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

Lecture – 24 Transfer Learning with pre-trained CNNs

In the last lecture, we looked at concepts behind convolution neural networks and use them for building an image classifier. We also learned that these models have a large number of parameters and in order to train them without overfitting we require a reasonably large amount of data. How do we use these models when we do not have a large amount of data? For example, say a manufacturing company wants to train a CNN model for recognizing faulty machine parts, and they do not have enough data, what can we do?

We learn techniques behind solving these kind of problems in this session. We will demonstrate these techniques by classifying cats and dogs.

(Refer Slide Time: 01:23)



A pre-trained model is at the centre of our strategy. A pre-trained model is a saved network that was previously trained on a large dataset, typically on a large scale image classification task. We either use the pre-trained model as it is or use transfer learning to customize this model for a given task. The intuition behind transfer learning is that if a model is trained on a large and general enough data set, this model will effectively serve as a general model of the visual world. We can take advantage of these learned feature maps without having to start training a large model from scratch on a large data set.

In this exercise we will try two ways to customize a pre-trained model.

(Refer Slide Time: 02:19)



CNNs are made up of one or more convolution and pooling layers, generally followed by dense layers. For example, we might have two convolution pooling layers followed by a convolution layer whose output is then fed into a dense layer to generate a label.

Now, the idea here is to use this CNN model which was trained on a large data set and use it for performing some other task. For example, we have a data set of machine parts and we want to recognize faulty machine parts. So, we want to build a CNN followed by feed forward neural network to get the label, which in this case is faulty or good. We know that a CNN model has a large number of parameters and if you do not have enough data points about machine parts and a label you are likely to overfit the CNN model.

Here, what we want to do is we want to take advantage of a pre-trained model which is trained on a different dataset from the dataset pertaining to our problem, and use it for solving the problem at hand. There are really two ways in which we can achieve this. (Refer Slide Time: 05:29)



Let us quickly introduce some generic structure to our CNNs. We have the convolutional neural network (CNN) part, which consists of convolutional layers and pooling layers. The other part, referred to as feed forward neural network (FFNN) consists of dense layers.

In the first of the two ways we can use our pretrained model, we will only use the convolutional part of the model. We use the convolutional part of the model for feature generation and pass the output of the CNN to a new network of dense layers. With the new data, we only train these new layers that were added and not the CNN. The weights of the CNN remain the same as the CNN part of the pretrained model.

In the second approach what we do is, we use a certain section of the CNN as it is in the pretrained model and retrain the rest of the network. The certain section is typically the first few layers of the network (refer to the image above). The training is done by freezing the weights in the earlier layers and tuning the rest of the weights using the training data at hand.

Now that you have understood how to use pre-trained model for building a custom model, let us look at the machine learning workflow involved in this particular process.

(Refer Slide Time: 12:05)

+ Code	+ Text & Copy to Drive 2. butto an input pipeline, in this case using Keras Imageuatavenerator 2. Compare www.model	··· Connecting 👻	🇨 Editing 🗸 🗸
5	 Load in our pretrained base model (and pretrained weights) Stack our classification layers on top 		
	4. Train our model 5. Evaluate model		
		$\uparrow \downarrow$	c) 🗘 🔋 :
0	fromfuture import absolute_import, division, print_function, unicode_literals		
	import os		
	import numpy as np		
	import matplotlib.pyplot as plt		
Q	jpip install tansorflow-gpu=>2.0.0-betal input: tensorflow, as tf of the difference on constant in a news Reference of the reference		
+ Da	ata preprocessing		
- Da	winload		
NP Us	TEL TensorFlow Datasets to load the cats and dogs dataset.		

So, we will first examine and understand the data which is exactly the same as traditional machine learning algorithms, then we build an input pipeline, then we build and compile our model. Only the model composition differs from the traditional machine learning algorithms.

In traditional machine learning algorithms, we define the model completely. In this case, we will have to load the pre-trained model and then add the classifier layer (as a dense layer) on top of it. Then, the remaining two steps (training and evaluation) are again very similar to our traditional machine learning algorithms.

So, let us start by importing all the necessary libraries. Let us also install and import the tensorflow package.

(Refer Slide Time: 13:29)



Let us load the data set using tfds package.

(Refer Slide Time: 13:34)

The	resulting tf.data.Dataset objects contain (image, label) pairs. Where the images have variable shape and 3 ch	annel	s, an	nd the	labe	is a	
SCal	al,	\uparrow	\downarrow	Θ	\$	Î	:
0	print(rew_train) print(rew_taildation) print(rew_test)						
θ	<pre><_OptionsDataset shapes: ((<u>None, None, 3</u>), ()), types: (tf.uint8, tf.int64)> <_OptionsDataset shapes: ((None, None, 3), ()), types: (tf.uint8, tf.int64)> <_OptionsDataset shapes: ((None, None, 3), ()), types: (tf.uint8, tf.int64)></pre>						
Sho	w the first two images and labels from the training set:						
[]	<pre>get_label_name =_imetadata.features['label'].int2str</pre>						
	<pre>for image, label in raw_train.take(2): plt.figure() plt.inshow(image) plt.title(get_label_name(label))</pre>						

The tfds.load method downloads and caches data and returns a tf.data.Dataset object. These objects provide powerful and efficient methods for manipulating data and piping it into our model. Since cats and dogs data set does not define standard split we use the subsplit function to divide it into train, validation and test with the split specified by the weighted parameter.

Here we use 80 percent data for training, 10 percent data for test and the remaining 10 percent data for validation.

The resulting dataset object contains images and label pairs. The images have variable shape and 3 channels, and label is a scalar quantity. Let us look at the first two images and add labels from the training set. So, iterate on the training set using take method on the tensor and we generate the string of the label using into string property.

(Refer Slide Time: 14:59)



You can see that this is a picture of a dog with label displayed at the top and there is a picture of the cat. You can see that the dog picture has height of 500 and width of over 350; whereas, the cat picture has height close to 400 and width of 500. So, you can see that all the pictures are not of the same size.

(Refer Slide Time: 15:28)



So, the first thing is to resize the image, so that we have the same input size for all the images. We will be using tf.image module to format the images for this task. We will also rescale the input channel to a range of minus 1 to plus 1. Here the desired size of image is 160 by 160. We will apply the function on each item in the dataset using the map method.

So, you can see that the format_example is applied on the training set, validation set and test set. We get three tensors - train, validation and test - that contains images of the same size and their labels.

(Refer Slide Time: 16:39)

ode ·	Text 🔹 Copy to Drive	Disk Editing
[9]	BATCH_SIZE = 32 SHUFFLE_BUFFER_SIZE = 1000	
[10]	train_batches = train.shuffle(SHUFFLE_BUFFER_SIZE).batch(BATCH_SIZE) validation_batches = validation.batch(BATCH_SIZE) test_batches = test.batch(BATCH_SIZE)	
Inspe	it a batch of data:	
0	for image_batch, label_batch in train_batches.take(1): pass	↑↓☺✿▮:
	image_batch.shape	
0	FensorShape([\$2, 160, 160, 3])	
Cre You v 1.4M syrin	It e the base model from the pre-trained convnets II create the base model from the MobileNet V2 model developed at Google. This is pre-trained on the Ima mages and 1000 classes of web images. ImageNet has a fairly arbitrary research training dataset with cat ge, but this base of knowledge will help us tell apart cats and dogs from our specific dataset. where to pick which layer of MobileNet V2 you will use for feature extraction. Obviously, the very last cla is of machine learning models go from bottom to top) is not very useful. Instead, you will follow the comm we take the fatter oncreation. This layer is called the "brotheneck layer". The bottheneck fatteres	ageNet dataset, a large dataset of tegories like jackfruit and ssification layer (on "top", as most mon practice to instead depend on

Let us shuffle and batch the data. We use batch size of 32, and for shuffling we define a buffer size of 1000. We shuffle only the training data and we batch all the three data sets by the batch size of 32. Let us inspect a batch of data from training batches. So, we can see that in the first training batch we have 32 images, each with height of 160 and width of 160 on the 3 channels. So, you can see that the image batch here is a 4D tensor.

Now that we have got in the data in the desired shape let us create the model. In model creation, there are two steps; first is to load the base convolution model and second, o add a classification layer on top of it.

Here we will create the base model from MobileNet version 2 developed at Google. This model was pre-trained on image data set which is a large data set of 1.4 million images from thousand classes of web images. ImageNet has categories like jack fruit and syringe, but we will use the image net classified here to classify cats versus dogs.

First you need to pick which layer of MobileNet you will use for feature extraction; obviously, very last classification layer is not going to be very useful for this task. Instead we will follow a common practice of extracting features at a layer just before the flatten operation. This layer is referred to as the bottleneck layer. A bottleneck feature retains generality as compared to final or top layer.

(Refer Slide Time: 18:53)



So, let us first instantiate MobileNet with a pre-loaded weights trained on ImageNet. We can do that using tf.keras.applications.MobileNetV2 function. We specify the input shape. We tell the model that we do not want to include the top layer or the classifier layer by specifying include_top argument to false and we specify that we want to use weights of MobileNetV2 when trained on the ImageNet dataset.

Since we specify include_top as false the network does not include the classification layer at the top which is ideal for feature extraction. The feature extractor converts each of 160x160x3 image into a 5x5x1280 block of features. Let us look at how this model looks like.

(Refer Slide Time: 19:58)

de	+ Text 🔺 Copy to Drive		V RAM I	▪ ✓ Editing
[12]] IMG_SHAPE = (IMG_SIZE, IMG_SIZE,	3)		
	<pre># Create the base model from the base_model = tf.keras.applicatio</pre>	<pre>pre-trained model MobileNet V2 ns.MobileNetV2(input_shape=IMG_SHA</pre>	PE,	
0	Downloading data from https:// 9412608/9406464 [=============	github.com/JonathanCMitchell/mol ======] - 0s 0us/step	pilenet_v2_keras/releases/download/v p	1.1/mobilenet_v2
	a.			
This	s feature extractor converts each 160x16	50x3 image to a 5x5x1280 block of fea	tures. See what it does to the example batch	of images:
				ψ 🙃 📅 🔒
0	base_model.summary()			v
-	N. J. J. H. L.J			
Θ	Model: "mobilenetv2_1.00_160"			
	Layer (type)	Output Shape Param #	Connected to	-
	input_1 (InputLayer)	[(None, 160, 160, 3) 0		=
	Conv1_pad (ZeroPadding2D)	(None, 161, 161, 3) 0	input_1[0][0]	-
an.	Conv1 (Conv2D)	(None, 80, 80, 32) 864	Conv1_pad[0][0]	_
(¥	Conv1 (BatchNormalization)	(None, 80, 80, 32) 128	Conv1[0][0]	-
NPT	Elonv1_relu (ReLU)	(None, 80, 80, 32) 0	bn_Conv1[0][0]	-

So, when we call summary on the model you get to see the complete architecture of the MobileNet V2. We can see that it takes a 4D tensor, where there are images of size 160x160 across 3 channels; that means, it takes coloured images of size 160x160.

(Refer Slide Time: 20:28)

2	block_16_expand (Conv2D)	(None,	5,	5,	960)	153600	block_15_add[0][0]	
)	block_16_expand_BN (BatchNormal	(None,	5,	5,	960)	3840	block_16_expand[0][0]	
	block_16_expand_relu (ReLU)	(None,	5,	5,	960)	0	<pre>block_16_expand_BN[0][0]</pre>	
	block_16_depthwise (DepthwiseCo	(None,	5,	5,	960)	8640	<pre>block_16_expand_relu[0][0]</pre>	
	block_16_depthwise_BN (BatchNor	(None,	5,	5,	960)	3840	<pre>block_16_depthwise[0][0]</pre>	
	block_16_depthwise_relu (ReLU)	(None,	5,	5,	960)	0	<pre>block_16_depthwise_BN[0][0]</pre>	
	block_16_project (Conv2D)	(None,	5,	5,	320)	307200	<pre>block_16_depthwise_relu[0][0]</pre>	
	block_16_project_BN (BatchNorma	(None,	5,	5,	320)	1280	<pre>block_16_project[0][0]</pre>	
	Conv_1 (Conv2D)	(None,	5,	5,	1280)	409600	<pre>block_16_project_BN[0][0]</pre>	
	Conv_1_bn (BatchNormalization)	(None,	5,	5,	1280)	5120	Conv_1[0][0]	
	out_relu (ReLU)	(None,	5,	5,	1280)	0	Conv_1_bn[0][0]	
	Total params: 2,257,984 Trainable params: 2,223,872 Non-trainable params: 34,112							

And, its final layer produces a 4D tensor, where we get 5x5 patches across 1280 channels. We also see the total number of parameters for this model. So, this model has got 2.2 million parameters.

(Refer Slide Time: 20:59)

	+ Text & Copy to Drive		*	1	βE	diting	g
[12]	Total params: 2,257,984						
[12]	Trainable params: 2,223,872						
0	Non-trainable params: 34,112						
		\uparrow	4	Ð	ф	Î	:
0	<pre>feature_batch = base_model(image_batch)</pre>	_			_		
	print(feature_batch.shape)						
0	(32, 5, 5, 1288)						
Fea You v level	ature extraction will freeze the convolutional base created from the previous step and use that as a feature extractor, add a classifier classifier.	on to	p of it	t and	l trai	n the	to
You v level	ature extraction will freeze the convolutional base created from the previous step and use that as a feature extractor, add a classifier classifier. b: ize the convolutional base	r on toj	p of it	t anc	l trai	n the	e to
Fee You v level Free It's in preve	ature extraction will freeze the convolutional base created from the previous step and use that as a feature extractor, add a classifier. U I I I I I I I I I I I I I I I I I I	on to ainat	p of i)le :)del's	t anc : Fa traii	l trai 1se) nable	n the I, you e flag	e to

Let us use the base model to generate the features.

So, you can see that on the image batch that we selected we computed the features for the image batch. You can see that for each of the 32 examples in the batch, we got a 3D tensor of size of shape (5, 5, 1280). We will freeze the convolution base created from the previous step and use that as a feature extractor. We add a classifier on top of it and train the top level classifier.

(Refer Slide Time: 21:48)



Let us see how to do that in the code. We use base_model.trainable attribute or property and set it to false. This makes sure that you freeze the convolution base before we compile and train the model. By freezing you prevent the weights in a given layer from being updated during the training. MobileNet has many layers, so, setting the entire models trainable flat to false will freeze all the layers.

Let us look at how the base model looks like.

(Refer Slide Time: 22:24)

-	+ Text 🍐 Copy to Drive						V RAM Disk	 Editing
0	block_16_project (Conv2D)	(None,	5,	5,	320)	307200	<pre>block_16_depthwise_relu[0][0]</pre>	
0	block_16_project_BN (BatchNorma	(None,	5,	5,	320)	1280	<pre>block_16_project[0][0]</pre>	
	Conv_1 (Conv2D)	(None,	5,	5,	1280)	409600	<pre>block_16_project_BN[0][0]</pre>	
	Conv_1_bn (BatchNormalization)	(None,	5,	5,	1280)	5120	Conv_1[0][0]	
	out_relu (ReLU)	(None,	5,	5,	1280)	0	Conv_1_bn[0][0]	
	Non-trainable params: 34,112 b			_				
[14]	Non-trainable params: 34,112 b	batch)						v
[14]	<pre>Non-trainable params: 34,112 b feature_batch = base_model(image_b print(feature_batch.shape)</pre>	batch)						•
[14]	<pre>Konstrainable parans: 24,112 b feature_batch = base_model(image_b print(feature_batch.shape) (32, 5, 5, 1288)</pre>	batch)						
[14] e	<pre>konstrainable parans: 24,122 b feature_batch = base_model(image_t print(feature_batch.shape) (32, 5, 5, 1280) ture extraction</pre>	batch)						
[14] • Fea	<pre>konstrainable parans: 24,122 b feature_batch = base_model(image_b print(feature_batch.shape) (32, 5, 5, 1288) (32, 5, 1288) (32, 5, 1288) (33, 5, 1288) (34, 12, 12, 12, 12, 12, 12, 12, 12, 12, 12</pre>	batch)	e pre	evic	bus step a	and use that	as a feature extractor, add a classifier on top c	f it and train the top-

I would like you to compare the trainable parameters after freezing. You can see that after freezing the trainable parameters become 0; that means we do not have to train any of the parameter of this network and all the parameter all the 2.2 million parameters become non-trainable.

(Refer Slide Time: 22:52)



In order to generate predictions from the block of the feature we average the special 5x5 block using a GlobalAveragePooling layer to convert the feature to a single vector of 1280

elements per image. So, we add the GlobalAveragePooling layer to the model. Let us look at the shape of the feature batch. You can see that each of the image each of the 32 images in the batch got converted into a 1D tensor containing 1280 numbers.

Finally we apply tf.keras.layers.dense layer to convert these features into a single prediction per image. We do not need an activation function here because this prediction will be treated as a raw prediction value. A positive number predicts class 1 and negative number predicts class 0.

(Refer Slide Time: 24:03)



So, let us stack the feature extractor and two layers using a tf.keras.Sequential model. So, we define the model to be tf.keras.Sequential model which has got a base model, which itself is a sequence of convolution and pooling layers according to the MobileNet architecture. You have a global average layer followed by a prediction layer. Prediction layer is a dense layer having a single unit. Let us compile the model before training it.

Since there are two classes we are using binary cross entropy loss and we are using RMSprop optimizer with a small learning rate.

(Refer Slide Time: 24:52)



Let us look at the summary of the model. You can see that the MobileNetV2 returns the 4D tensor of shape (None, 5, 5, 1280) and it has got 2.2 million parameters. The global average pooling returns 1280 numbers per image and finally, we have a dense layer with a single unit.

We can see that this dense layer receive 1280 inputs plus 1 bias unit makes 1281 parameters in the dense layer. So, you can see that the total number of parameters are more than 2.2 million, out of which most of the parameters are non-trainable and we have to only train 1281 parameters corresponding to the output layer.

(Refer Slide Time: 25:44)

▼ Trai	n the model	
After	training for 10 epochs, you should see ~96% accuracy.	
[] 4	<pre>num_train, num yal, num_test = (metadata.splits['train'].num_examples*weight/10 for weight in SPLIT_WEIGHTS)</pre>	
[]	<pre>initial_epochs = 10 steps_per_epoch = round(num_train)//BATCH_SIZE validation_steps = 20 loss0,accuracy0 = model.evaluate(validation_batches, steps = validation_steps)</pre>	
[]	<pre>nrint/"initial loss: (: 261% format/loss@))</pre>	
	<pre>print("initial accuracy: {:.2f}".format(accuracy0))</pre>	
[]	history = model.fit(train_batches, epochs=initial_epochs, validation_data=validation_batches)	
- Lea	rning curves	
Lette	take a look at the learning curves of the training and validation accuracy/loss when using the Mo	bileNet V2 base model as a fixed feature

Let us train the model for 10 epochs.

(Refer Slide Time: 25:47)

+ Code	+ Text 🏾 💩 Copy to Drive		V RA Dis	sk 🗖 🔹	🖍 Editing 🔷 🗸
[27]	Epoch 9/10 582/582 [======] - Epoch 10/10 582/582 [======] -	75s 129ms/step - loss: 0 74s 127ms/step - loss: 0	.4289 - accuracy: 0.94 .4107 - accuracy: 0.94	428 - val_loss: 440 - val_loss:	0.3677 - va: 0.3425 - va:
		CODE TEXT			,
	rning curves				
Let's extra	take a look at the learning curves of the training and v actor.	alidation accuracy/loss when u	ising the MobileNet V2 bas	e model as a fixed	feature
[28]	<pre>acc = history.history['accuracy'] val_acc = history.history['val_accuracy']</pre>				
	<pre>loss = history.history['loss'] val_loss = history.history['val_loss']</pre>				
	<pre>plt.figure(figsize=(8, 8)) plt.subplot(2, 1, 1) plt.plot(acc, label='Training Accuracy') plt.plot(val_acc, label='Validation Accuracy' plt.legend(loc='lower right') plt.vlot('iseruscu'')</pre>				
	<pre>plt.ylade(Accuracy) plt.ylim([min(plt.ylim()),1]) plt.title('Training and Validation Accuracy')</pre>				
(*	<pre>plt.subplot(2, 1, 2) plt.plot(loss, label='Training Loss') plt.plot(val_loss, label='Validation Loss') plt.loged(logs'upper picht')</pre>				
NPTI	<pre>plt.ylabel('Cross Entropy') plt.ylim([0,1.0]) plt.tild('Training and Validation Loss')</pre>				

You can see that after 10th epoch we cross the accuracy of 94 percent. Let us take a look at the learning curves of the training and validation accuracy or loss when using MobileNet base model as the fixed feature extractor.

(Refer Slide Time: 26:10)



You can see that training and validation loss and accuracies are quite close by after the 10th epoch. You must be wondering why validation matrices are better than the training matrices. The main factor is because layers like tf.keras.layers.BatchNormalization and dropout affect accuracy during the training. They are turned off when calculating the validation loss. To a lesser extent it is also because training matrices report average for an epoch while validation matrix are evaluated after every epoch. So, validation matrix see a model that has trained slightly longer.

(Refer Slide Time: 26:50)



In our feature extraction experiment, we were only training a few layers on top of MobileNet base model. The weights of the pre-trained network were not updated during the training. The weights were frozen for the base model. One idea to improve the performance can be to fine tune the weights of the top layer of the pre-trained model alongside the classifier layer.

So, now, what we are going to do here is we are going to we are going to unfreeze these layers of the base network and train them along with the classifier layer. Let us see how to specify unfreezing of these layers in the code.

(Refer Slide Time: 28:00)



So, we set the trainable property of the base model to true and what we do is we specify the layers after which you want to unfreeze the network. So, here you want to find tune from 100th layer onwards. So, we specify the fine_tune_at variable and anything after that will be trained and anything before this fine_tune_at will be fixed. So, we set the trainable value of the layers before 100th layer as false. So, there are 155 layers in the base model out of which we retrain the top 55 layers.

It is important to note that as we go deeper in the network, network become more specialized to patterns learned from the training data. The initial convolution layer learns general patterns or simpler patterns like edges or corners, but as we go deeper and deeper the models become specialized to recognize patterns from the training data.

So, if we are using CNNs as the pre-trained model for some other task it is important to cut it somewhere in the middle if the task is very very different from the training data. For example, if the training data contains pictures of cats and dogs whereas, that the training data for the new task is about machine parts we might have to start unfreezing quite early in the network, because the original network which was trained on cats and dogs will start recognizing patterns related to cats and dog as it goes deeper.

(Refer Slide Time: 30:20)

		•	JL G		:
D	<pre># Let's take a look to see how many layers are in the base model print("Number of layers in the base model: ", len(base_model.layers))</pre>	1	V C	 Ì	
	<pre># Fine tune from this layer onwards fine_tune_at = 100</pre>				
	<pre># Freeze all the layers before the `fine tune at` layer for layer in base_model.layers[:fine_tune_at]: layer.trainable = False</pre>				
9	Number of layers in the base model: 155				
om	pile the model pile the model using a much lower training rate.				
Com Com 31]	<pre>pile the model pile the model using a much lower training rate. model.compile(loss='binary_crossentropy',</pre>				
Com Com 31] 32]	<pre>pile the model pile the model using a much lower training rate. model.compile(loss-'binary_crossentropy',</pre>				
com com 31] 32]	<pre>pile the model pile the model using a much lower training rate. model.compile(loss='binary_crossentropy',</pre>				

So, let us use the same setting we use binary towards entropy loss and we use RMSprop as an optimizer. Let us look at the model summary after unfreezing some layers of the network.

(Refer Slide Time: 30:35)

COILI	prie trie mouer usilig a much lowe	r training rate.			
[31]	<pre>model.compile(loss='binary_ optimizer = tr metrics=['accu</pre>	<pre>crossentropy', f.keras.optimizers.RMSp uracy'])</pre>	rop(lr=base_learning_rate/10),	
[32]	model.summary()				
0	Model: "sequential"				
	Layer (type)	Output Shape	Param #		
	mobilenetv2_1.00_160 (Mode	1) (None, 5, 5, 1280)	2257984		
	global_average_pooling2d (Gl (None, 1280)	0		
	dense (Dense) Total params: 2,259,265 Trainable params: 1,863,87 Non-trainable params: 395,	(None, 1) 3 392	1281		
*	len(model.trainable_variable	25)			

Now, you can see that out of 2.2 million total parameters, the non-trainable parameters are reduced to 395000 and the rest of the parameters are trainable parameters. Let us compare this model summary with the model summary from the previous exercise.

So, you can see that in the previous exercise we only had 1281 trainable parameters and all of the 2.2 million parameters of the base convolution neural network were non-trainable, but since we have unfrozen some of the layers, the number of parameters have gone up.

We should note that we should only attempt fine tuning after training the top level classifier with the pre-trained model set to non-trainable. If you add a randomly initialized classifier on the top of the pre-trained model and attempt to train all the layers jointly, the magnitude of the gradient updates will be too large and our pre-trained model will forget what it has learned.

(Refer Slide Time: 31:55)

+ Code + Text 🕼 Copy to Drive		V RAM Disk	🔹 🧨 Editing 🗸 🗸
<pre>> [33] len(model.trainable_variables)</pre>			
9 58			
▼ Continue Train the model			
If you trained to convergence earlier, this will get yo	you a few percent more accuracy.	1	\ ↓ @ ☆ î :
<pre>fine_tune_epochs = 10 total_epochs = initial_epochs + fine_</pre>	_tune_epochs		
history_time = model.tit(train_batches epochs=total_ initial_epoch validation_da	rs, _epochs, h = <u>initial_epochs</u> , lata=validation_batches)		
Epoch 11/20 582/582 [====================================	=====] - 105s 181ms/step - loss: 0.3738	- accuracy: 0.9546 - v	/al_loss: 0.0000e+00
582/582 [====================================	=====] - 87s 150ms/step - loss: 0.2919	- accuracy: 0.9767 - va	al_loss: 0.1964 - va:
582/582 [====================================	=====] - 86s 148ms/step - loss: 0.2542 =====] - 86s 147ms/step - loss: 0.2381	- accuracy: 0.9821 - va - accuracy: 0.9839 - va	al_loss: 0.1745 - va: al_loss: 0.1871 - va:
PTE 82/582 [Epoch 16/20	=====] - 85s 146ms/step - loss: 0.2267	- accuracy: 0.9846 - va	al_loss: 0.1822 - va:

Let us train the model for 10 epochs. What we do is we start the training of the model where we had stopped in the previous run. So, we had stopped on 10th epoch. So, we want to initialize the model with the 10th epoch and resume the training from that. So, you can see that when we run this particular fit function or the model training we started training from 11th epoch.

(Refer Slide Time: 32:26)

 58 Ep 58 Ep 	!/582 [====================================] -	0.0-									
Ep 58 Ep	ch 14/20		865	148ms/step	- loss:	0.2542	- accur	acy: 0.9	821 -	val_loss:	0.1745	- va
58 Ep	1 000											
Ep	/ 582 [=======================] -	86s	147ms/step	- loss:	0.2381	- accur	acy: 0.9	839 -	val_loss:	0.1871	- va
	ch 15/20											
58	./582 [====================================] -	85s	146ms/step	- loss:	0.2267	- accur	acy: 0.9	846 -	val_loss:	0.1822	- va
Ep	ch 16/20											
58	/582 [====================================		85s	146ms/step	- loss:	0.2195	- accur	acy: 0.9	854 -	val_loss:	0.2099	- va
Ep	ch 17/20	1	05.	146-00/0400	1	0 2122			000		0 1000	
50	/582 [====================================		855	146ms/scep	- 1055:	0.2125	- accur	acy: 0.9	858 -	val_ioss:	0.1900	- va
LP EQ	CN 18/20		95c	1/6mc/cton	- 1055	0 2001	- 2000	acu: 0 0	963 -	val locc:	0 2051	- 1/2
En	./362 [=======		000	1401157 5 cep	- 1055.	0.2034	- decui	dcy: 0.5	005 -	Vd1_1055.	0.2001	- va
58	//582 [====================================	- [85s	146ms/step	- loss:	9,2077	- accur	ary: 0.9	864 -	val loss:	9,2071	- va
Ep	och 20/20		0.55	1401101 0 000		01201.		ucj. 0.1	001	vur_10011	0.2072	
58	2/582 [====================================] -	85s	146ms/step	- loss:	0.2020	- accur	acy: 0.9	865 -	val loss:	0.2033	- va
				2 4						-		
1		_										,

Now, the accuracy has gone up to 99 percent; it has crossed 98 percent which is more than 4 percentage points from the previous exercise where we used MobileNet merely as a feature extraction mechanism.

So, in this module we learned how to use a pre-trained machine learning model as a feature extractor and we also learned how to fine tune a pre-trained model. Hope you had fun learning these concepts. In the next session, we will use TensorFlow hub for loading some of the pre-trained models and using them for feature extraction as well as for fine tuning.