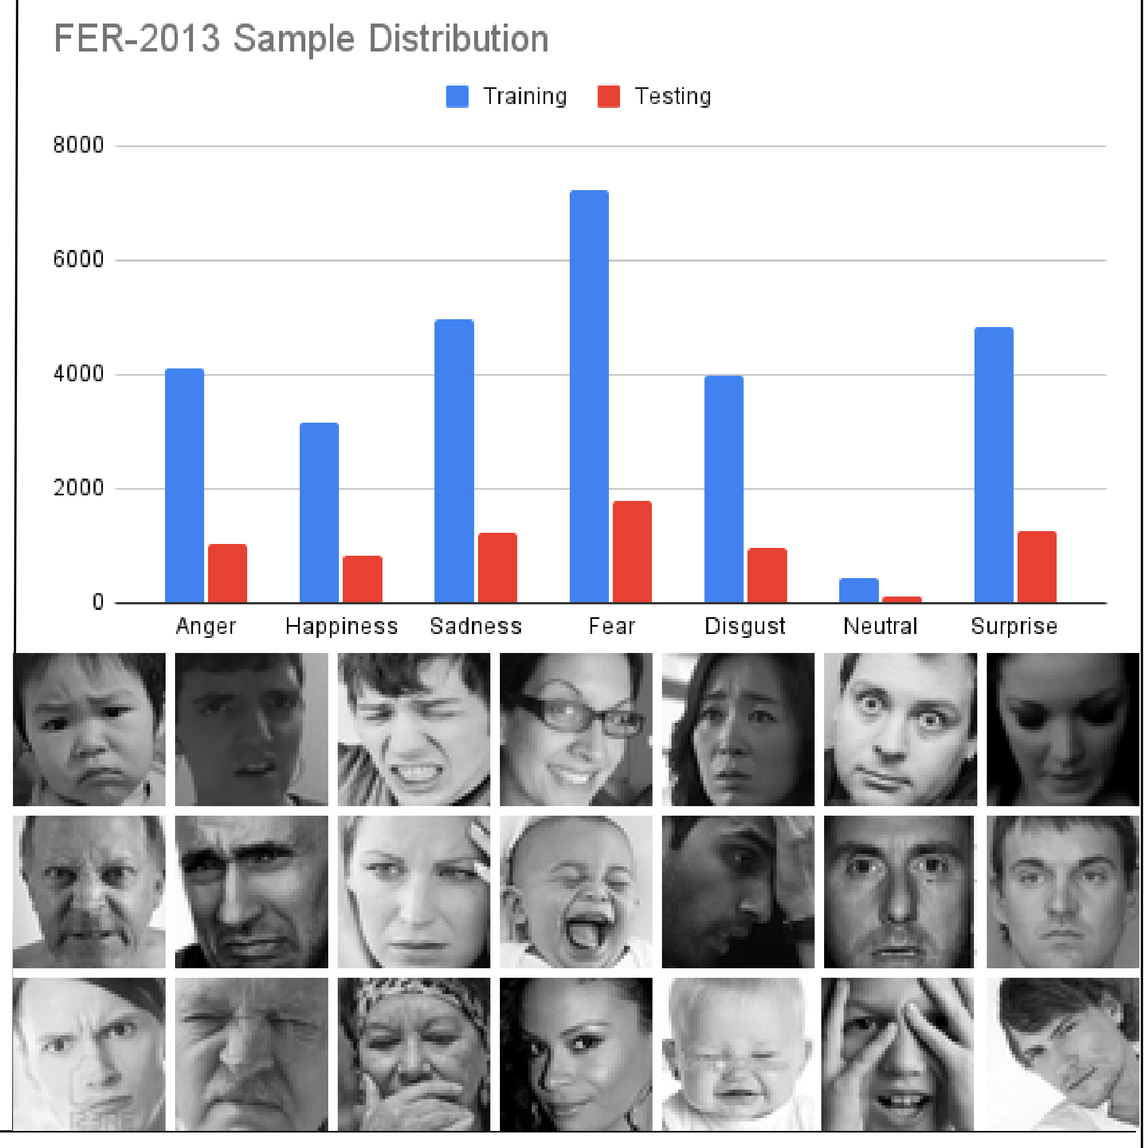

RESULTS:

Train Accuracy: 91.03
Test Accuracy: 66.58

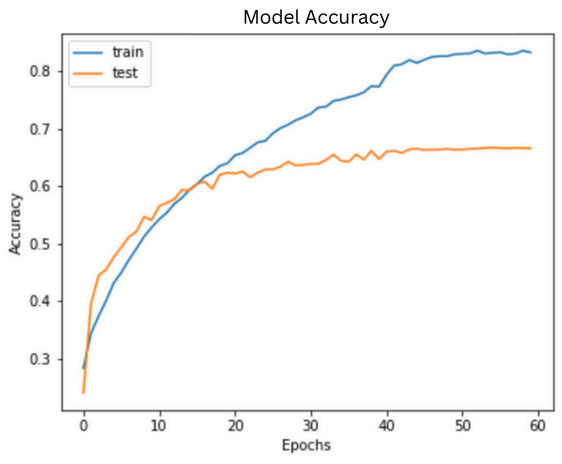

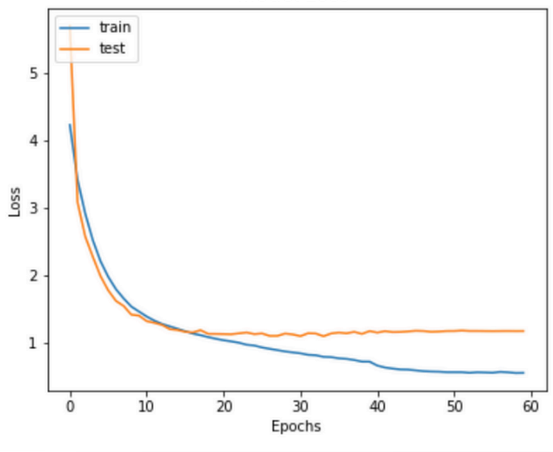

CONCLUSION:

- Overall, we were able to create a **successful model** that achieved a high training accuracy score, but it **could be improved** when it comes to the testing accuracy score
- Looking back, there were several potential sources of error
  - There was **high variation** in the sample size for each emotion class, resulting in varying results for different classes
    - There was a **statistically significant difference** between the number of happiness samples and disgust samples classified correctly
  - The model overfit the training data, which is the reason for such a high difference between training and testing accuracy

NEXT STEPS:

- Going forward, we would like to try training this FER model on other data sets
- While there was some inclusivity in the data, the FER-2013 dataset **lacked diversity**, possibly resulting in biased results
  - **Would the results be the same on minority samples?**
- When it comes to emotions and expressions, some facial features are more important in some emotions than others, such as eyes indicating happiness or nose indicating disgust
  - **What distinguishes each emotion from one another and how does that affects the class's associated accuracy?**