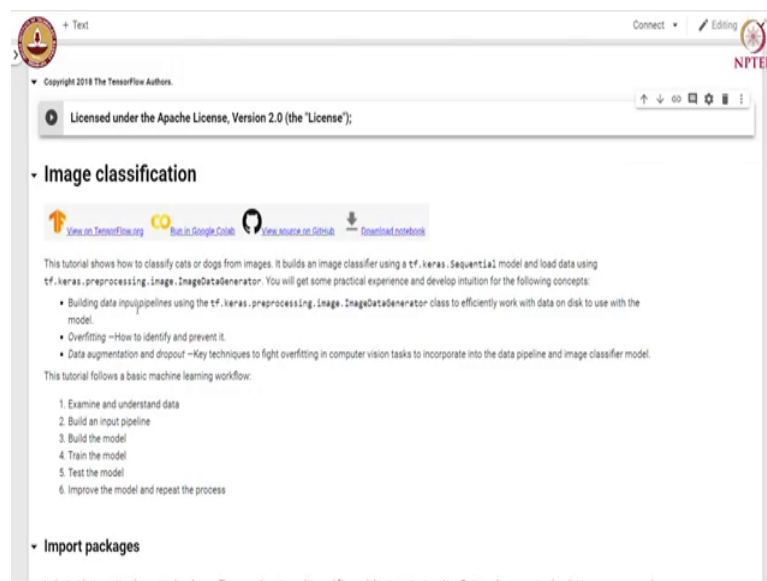


**Practical Machine Learning**  
**Dr. Ashish Tendulkar**  
**Department of Computer Science and Engineering**  
**Indian Institute of Technology, Madras**

**Lecture - 24**  
**Image Classification and Visualization**

(Refer Slide Time: 00:13)



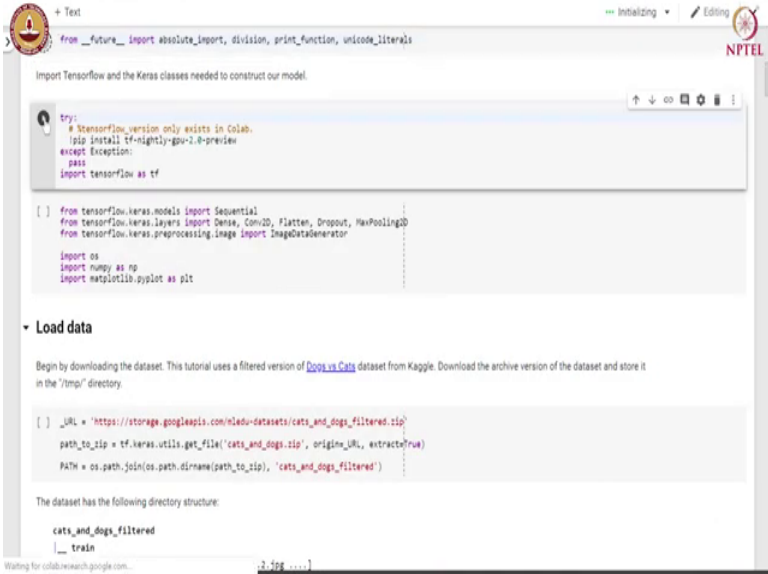
In the previous session, we studied CNNs we also learnt how to build CNN models with transfer learning. In this session, we will build Image Classification models from scratch and we will use bunch of strategies that are employed in practice while building image classification model and we will also visualize what the CNN is learning by looking at the activations after each layer.

So, we will follow a basic machine learning work flow where we will examine and understand the data. We will build the input pipeline to bring the data to the training. We will build the model, they train it, test it and then will improve the performance of the model and repeat the process.

We will get some practical experience and develop intuitions for building input pipelines for images using image data generator class. We will also study how to identify over

fitting and prevent it and we will also learn key concepts like data augmentation and dropout.

(Refer Slide Time: 01:50)



```
from __future__ import absolute_import, division, print_function, unicode_literals

import tensorflow and the Keras classes needed to construct our model.

try:
    # tensorflow_version only exists in Colab.
    !pip install tf-nightly-gpu-2.0-preview
except Exception:
    pass
import tensorflow as tf

[ ] from tensorflow.keras.models import Sequential
    from tensorflow.keras.layers import Dense, Conv2D, Flatten, Dropout, MaxPooling2D
    from tensorflow.keras.preprocessing.image import ImageDataGenerator

import os
import numpy as np
import matplotlib.pyplot as plt
```

**Load data**

Begin by downloading the dataset. This tutorial uses a filtered version of [Dogs vs Cats](#) dataset from Kaggle. Download the archive version of the dataset and store it in the "/tmp/" directory.

```
[ ] _url = 'https://storage.googleapis.com/mledu-datasets/cats_and_dogs_filtered.zip'
    path_to_zip = tf.keras.utils.get_file('cats_and_dogs.zip', origin=_url, extract=True)
    PATH = os.path.join(os.path.dirname(path_to_zip), 'cats_and_dogs_filtered')
```

The dataset has the following directory structure:

```
cats_and_dogs_filtered
├── train
```

Waiting for colab.research.google.com... [2.788 sec]

Let us install TensorFlow 2.0. Let us import image data generator and other libraries like dense convolution 2D flattened Dropout and MaxPooling 2D from Keras layers and we will also import sequential for building model. We use matplotlib.pyplot for plotting the performance of the model. We use dogs versus cats dataset from Kaggle competition.

(Refer Slide Time: 02:36)



```
_url = 'https://storage.googleapis.com/mledu-datasets/cats_and_dogs_filtered.zip'
path_to_zip = tf.keras.utils.get_file('cats_and_dogs.zip', origin=_url, extract=True)
PATH = os.path.join(os.path.dirname(path_to_zip), 'cats_and_dogs_filtered')
```

Downloading data from [https://storage.googleapis.com/mledu-datasets/cats\\_and\\_dogs\\_filtered.zip](https://storage.googleapis.com/mledu-datasets/cats_and_dogs_filtered.zip)  
68600000/686062136 [\*\*\*\*\*] - 1s 0us/step

The dataset has the following directory structure:

```
cats_and_dogs_filtered
├── train
│   ├── cats: [cat.0.jpg, cat.1.jpg, cat.2.jpg ...]
│   └── dogs: [dog.0.jpg, dog.1.jpg, dog.2.jpg ...]
└── validation
    ├── cats: [cat.2000.jpg, cat.2001.jpg, cat.2002.jpg ...]
    └── dogs: [dog.2000.jpg, dog.2001.jpg, dog.2002.jpg ...]
```

After extracting its contents, assign variables with the proper file path for the training and validation set.

```
[ ] train_dir = os.path.join(PATH, 'train')
    validation_dir = os.path.join(PATH, 'validation')

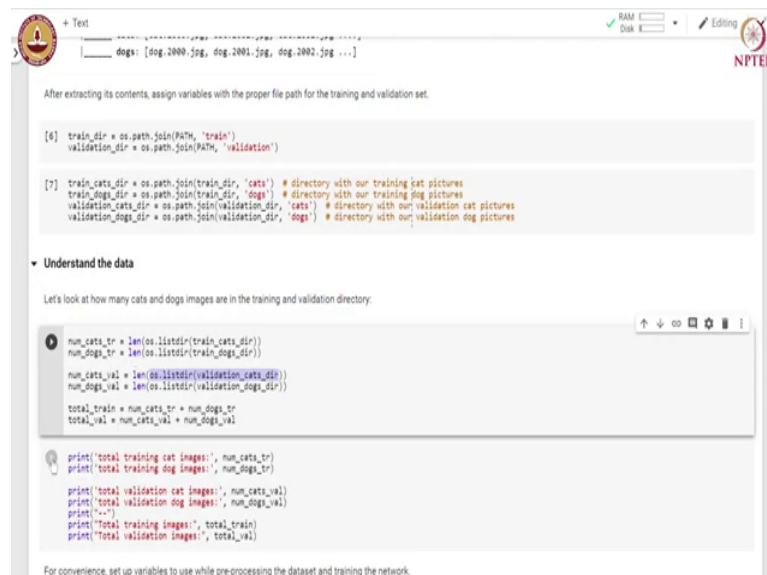
[ ] train_cats_dir = os.path.join(train_dir, 'cats') # directory with our training cat pictures
    train_dogs_dir = os.path.join(train_dir, 'dogs') # directory with our training dog pictures
    validation_cats_dir = os.path.join(validation_dir, 'cats') # directory with our validation cat pictures
    validation_dogs_dir = os.path.join(validation_dir, 'dogs') # directory with our validation dog pictures
```

**Understand the data**

The dataset as the following structure. There is a top level directory called cats and dogs filtered, then we have training and validation dataset. Within training dataset, there are two sub directories cats and dogs, the validation directory structure also follows the same. The validation also has two sub directories cats and dogs. Within cats directory we have images of the cats stored in jpeg format and each file has a name cat.id.jpeg and dog.id.jpeg.

So, first 2000 examples are used as training and the remaining examples are used for validation. After extracting the content, we assign variables with proper file paths for training and validation sets.

(Refer Slide Time: 03:46)



```
+ Text
dogs: [dog.2000.jpg, dog.2001.jpg, dog.2002.jpg ...]

After extracting its contents, assign variables with the proper file path for the training and validation set.

[6] train_dir = os.path.join(PATH, 'train')
    validation_dir = os.path.join(PATH, 'validation')

[7] train_cats_dir = os.path.join(train_dir, 'cats') # directory with our training cat pictures
    train_dogs_dir = os.path.join(train_dir, 'dogs') # directory with our training dog pictures
    validation_cats_dir = os.path.join(validation_dir, 'cats') # directory with our validation cat pictures
    validation_dogs_dir = os.path.join(validation_dir, 'dogs') # directory with our validation dog pictures

▼ Understand the data

Let's look at how many cats and dogs images are in the training and validation directory:

num_cats_tr = len(os.listdir(train_cats_dir))
num_dogs_tr = len(os.listdir(train_dogs_dir))
num_cats_val = len(os.listdir(validation_cats_dir))
num_dogs_val = len(os.listdir(validation_dogs_dir))
total_train = num_cats_tr + num_dogs_tr
total_val = num_cats_val + num_dogs_val

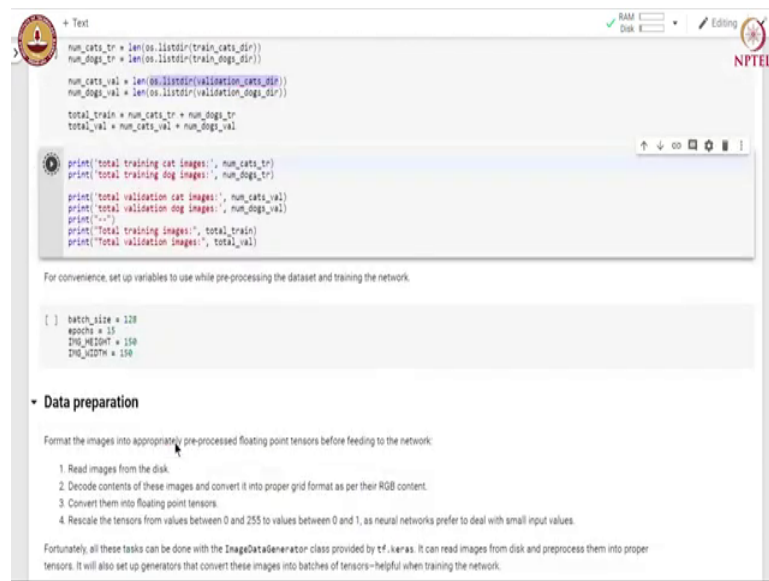
print('total training cat images:', num_cats_tr)
print('total training dog images:', num_dogs_tr)

print('total validation cat images:', num_cats_val)
print('total validation dog images:', num_dogs_val)
print('--')
print('Total training images:', total_train)
print('Total validation images:', total_val)

For convenience, set up variables to use while pre-processing the dataset and training the network.
```

Then we construct paths for training and validation directories for cats and dogs. Let us look at how many cats and dog images are there in training and validation directory. We use os.listdir command to list the content of the directory and take the length of this directory listing to calculate the number of cats and dogs in training and validation.

(Refer Slide Time: 04:21)



```
num_cats_tr = len(os.listdir(train_cats_dir))
num_dogs_tr = len(os.listdir(train_dogs_dir))

num_cats_val = len(os.listdir(validation_cats_dir))
num_dogs_val = len(os.listdir(validation_dogs_dir))

total_train = num_cats_tr + num_dogs_tr
total_val = num_cats_val + num_dogs_val

print('total training cat images:', num_cats_tr)
print('total training dog images:', num_dogs_tr)

print('total validation cat images:', num_cats_val)
print('total validation dog images:', num_dogs_val)
print('-')
print('Total training images:', total_train)
print('Total validation images:', total_val)
```

For convenience, set up variables to use while pre-processing the dataset and training the network.

```
[ ] batch_size = 128
    epochs = 15
    IMG_HEIGHT = 150
    IMG_WIDTH = 150
```

**Data preparation**

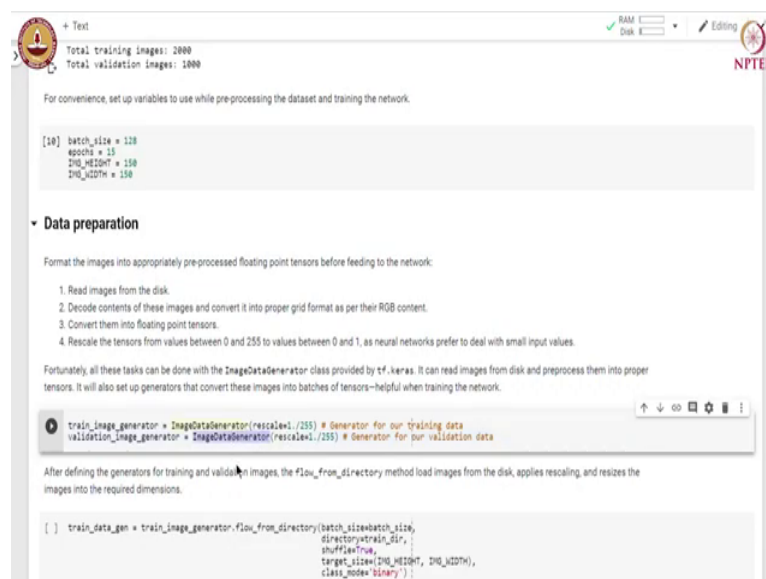
Format the images into appropriately pre-processed floating point tensors before feeding to the network:

1. Read images from the disk.
2. Decode contents of these images and convert it into proper grid format as per their RGB content.
3. Convert them into floating point tensors.
4. Rescale the tensors from values between 0 and 255 to values between 0 and 1, as neural networks prefer to deal with small input values.

Fortunately, all these tasks can be done with the `ImageDataGenerator` class provided by `tf.keras`. It can read images from disk and preprocess them into proper tensors. It will also set up generators that convert these images into batches of tensors—helpful when training the network.

You can see that we are using 2000 images for training and 1000 images total for validation. 1000 cat images are used for training and 500 cat images are used for validation. The same proportion of images are used for training and validation from dog class. So, let us setup some variables like batch size, epochs and the height and width of the image.

(Refer Slide Time: 05:09)



```
Total training Images: 2000
Total validation Images: 1000
```

For convenience, set up variables to use while pre-processing the dataset and training the network.

```
[10] batch_size = 128
    epochs = 15
    IMG_HEIGHT = 150
    IMG_WIDTH = 150
```

**Data preparation**

Format the images into appropriately pre-processed floating point tensors before feeding to the network:

1. Read images from the disk.
2. Decode contents of these images and convert it into proper grid format as per their RGB content.
3. Convert them into floating point tensors.
4. Rescale the tensors from values between 0 and 255 to values between 0 and 1, as neural networks prefer to deal with small input values.

Fortunately, all these tasks can be done with the `ImageDataGenerator` class provided by `tf.keras`. It can read images from disk and preprocess them into proper tensors. It will also set up generators that convert these images into batches of tensors—helpful when training the network.

```
train_image_generator = ImageDataGenerator(rescale=1./255) # Generator for our training data
validation_image_generator = ImageDataGenerator(rescale=1./255) # Generator for our validation data
```

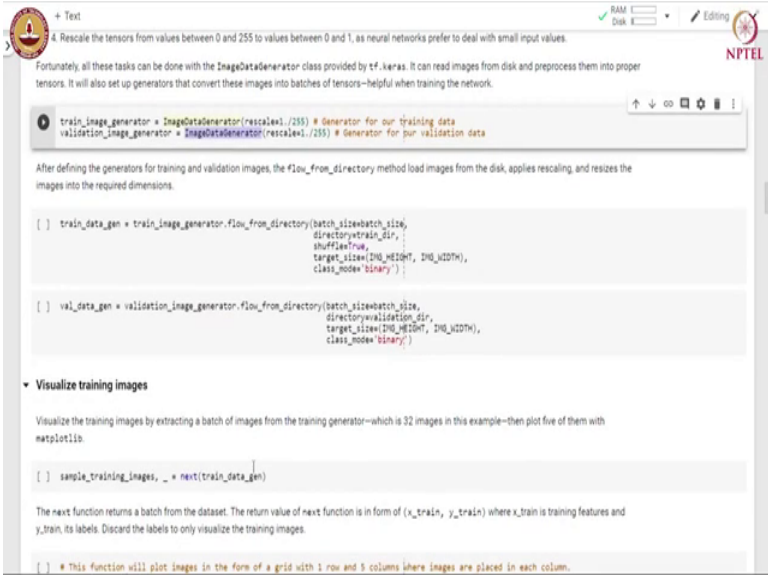
After defining the generators for training and validation images, the `flow_from_directory` method load images from the disk, applies rescaling, and resizes the images into the required dimensions.

```
[ ] train_data_gen = train_image_generator.flow_from_directory(batch_size=batch_size,
    directory=train_dir,
    shuffle=True,
    target_size=(IMG_HEIGHT, IMG_WIDTH),
    class_mode='binary')
```

Let us prepare the data for training. So, we perform the following steps we will first read the images from the disk, then we will decode the content of this images and convert them into proper grid format as per their RGB content. Then we convert them into floating point tensors and then we rescale this tensors from values between 0 and 255 to values between 0 and 1.

So, all this task are done by image data generator class provided by tf.keras. It can read images from the disc and preprocess them into proper tensors. It will also setup generators that convert this images into batches of tensors which is very helpful during the training. So, we setup image data generator for training and validation set.

(Refer Slide Time: 06:22)



```
4. Rescale the tensors from values between 0 and 255 to values between 0 and 1, as neural networks prefer to deal with small input values.

Fortunately, all these tasks can be done with the ImageDataGenerator class provided by tf.keras. It can read images from disk and preprocess them into proper tensors. It will also set up generators that convert these images into batches of tensors—helpful when training the network.

train_image_generator = ImageDataGenerator(rescale=1./255) # Generator for our training data
validation_image_generator = ImageDataGenerator(rescale=1./255) # Generator for our validation data

After defining the generators for training and validation images, the flow_from_directory method load images from the disk, applies rescaling, and resizes the images into the required dimensions.

[ ] train_data_gen = train_image_generator.flow_from_directory(batch_size=batch_size,
                                                             directory=train_dir,
                                                             shuffle=True,
                                                             target_size=(IMG_HEIGHT, IMG_WIDTH),
                                                             class_mode='binary')

[ ] val_data_gen = validation_image_generator.flow_from_directory(batch_size=batch_size,
                                                                directory=validation_dir,
                                                                target_size=(IMG_HEIGHT, IMG_WIDTH),
                                                                class_mode='binary')

▼ Visualize training images

Visualize the training images by extracting a batch of images from the training generator—which is 32 images in this example—then plot five of them with matplotlib.

[ ] sample_training_images, _ = next(train_data_gen)

The next function returns a batch from the dataset. The return value of next function is in form of (x_train, y_train) where x_train is training features and y_train, its labels. Discard the labels to only visualize the training images.

[ ] # This function will plot images in the form of a grid with 1 row and 5 columns where images are placed in each column.
```

After defining these generators, we use flow from directory method to load images from the disc, apply rescaling and resizing of the image into required dimension. So, here the target image size is 150 x 150 and we want to shuffle the training data. We do not shuffle the data in the validation set. We also specify the batch size and the directories where data is stored.

(Refer Slide Time: 07:02)

```
target_size=(IMG_HEIGHT, IMG_WIDTH),
class_mode='binary')

Visualize training images

Visualize the training images by extracting a batch of images from the training generator—which is 32 images in this example—then plot five of them with matplotlib.

[ ] sample_training_images, _ = next(train_data_gen)

The next function returns a batch from the dataset. The return value of next function is in form of (x_train, y_train) where x_train is training features and y_train, its labels. Discard the labels to only visualize the training images.

[ ] # This function will plot images in the form of a grid with 1 row and 5 columns where images are placed in each column.
def plot_images(images_arr):
    fig, axes = plt.subplots(1, 5, figsize=(20,20))
    axes = axes.flatten()
    for img, ax in zip(images_arr, axes):
        ax.imshow(img)
        ax.axis('off')
    plt.tight_layout()
    plt.show()

[ ] plot_images(sample_training_images[5])

Create the model

The model consists of three convolution blocks with a max pool layer in each of them. There's a fully connected layer with 512 units on top of it that is activated by a relu activation function. The model outputs class probabilities based on binary classification by the sigmoid activation function.

[ ] model = Sequential([
    Conv2D(16, 3, padding='same', activation='relu', input_shape=(IMG_HEIGHT, IMG_WIDTH, 3)),
    MaxPooling2D(),
```

Let us visualize the training images by extracting a batch of images from the image generator. Let us plot five of these images.

(Refer Slide Time: 07:28)

```
# This function will plot images in the form of a grid with 1 row and 5 columns where images are placed in each column.
def plot_images(images_arr):
    fig, axes = plt.subplots(1, 5, figsize=(20,20))
    axes = axes.flatten()
    for img, ax in zip(images_arr, axes):
        ax.imshow(img)
        ax.axis('off')
    plt.tight_layout()
    plt.show()

plot_images(sample_training_images[5])

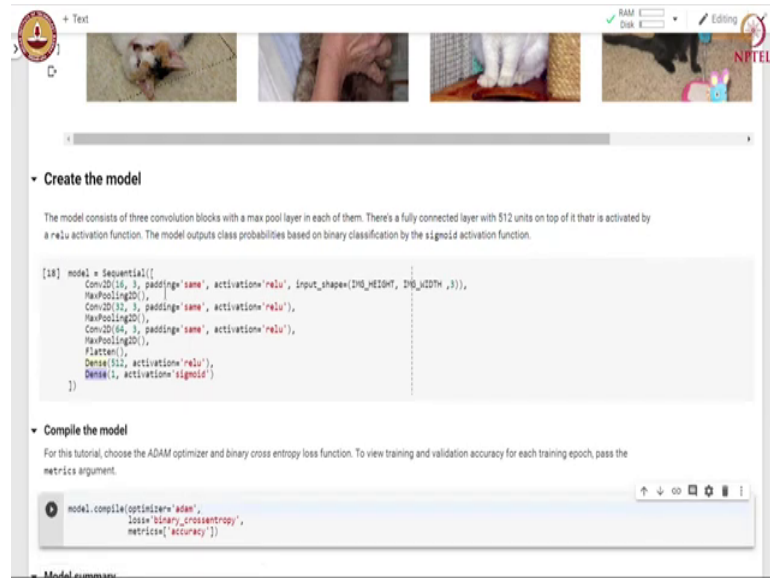
Create the model

The model consists of three convolution blocks with a max pool layer in each of them. There's a fully connected layer with 512 units on top of it that is activated by a relu activation function. The model outputs class probabilities based on binary classification by the sigmoid activation function.

[ ] model = Sequential([
```

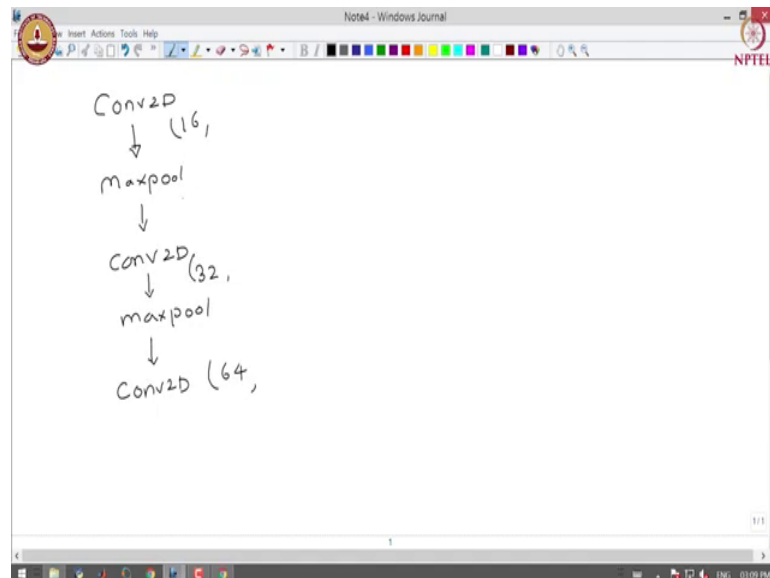
We use matplotlib for plotting these images.

(Refer Slide Time: 07:45)



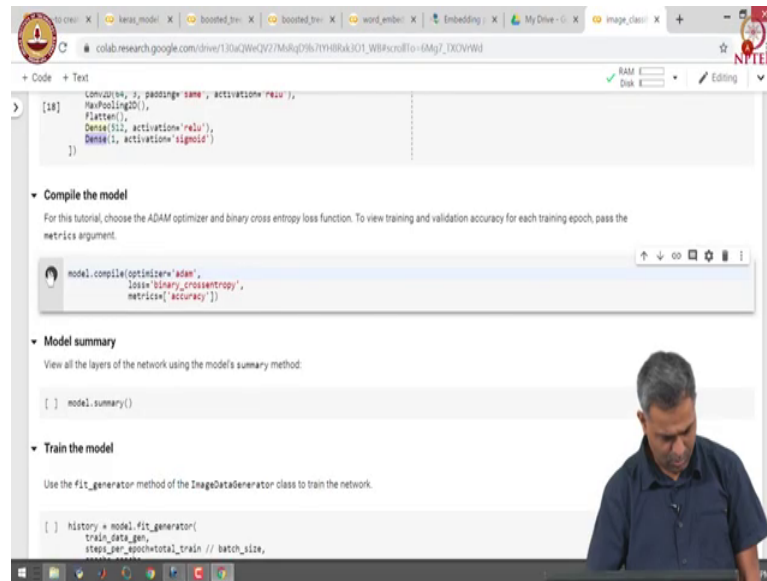
Let us create a model a CNN model for classifying cats and dogs. So, we use three convolution blocks with a MaxPool layer in each one of them. Then we use a fully connected layer with 512 units on top of that with Relu activation. And the model outputs class probabilities based on binary classification by sigmoid activation function in the output layer. Let us look at the structure of the model.

(Refer Slide Time: 08:47)



So, we use a Conv2D followed by MaxPool, then another convolution layer.

(Refer Slide Time: 09:37)



```
[18]: conv2d_4 (Conv2D) (None, 28, 28, 32) 4640
      max_pooling2d_1 (MaxPooling2D) (None, 14, 14, 32) 0
      flatten_1 (Flatten) (None, 4096) 0
      dense_1 (Dense) (None, 10) 105
      dense_2 (Dense) (None, 1) 1
      sigmoid_1 (Sigmoid) (None, 1) 0
      Total params: 47,537
      Trainable params: 47,537
      Non-trainable params: 0
```

**Compile the model**

For this tutorial, choose the Adam optimizer and binary cross entropy loss function. To view training and validation accuracy for each training epoch, pass the metrics argument.

```
model.compile(optimizer='adam',
              loss='binary_crossentropy',
              metrics=['accuracy'])
```

**Model summary**

View all the layers of the network using the model's summary method:

```
[ ] model.summary()
```

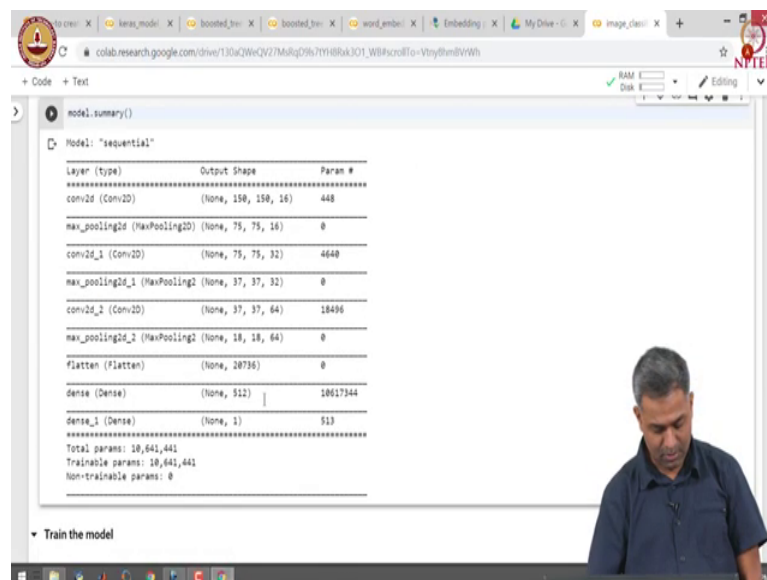
**Train the model**

Use the `fit_generator` method of the `ImageDataGenerator` class to train the network.

```
[ ] history = model.fit_generator(
    train_data_gen,
    steps_per_epoch=train_data_gen.n,
    validation_data_gen=validation_data_gen,
    validation_steps=validation_data_gen.n,
    epochs=10)
```

Reduce *binary\_crossentropy* as a loss with *Adam* as an optimizer and *accuracy* as a metric to track.

(Refer Slide Time: 09:52)



```
Model: "sequential"
Layer (type) Output Shape Param #
-----
conv2d_1 (Conv2D) (None, 28, 28, 32) 4640
max_pooling2d_1 (MaxPooling2D) (None, 14, 14, 32) 0
conv2d_2 (Conv2D) (None, 14, 14, 32) 4640
max_pooling2d_2 (MaxPooling2D) (None, 7, 7, 32) 0
conv2d_3 (Conv2D) (None, 7, 7, 64) 18496
max_pooling2d_3 (MaxPooling2D) (None, 3, 3, 64) 0
flatten_1 (Flatten) (None, 576) 0
dense_1 (Dense) (None, 10) 5801
dense_2 (Dense) (None, 1) 1
sigmoid_1 (Sigmoid) (None, 1) 0
Total params: 10,641,441
Trainable params: 10,641,441
Non-trainable params: 0
```

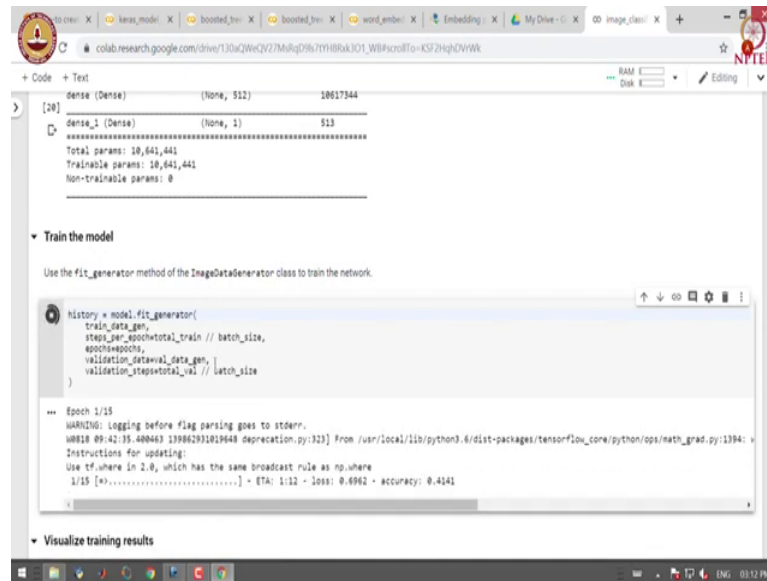
**Train the model**

```
[ ] history = model.fit_generator(
    train_data_gen,
    steps_per_epoch=train_data_gen.n,
    validation_data_gen=validation_data_gen,
    validation_steps=validation_data_gen.n,
    epochs=10)
```

So, you can see that we have modelled with more than 10 million parameters.



(Refer Slide Time: 10:25)



The screenshot shows a Google Colab notebook with the following content:

```
dense (Dense) (None, 512) 10617344
[20] dense_1 (Dense) (None, 1) 512
-----
Total params: 10,641,441
Trainable params: 10,641,441
Non-trainable params: 0
```

**Train the model**

Use the `fit_generator` method of the `ImageDataGenerator` class to train the network.

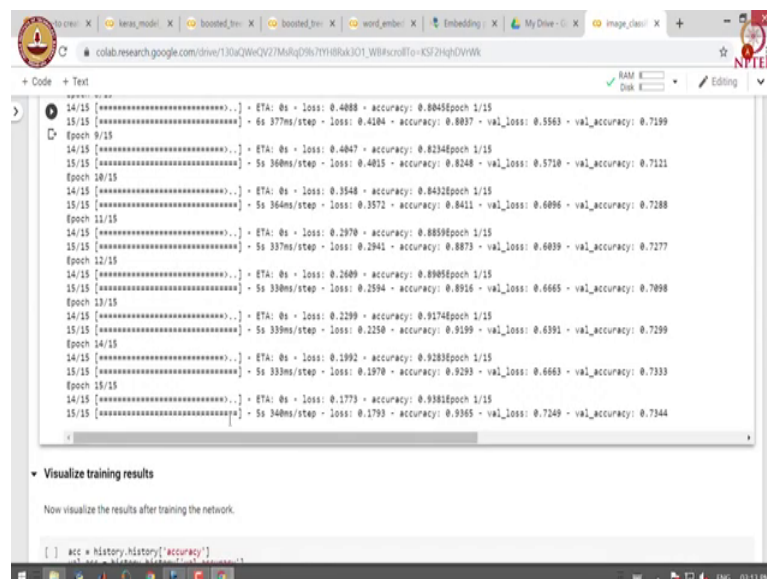
```
history = model.fit_generator(
    train_data_gen,
    steps_per_epoch=total_train // batch_size,
    epochs=epochs,
    validation_data=val_data_gen,
    validation_steps=total_val // batch_size
)
```

Epoch 1/15  
WARNING: Logging before flag parsing goes to stderr.  
WARNING: 09:42:35.400463 139862931819648 deprecation.py:323] From /usr/local/lib/python3.6/dist-packages/tensorflow\_core/python/ops/math\_grad.py:1394: instructions for updating:  
Use tf.where in 2.0, which has the same broadcast rule as np.where  
1/15 [0]..... ETA: 1:12 - loss: 0.6962 - accuracy: 0.4141

**Visualize training results**

Let us train the model for 15 epochs.

(Refer Slide Time: 10:35)



The screenshot shows the same Google Colab notebook after 15 epochs of training. The output displays the following metrics for each epoch:

Epoch	ETA	loss	accuracy	val_loss	val_accuracy
1/15	0s	0.4088	0.8045	0.5563	0.7199
2/15	6s	0.377ms/step	0.4104	0.8037	0.7199
3/15	0s	0.4047	0.8234	0.5710	0.7121
4/15	360ms/step	0.4015	0.8248	0.5710	0.7121
5/15	0s	0.3548	0.8432	0.6096	0.7288
6/15	364ms/step	0.3572	0.8411	0.6096	0.7288
7/15	0s	0.2970	0.8859	0.6039	0.7277
8/15	337ms/step	0.2941	0.8873	0.6039	0.7277
9/15	0s	0.2609	0.8909	0.6665	0.7090
10/15	330ms/step	0.2594	0.8916	0.6665	0.7090
11/15	0s	0.2299	0.9174	0.6391	0.7299
12/15	339ms/step	0.2250	0.9199	0.6391	0.7299
13/15	0s	0.1992	0.9283	0.6663	0.7333
14/15	333ms/step	0.1970	0.9293	0.6663	0.7333
15/15	0s	0.1773	0.9381	0.7249	0.7344
16/15	340ms/step	0.1793	0.9365	0.7249	0.7344

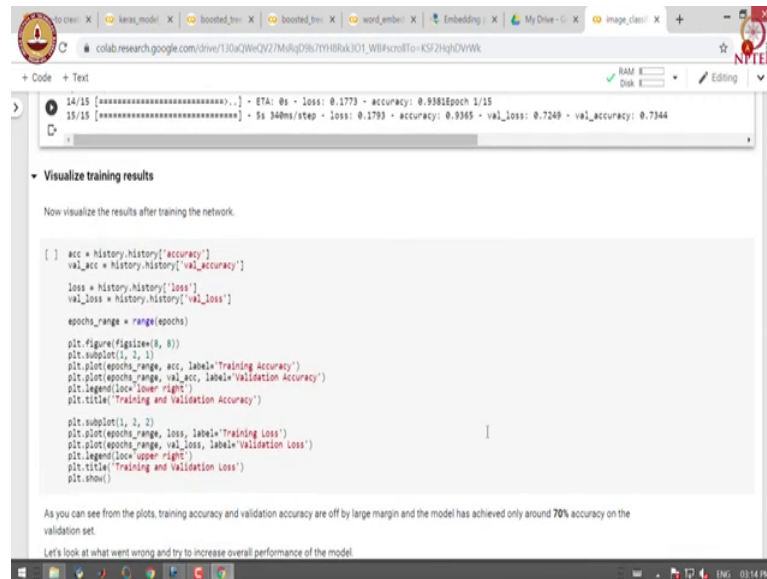
**Visualize training results**

Now visualize the results after training the network.

```
acc = history.history['accuracy']
```

So after 15 iteration we got accuracy closed to 94 % on the training set and 73 % on the validation set.

(Refer Slide Time: 10:48)



The screenshot shows a Google Colab notebook interface. At the top, the terminal displays training progress for epochs 14 and 15. Below the terminal, a code cell is titled "Visualize training results". The code in the cell uses the `History` object to extract accuracy and loss data, then uses `matplotlib` to create two subplots: "Training and Validation Accuracy" and "Training and Validation Loss". The accuracy plot shows training accuracy (blue line) rising sharply while validation accuracy (orange line) plateaus around 0.7. The loss plot shows training loss (blue line) decreasing steadily while validation loss (orange line) initially decreases and then starts to increase after approximately 50 epochs, indicating overfitting.

```
[ ] acc = history.history['accuracy']
val_acc = history.history['val_accuracy']

loss = history.history['loss']
val_loss = history.history['val_loss']

epochs_range = range(epochs)

plt.figure(figsize=(8, 8))
plt.subplot(1, 2, 1)
plt.plot(epochs_range, acc, label='Training Accuracy')
plt.plot(epochs_range, val_acc, label='Validation Accuracy')
plt.legend(loc='lower right')
plt.title('Training and Validation Accuracy')

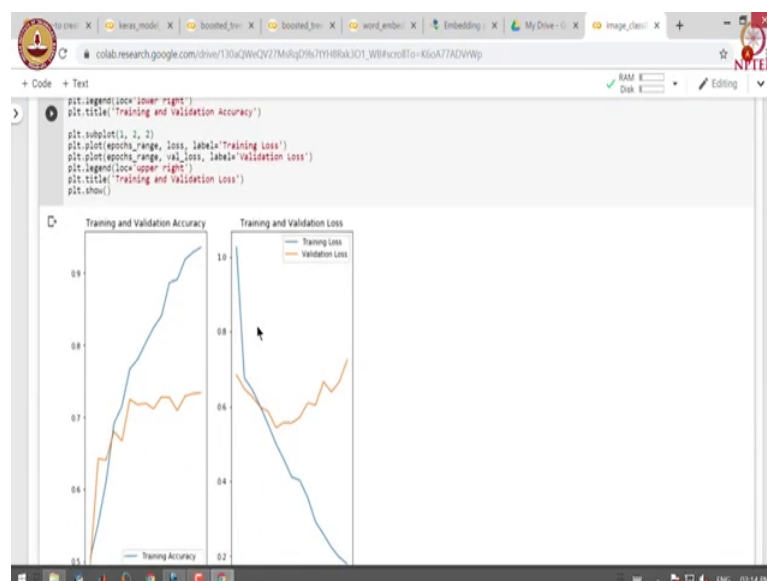
plt.subplot(1, 2, 2)
plt.plot(epochs_range, loss, label='Training Loss')
plt.plot(epochs_range, val_loss, label='Validation Loss')
plt.legend(loc='upper right')
plt.title('Training and Validation Loss')
plt.show()
```

As you can see from the plots, training accuracy and validation accuracy are off by large margin and the model has achieved only around 70% accuracy on the validation set.

Let's look at what went wrong and try to increase overall performance of the model.

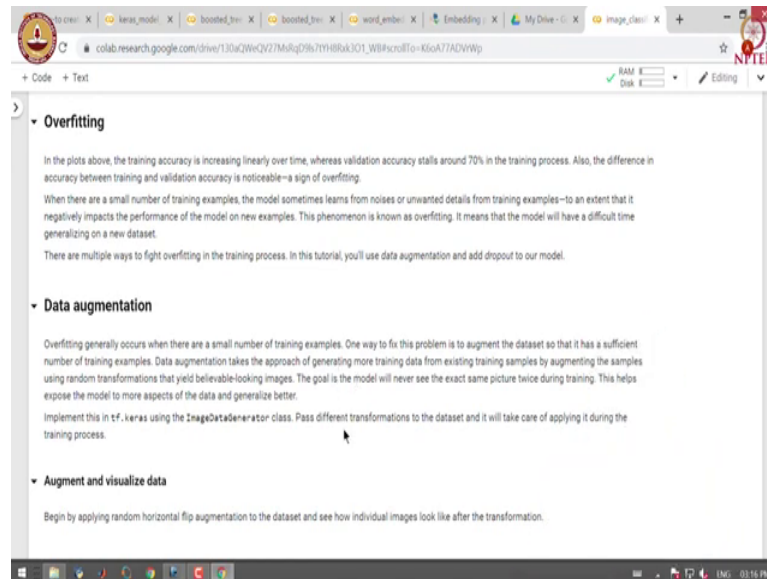
Let us visualize the training and validation accuracy across epochs.

(Refer Slide Time: 10:56)



You can see that as we trained for more epochs, the training the training accuracy kept raising whereas, the validation accuracy plateaued after some time. We also see similar trend in the training and validation loss where you found that training loss was constantly decreasing as you train for more epochs, but validation loss initially declined, but then grows steadily after certain epochs after about 50 epochs.

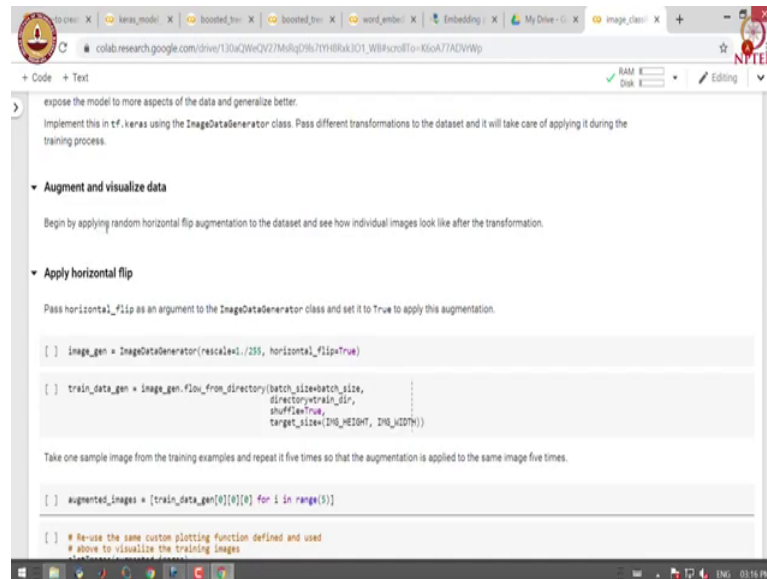
(Refer Slide Time: 11:37)



So, this points to the fact that the model is over fitting. So, you have to explore strategies to increase the performance of the model by reducing the over-fitting. So, we will use data augmentation and dropout as two strategies for fighting over-fitting problem. In data augmentation, we will take the existing images and perform certain transformation on them to gather more data.

So, what kind of transformations we do? We can rotate, translate, change the scale of the image to create more and more examples of the image such that model is less likely to over-fit with more data.

(Refer Slide Time: 12:39)



The screenshot shows a Google Colab notebook with the following content:

- Text:** "expose the model to more aspects of the data and generalize better. Implement this in tf.keras using the `ImageDataGenerator` class. Pass different transformations to the dataset and it will take care of applying it during the training process."
- Section: Augment and visualize data**
  - Text:** "Begin by applying random horizontal flip augmentation to the dataset and see how individual images look like after the transformation."
  - Section: Apply horizontal flip**
    - Text:** "Pass `horizontal_flip` as an argument to the `ImageDataGenerator` class and set it to `True` to apply this augmentation."
    - Code:**

```
[ ] image_gen = ImageDataGenerator(rescale=1./255, horizontal_flip=True)

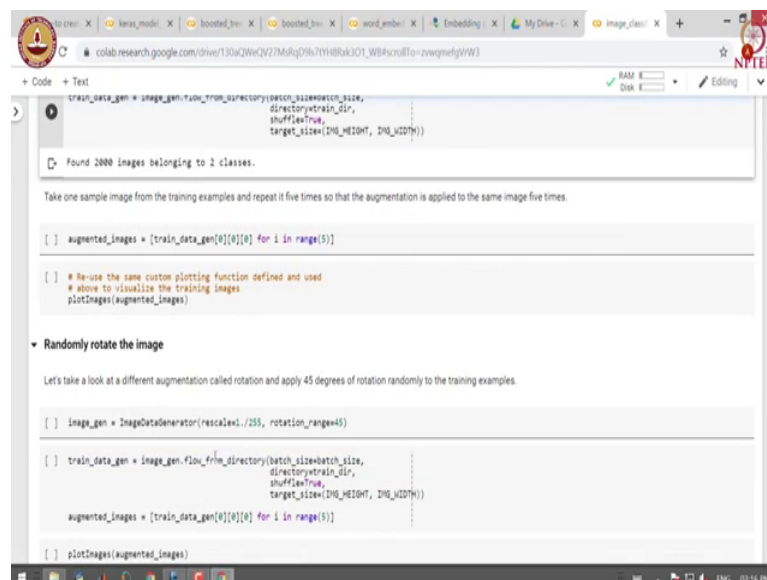
[ ] train_data_gen = image_gen.flow_from_directory(batch_size=batch_size,
                                                  directory=train_dir,
                                                  shuffle=True,
                                                  target_size=(IMG_HEIGHT, IMG_WIDTH))
```
    - Text:** "Take one sample image from the training examples and repeat it five times so that the augmentation is applied to the same image five times."
    - Code:**

```
[ ] augmented_images = [train_data_gen[0][0] for i in range(5)]

[ ] # Re-use the same custom plotting function defined and used
    # above to visualize the training images
    plotImages(augmented_images)
```

So, let us try to augment the data and visualize it. We can perform random horizontal flip augmentation to the dataset. We can use `ImageDataGenerator` with `horizontal flip = True` and let us generate the data.

(Refer Slide Time: 13:09)



The screenshot shows a Google Colab notebook with the following content:

- Code:**

```
train_data_gen = image_gen.flow_from_directory(batch_size=batch_size,
                                              directory=train_dir,
                                              shuffle=True,
                                              target_size=(IMG_HEIGHT, IMG_WIDTH))
```
- Text:** "Found 2880 images belonging to 2 classes."
- Text:** "Take one sample image from the training examples and repeat it five times so that the augmentation is applied to the same image five times."
- Code:**

```
[ ] augmented_images = [train_data_gen[0][0] for i in range(5)]

[ ] # Re-use the same custom plotting function defined and used
    # above to visualize the training images
    plotImages(augmented_images)
```
- Section: Randomly rotate the image**
  - Text:** "Let's take a look at a different augmentation called rotation and apply 45 degrees of rotation randomly to the training examples."
  - Code:**

```
[ ] image_gen = ImageDataGenerator(rescale=1./255, rotation_range=45)

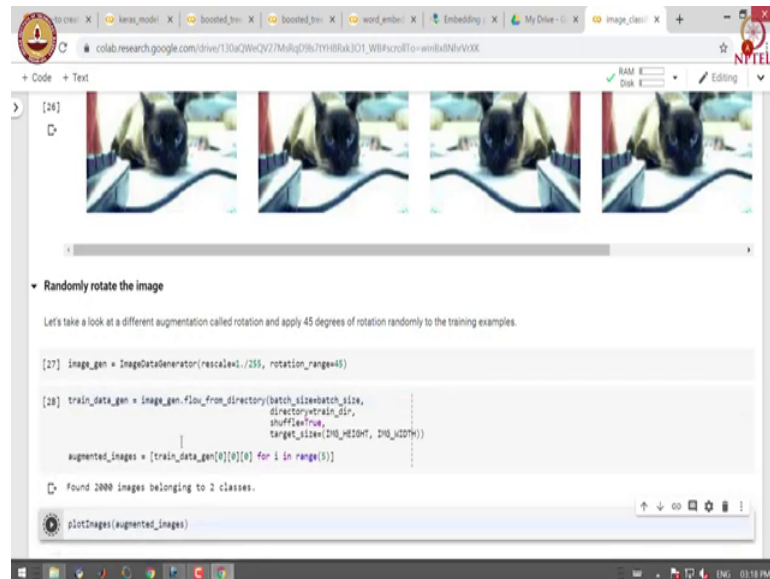
[ ] train_data_gen = image_gen.flow_from_directory(batch_size=batch_size,
                                                  directory=train_dir,
                                                  shuffle=True,
                                                  target_size=(IMG_HEIGHT, IMG_WIDTH))

    augmented_images = [train_data_gen[0][0] for i in range(5)]

[ ] plotImages(augmented_images)
```

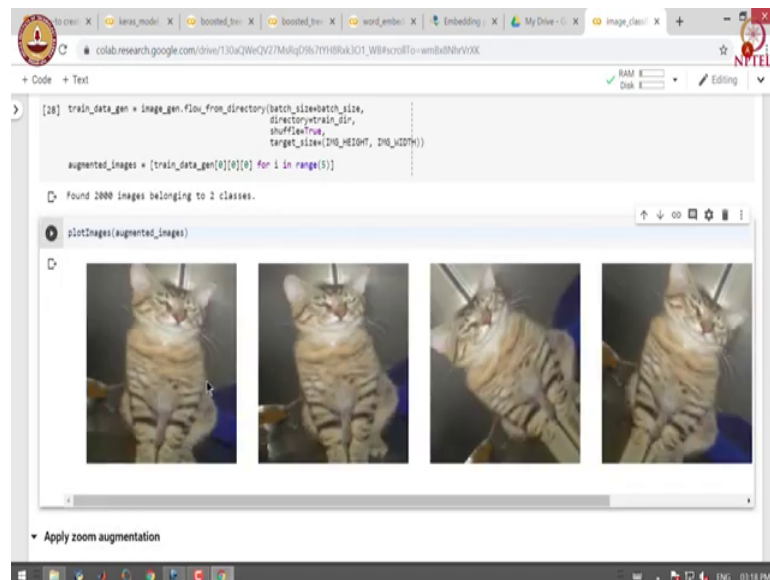
Let us take one example and see how the data got augmented with horizontal flipping.

(Refer Slide Time: 13:22)



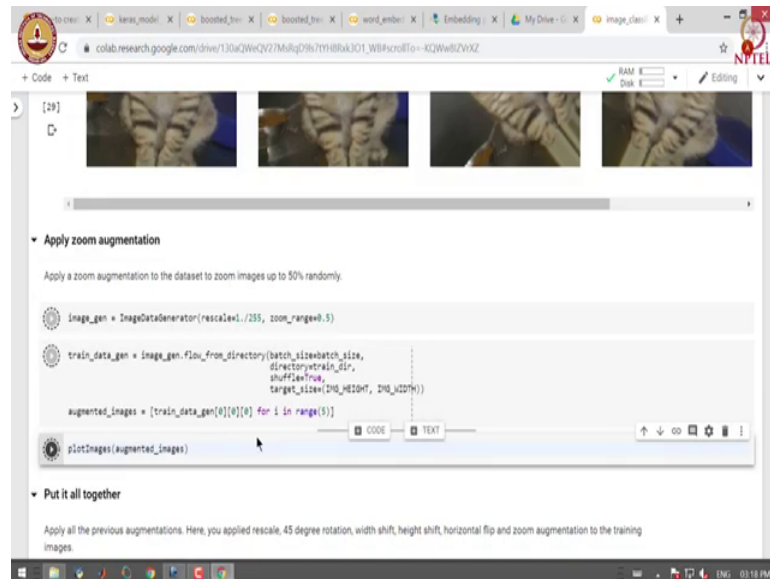
So, after horizontal flipping, we got more examples that were generated from a single example. We can also do a different kind of augmentation by rotation. Let us apply 45 degree rotation randomly to the training examples.

(Refer Slide Time: 14:17)



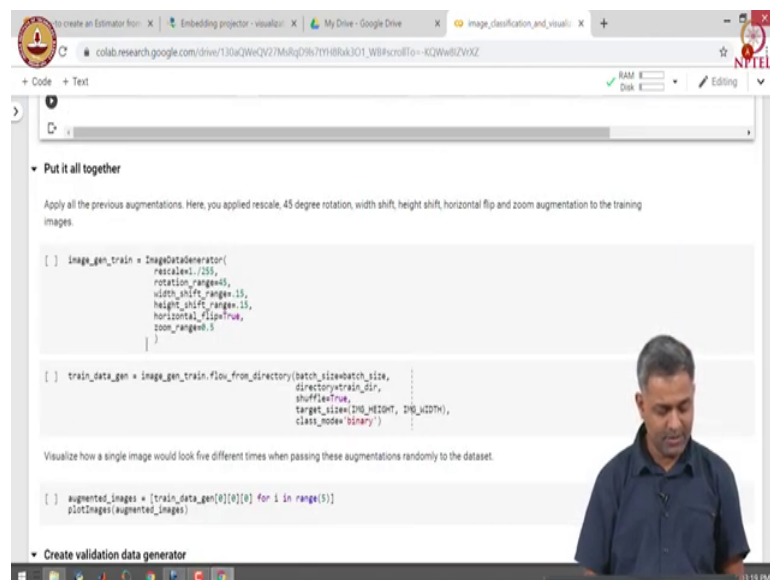
Now, you can see that the same picture of the cat was rotated in different orientations with different angles and by applying this particular data augmentation strategy, we created more pictures from a single picture.

(Refer Slide Time: 14:45)



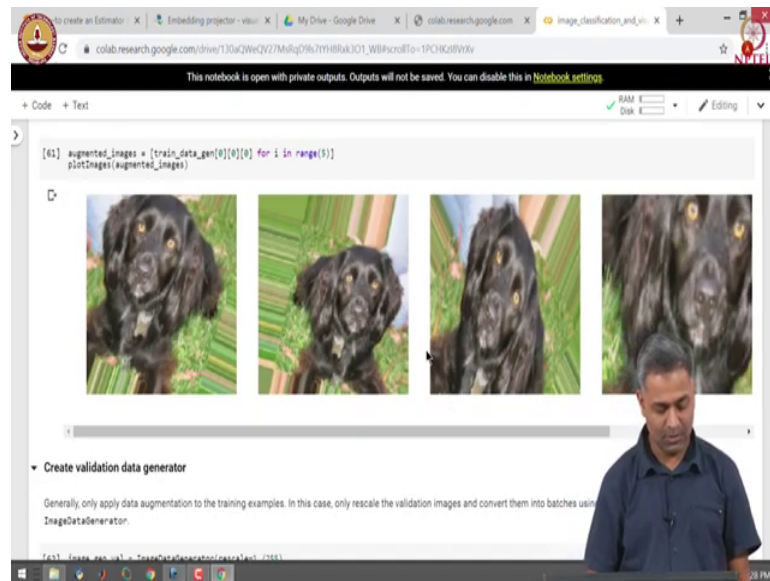
In addition to that we can also apply zoom augmentation by specifying the zoom range. Let us see how zoom augmentation look like.

(Refer Slide Time: 14:59)



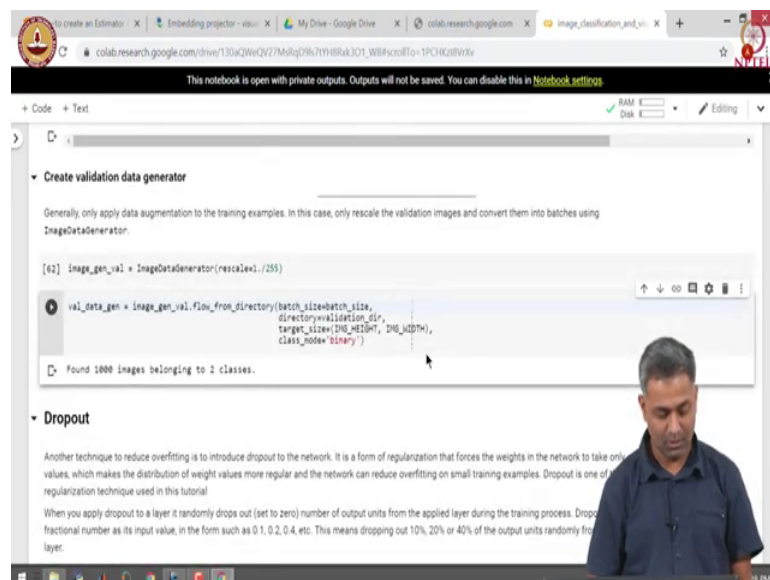
Let us put together all this augmentations strategies. So, we use ImageDataGenerator and apply rescaling, then rotation of 45 degrees, width shift, height shift and horizontal flip and the zoom augmentation to the training images.

(Refer Slide Time: 15:50)



So we apply all this augmentation techniques to the trainings set and obtain more data by augmenting the original images. These are some of the augmented pictures of a dog where we have the original picture and these are the pictures that were obtained by applying various data augmentations strategies.

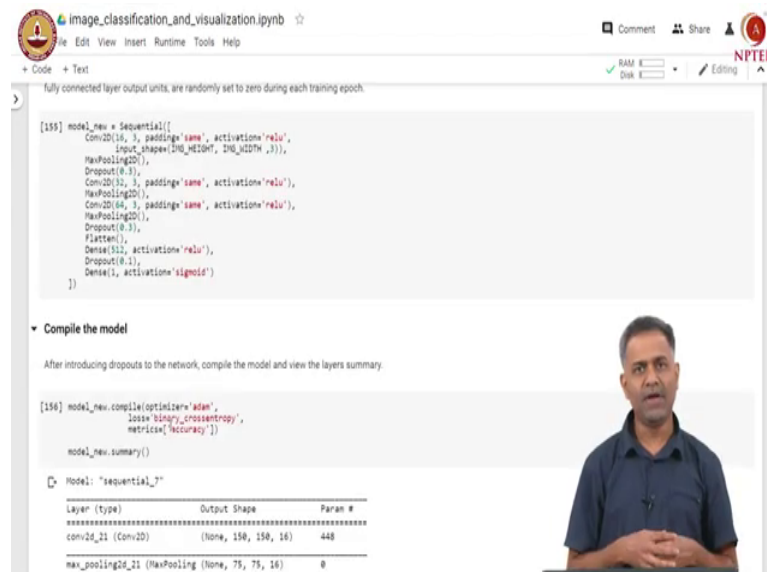
(Refer Slide Time: 16:14)



We apply the data augmentation only to the training examples, we do not apply data augmentation on the validation examples except rescaling them. And we also convert the validation examples into batches using ImageDataGenerator.

Another technique to reduce over fitting is dropout.

(Refer Slide Time: 16:44)



```
[155] model_new = Sequential([
    Conv2D(16, 3, padding='same', activation='relu',
        input_shape=(256, 256, 3)),
    MaxPooling2D(),
    Dropout(0.3),
    Conv2D(32, 3, padding='same', activation='relu'),
    MaxPooling2D(),
    Conv2D(64, 3, padding='same', activation='relu'),
    MaxPooling2D(),
    Dropout(0.3),
    Flatten(),
    Dense(512, activation='relu'),
    Dropout(0.1),
    Dense(1, activation='sigmoid')
])
```

Compile the model

After introducing dropouts to the network, compile the model and view the layers summary.

```
[156] model_new.compile(optimizer='adam',
    loss='binary_crossentropy',
    metrics=['accuracy'])

model_new.summary()
```

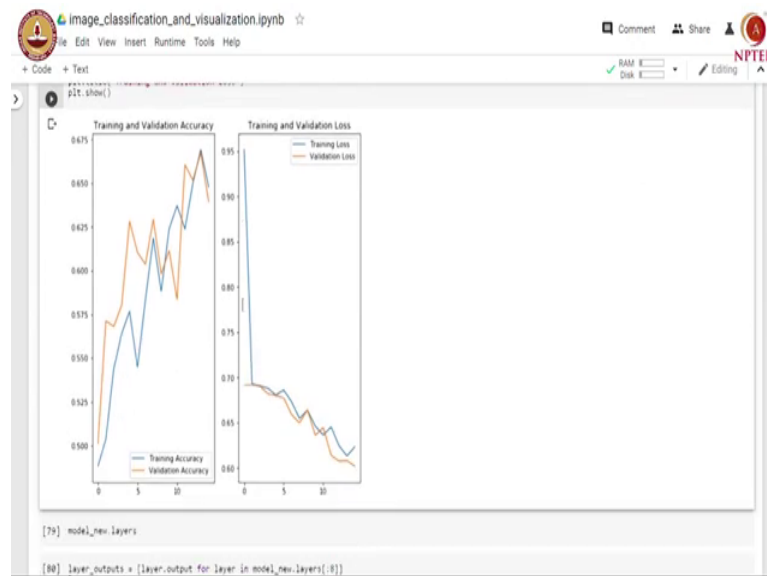
Model: "sequential\_7"

Layer (type)	Output Shape	Param #
=====		
conv2d_11 (Conv2D)	(None, 128, 128, 16)	448
=====		
max_pooling2d_11 (MaxPooling2D)	(None, 64, 64, 16)	0

We had dropout to the model after augmenting the data through various augmentation schemes. We had dropout of 0.3 each after couple of max pooling layers and Dropout of 0.1 after the fully connected layer. The effect of the dropout is that the randomly 30 % of max-pooling nodes and 10 percent of fully connected nodes are set to 0 during each epoch. We compile the model and train it and we can see that at the end of end of 15 epoch. We have a training accuracy of 64 % and validation accuracy of 63 %.



(Refer Slide Time: 17:32)

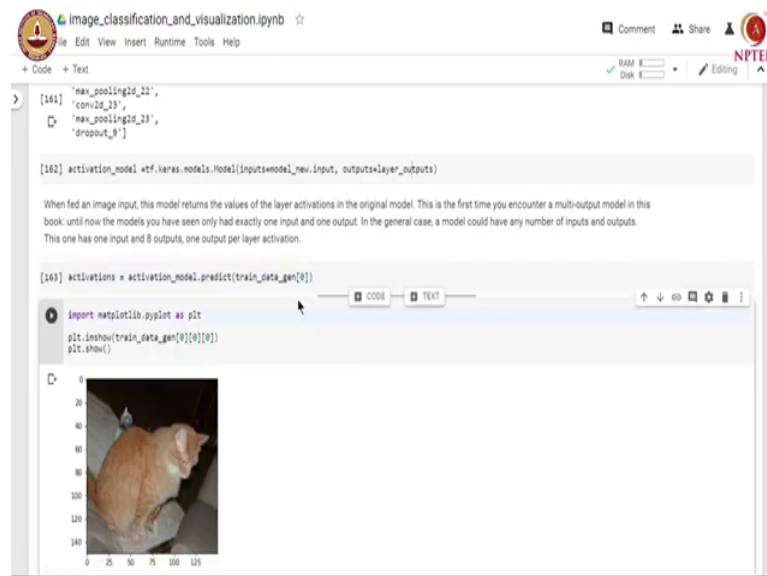


And if we plot the losses, you can see that both the training and validation loss is trending together and there is a lesser over fitting than the earlier model.

(Refer Slide Time: 17:50)

Let us visualize what this particular CNN is learning. So, we will take the output of the model after every layer. So, there are eight different layers convolution followed by pooling, then there is a dropout then, convolution pooling, convolution pooling and dropout.

(Refer Slide Time: 18:27)

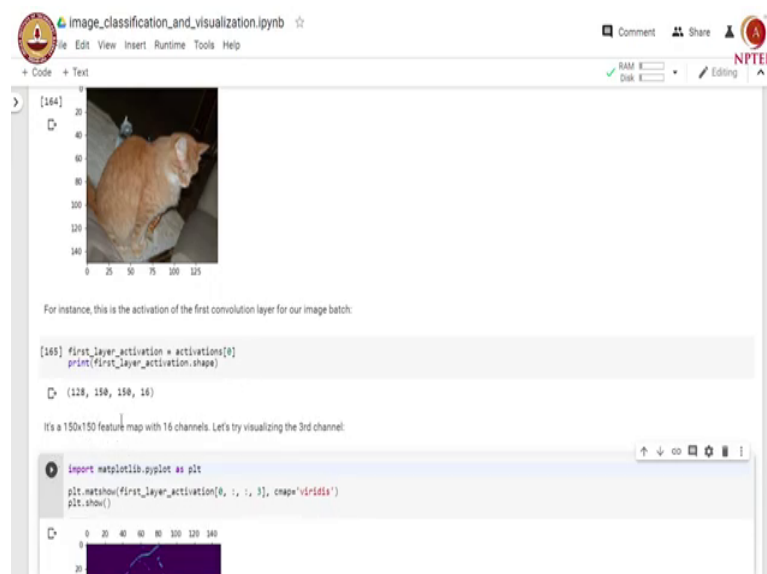


```
[161] 'max_pooling2d_12',  
      'conv2d_13',  
      'max_pooling2d_13',  
      'dropout_9']  
  
[162] activation_model = tf.keras.models.Model(inputs=model.input, outputs=layer_outputs)  
  
When fed an image input, this model returns the values of the layer activations in the original model. This is the first time you encounter a multi-output model in this book: until now the models you have seen only had exactly one input and one output. In the general case, a model could have any number of inputs and outputs. This one has one input and 8 outputs, one output per layer activation.  
  
[163] activations = activation_model.predict(train_data_gen[0])  
  
import matplotlib.pyplot as plt  
plt.imshow(train_data_gen[0][0][0])  
plt.show()
```

The screenshot shows a Jupyter Notebook interface with the title 'image\_classification\_and\_visualization.ipynb'. The code cell contains the following Python code: [161] 'max\_pooling2d\_12', 'conv2d\_13', 'max\_pooling2d\_13', 'dropout\_9'] [162] activation\_model = tf.keras.models.Model(inputs=model.input, outputs=layer\_outputs) When fed an image input, this model returns the values of the layer activations in the original model. This is the first time you encounter a multi-output model in this book: until now the models you have seen only had exactly one input and one output. In the general case, a model could have any number of inputs and outputs. This one has one input and 8 outputs, one output per layer activation. [163] activations = activation\_model.predict(train\_data\_gen[0]) import matplotlib.pyplot as plt plt.imshow(train\_data\_gen[0][0][0]) plt.show(). Below the code, there is a plot of a cat's face, which is the first image in the batch. The plot has x and y axes ranging from 0 to 125.

We get the output after every layer. So, this is different from the models that we have seen so far. This is the first time we are encountering a multi output model until now. We have models that were giving exactly one output at the end of the final layer here, we are gathering output after every layer. Let us look at the predictions coming from the activation model and let us plot those activations.

(Refer Slide Time: 19:07)



```
[164] first_layer_activation = activations[0]  
print(first_layer_activation.shape)  
  
(128, 128, 16)  
  
It's a 128x128 feature map with 16 channels. Let's try visualizing the 3rd channel:  
  
import matplotlib.pyplot as plt  
plt.matshow(first_layer_activation[0, :, :3], cmap='viridis')  
plt.show()
```

The screenshot shows a Jupyter Notebook interface with the title 'image\_classification\_and\_visualization.ipynb'. The code cell contains the following Python code: [164] first\_layer\_activation = activations[0] print(first\_layer\_activation.shape) (128, 128, 16) It's a 128x128 feature map with 16 channels. Let's try visualizing the 3rd channel: import matplotlib.pyplot as plt plt.matshow(first\_layer\_activation[0, :, :3], cmap='viridis') plt.show(). Below the code, there is a plot of the first layer's activation for the first image in the batch. The plot shows a heatmap of the first layer's activation for the first image in the batch. The plot has x and y axes ranging from 0 to 125.

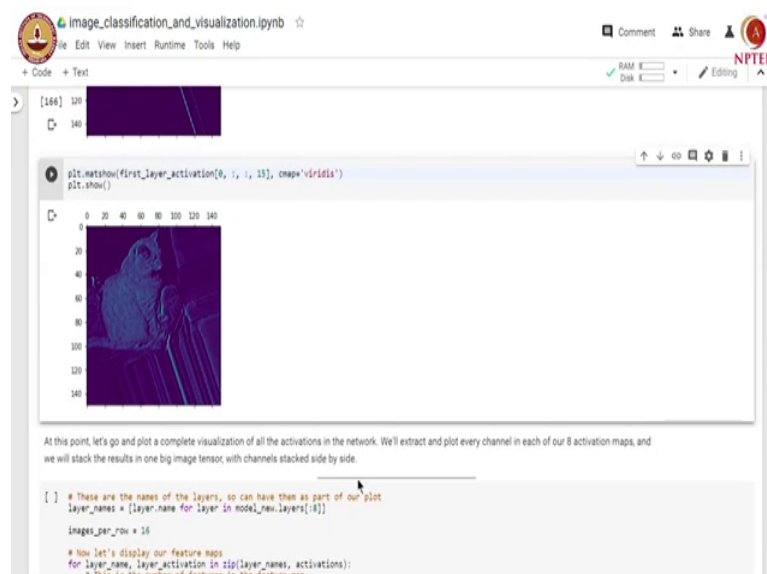
So, this is the input this is the first image that is input to us and the first level activation is essentially 150 x 150 feature map with 16 channels.

(Refer Slide Time: 19:23)



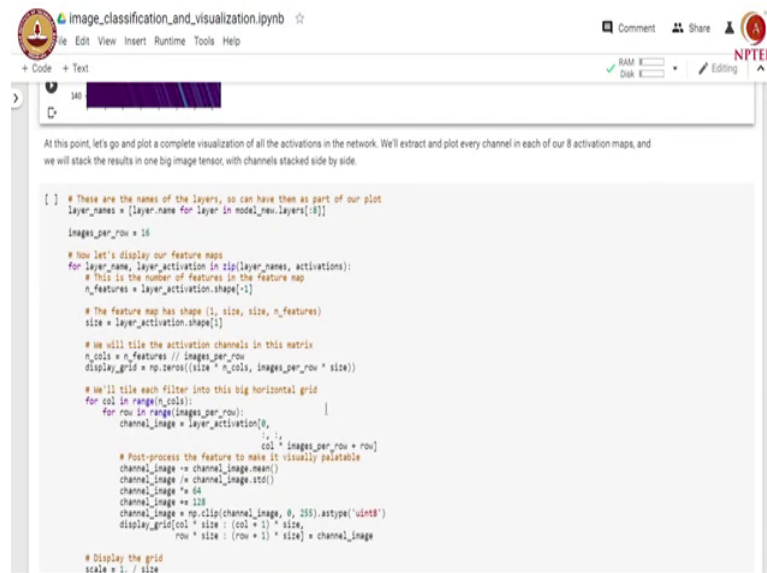
So, let us look at the third channel. You can see that the third channel is roughly detecting the edges from the cat picture.

(Refer Slide Time: 19:36)



If you look at the 15th channel, it is identifying probably the overall shape of the image. At this point let us go and plot a complete visualization of all the activations in the network.

(Refer Slide Time: 19:54)



```
[ ] # These are the names of the layers, so can have them as part of our plot
layer_names = [layer.name for layer in model_new.layers[1:]]

images_per_row = 16

# Now let's display our feature maps
for layer_name, layer_activation in zip(layer_names, activations):
    # This is the number of features in the feature map
    n_features = layer_activation.shape[-1]

    # The feature map has shape (1, size, size, n_features)
    size = layer_activation.shape[1]

    # We will tile the activation channels in this matrix
    n_cols = n_features // images_per_row
    display_grid = np.zeros((size * n_cols, images_per_row * size))

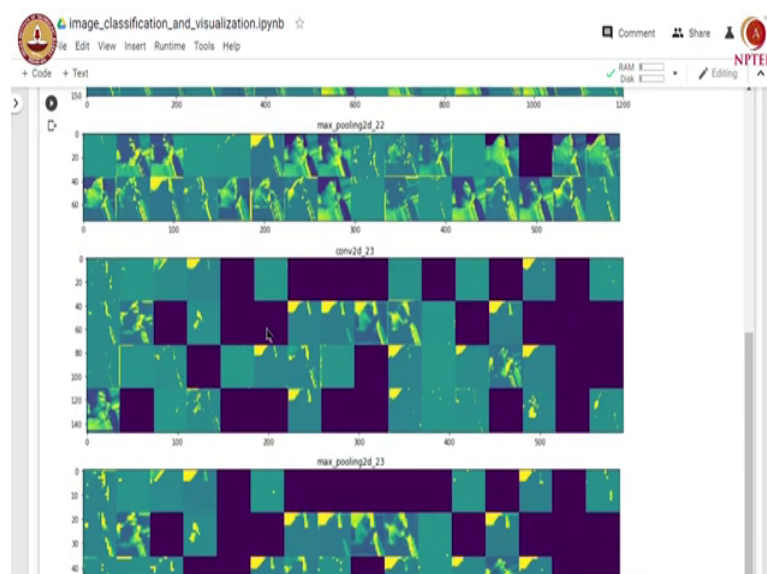
    # We'll tile each filter into this big horizontal grid
    for row in range(images_per_row):
        for col in range(n_cols):
            channel_image = layer_activation[0, :, :, col * images_per_row + row]

            # Post-process the feature to make it visually palatable
            channel_image -= channel_image.mean()
            channel_image /= channel_image.std()
            channel_image *= 64
            channel_image += 128
            channel_image = np.clip(channel_image, 0, 255).astype('uint8')
            display_grid[col * size : (col + 1) * size, row * size : (row + 1) * size] = channel_image

    # Display the grid
    scale = 1. / size
```

We will extract and plot every channel in each of our 8 activation maps and we will stack the results in one big image tensor with channels stacked side by side.

(Refer Slide Time: 20:06)



So, now you can see that the first convolution layer seems to be detecting edges and as we go deeper and deeper in the network. We are detecting different kind of features from the images. You can see that there are lot of empty spots after the second convolution and max pooling operation.

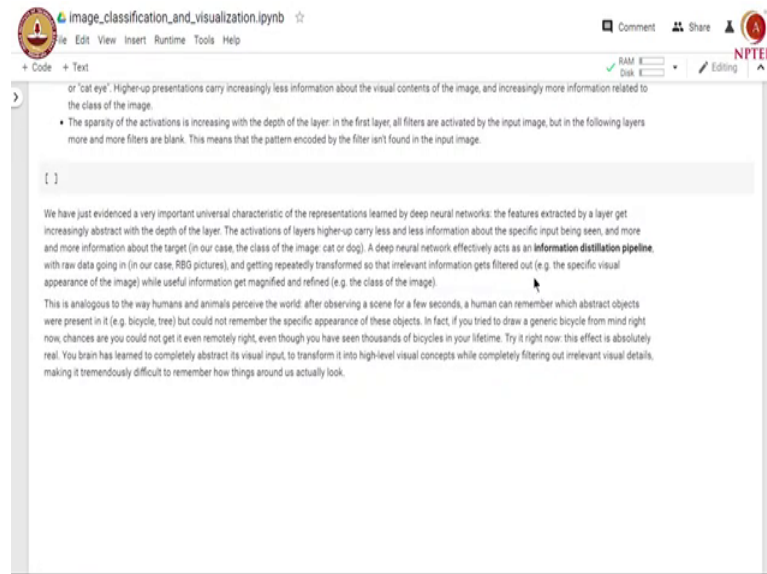
(Refer Slide Time: 20:38)



There are few remarkable things to note here. The first layer acts as a collection of various edge detectors. At that stage the activations are still returning almost all the information present in the initial picture. You can see that lot of features from the input are being preserved. As we go higher up the activations become increasingly abstract and less visually interpretable. They start encoding higher level concepts such as cat ear and cat eye higher of representations carry increasingly less information about a visual content of the image and increase in the more information related to the class of the image.

The sparsity of activations is increasing with the depth of the layer. The first layer almost all filters are activated by the input image, but as we go deeper and deeper in the network there are more and more blank filters. This means that patterns encoded by the filter is not found in the input image. So, this is a very nice way of visualizing how convolution convolutional neural network is learning patterns in the image.

(Refer Slide Time: 22:14)



So, we showed a very important universal characteristics of representations learned by deep neural network. The feature the features extracted by layer get increasingly abstract with depth of the layer. The activations of layers higher of carry less and less information about specific input being seen, but more and more information about the target class. A deep neural network effectively acts as an information distillation pipeline with raw data going in and getting repetitively transformed so, that irrelevant information gets filtered out and useful information gets magnified and refined.