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

Lecture – 13 Building Data Pipelines for Tensor flow-Part 2

[FL] In this session, we will study how to construct dataset from different formats. We will cover formats like CSV or in memory structures from NumPy and Pandas we will also study how to construct dataset objects for images and text.

(Refer Slide Time: 00:35)



So, let us start with CSV data. So, the data that we use in this session is taken from Titanic passenger list where you want to build a model to predict the likelihood of a passenger survival based on their characteristics like age, gender, the ticket class, whether the passenger was travelling alone etc.

(Refer Slide Time: 01:17)



Install Tensor Flow 2.0 and import essential libraries.

(Refer Slide Time: 01:43)



Then, download the train and test data from titanic dataset. The titanic dataset is available in *tf.dataset*. We set the precision to 3.

(Refer Slide Time: 01:59)



Let us look at first few examples from CSV file. So, you can see that the label survived is the first column it is either 0 or 1. Then, we have features or characteristics like gender or sex, age, number of siblings or spouses that were travelling together, parch, fare, class, deck, embark_town and whether the traveler was alone.

We loaded the data using Pandas and pass the NumPy arrays to Tensor Flow. If you want to scale if you want to scale up to a large set of files then we will have to use make_csv_dataset() function which is currently in the experimental version of TensorFlow.

(Refer Slide Time: 03:13)



We separate the label column and as we said there are two labels 0 or 1. Let us read the file and create a dataset. Here we demonstrate the make_csv_dataset() function which can be used for large datasets the current dataset is pretty small, but still we decide to demonstrate make_csv_dataset() function with Titanic dataset because this will be applicable for large CSV files as well. Here we are taking a small batch size internationally, so that we can examine the examples in each batch. We specify the label name, specify how null values are predicted how null values are specified with question mark here in this case and number of epochs we set it to 1.

We construct the training data by giving the training file to get_dataset() function and test data by providing the path of the test file to the get_dataset()function. The get_dataset() function makes a dataset from CSV file and returns the dataset.

(Refer Slide Time: 05:29)



Each item in the dataset is a batch represented as a tuple of examples and labels. The data from the examples is organized in column based tensors each with as many elements as the batch size.

(Refer Slide Time: 06:07)

ode	+ Text 🍐 Copy to l	Drive ✓ RAM ⊥.interieave(map_tunc, cycie_iengtn, piock_iengtn, num_paraiiei_	_calls=tt.da	Editing ta.experimer
0	4			
[8]	<pre>def show_batch(datas for batch, label i for key, value i print("{:20s}:</pre>	et): n dataset.take(1): n batch.items(): {}".format(key,value.numpy()))		
Each	n item in the dataset is a l	patch, represented as a tuple of (many examples, many labels). The data from th	ie examples is	organized in
colu It mi	mn-based tensors (rather ight help to see this yours	than row-based tensors), each with as many elements as the batch size (12 in elf.	this case).	
colu It mi	mn-based tensors (rather ght help to see this yours show_batch(raw_train	than row-based tensors), each with as many elements as the batch size (12 in telf. _data)	this case). ↑ ↓ ⊂	ə 🌣 🧻 :
colu It mi	mn-based tensors (rather ght help to see this yours show_batch(raw_train sex age n_siblings_spouses parch fare class deck embark_town bisu_thematon 1	<pre>than row-based tensors), each with as many elements as the batch size (12 in elf. </pre>	this case).	a 🌣 🏚 🔋 :

Let us look at one of the batch. You can see that since you set out the batch size to 5 for each of the feature you have exactly 5 values. Each feature is organized as a column based tensor. So, you can see that there is one tenser for feature, one tenser for sex, another for age and so on, and in each of the tensor you will find exactly five values.



(Refer Slide Time: 06:43)

We can specify the columns that we want to use to construct a dataset and pass that in the *select_columns* argument of the get_dataset() function that we wrote earlier.

(Refer Slide Time: 07:09)



And, since we choose only hand full of the features we can see that there are only those many features that were selected are shown here.

(Refer Slide Time: 07:31)

and the second s	1	en fala least 1 a geografication y x + + - a geografication y x + - a
>	-to	de + Text
	•	Data preprocessing
		A CSV file can contain a variety of data types. Typically you want to convert from those mixed types to a fixed length vector before feeding the data into your model.
		TensorFlow has a built-in system for describing common input conversions: tf.feature_column, see this tutorial for details.
		You can preprocess your data using any tool you like (like <u>nltk</u> or <u>sklearn</u>), and just pass the processed output to TensorFlow.
		The primary advantage of doing the preprocessing inside your model is that when you export the model it includes the preprocessing. This way you can pass the raw data directly to your model.
	•	Continuous data
		If your data is already in an apropriate numeric format, you can pack the data into a vector before passing it off to the model:
		<pre>[] SELECT_COLUMNS = ['survived', 'age', 'n_siblings_spouses', 'parch', 'fare'] DEFAULTS = [0, 0.0, 0.0, 0.0, 0.0] temp dataset = get dataset(train file path,</pre>
ŧ	Ħ	C 🛤 🕭 🖉 🖉 🖉 🖉 🖉 🖉

A CSV file can contain a variety of data types and typically you want to convert the mix type to a fix length vector before feeding the data into the model. So, TensorFlow uses a construct called *tf.feature column* for such kind of convergence.

(Refer Slide Time: 08:01)

ode	+ Text 🍐 Copy to	Drive	ic format	VOLL COD DO	ick the d	ata into a i	vector hefo	RAM Disk Disk	off to the	▼ del	/ Ed	diting	NI	' Ţ E
ii yoo	in data io aneday in an a		io ronnut,	you oun po		ata into a		e pubbling it	on to th	, mouci				
[12]	SELECT_COLUMNS = [': DEFAULTS = [0, 0.0, temp_dataset = get_c show_batch(temp_data	survived', 'ag 0.0, 0.0, 0.0 dataset(train_ select column aset)	e', 'n_si] file_path _columns= _defaults	<mark>iblings_sp</mark> n, =SELECT_CO s = DEFAUL	ouses', LUMNS, TS)	'parch',	'fare']							
0	age n_siblings_spouses parch fare	: [19. 52. : [0. 0. 4. : [0. 0. 1. : [0.	4. 28. 0. 1.] 0. 0.] 30.5	17.] 29.125	7.75	108.9]								
[13]	example_batch, labe	ls_batch = nex	t(iter(te	emp_datase	t))									
Here	s a simple function that	will pack togeth	er all the	columns:										
0	<pre>def pack(features, 1 return tf.stack(1)</pre>	label): ist(features.v	alues()),	, axis=-1)	, label				1	↓ e	\$	Î	:	

If a data is already in a numeric format we can pack the data into a vector before passing it out of the model. So, we select columns which are numeric is specify by the default values and we batch the dataset. We can look at example batch and the label batch. We will pack together all the columns and we apply this to each element of the dataset. We apply the pack() function to each element in the dataset.

		a oopy to	Diric					♥ Di	sk 📖		unung
0	packed_dat	aset = ter	np_data	set.map(pack)					<u>↑</u> \	00	
	for featur print(fe print() print(la	es, label: atures.num bels.numpy	s in pa npy()) /())	cked_dataset.t	ake(1):						
θ	[[19. [52. [4. [28. [17. [0 1 0 0 1	0. 0. 4. 0. 1.	0. 0. 1. 0.	0.] 30.5] 29.125] 7.75] 108.9]]							
If you this i	ı have mixed ncurs some c	datatypes ; werhead ar	you may nd shoul	want to separat d be avoided un	e out these less really n	simple-numeri ecessary. Switc	c fields. The t ch back to the	f.feature_ mixed datas	column api c set:	an handle	them, b
[]	show_batch	(raw_trai	n_data)								
	averal a ba	tch laha	le hate	h = nevt(iten)	temn datası	et))					

(Refer Slide Time: 08:59)

We can see that we have a feature metric or a 2D tensor of features and a 1D tensor of labels. If we have mixed data types we may want to separate out this simple numeric features. The *tf.feature_column* API can handle them, but this incurs some overhead and should be avoided unless really necessary.

(Refer Slide Time: 09:35)



(Refer Slide Time: 09:43)



Let us define a more general preprocessor that selects a list of numeric features and packs them into a single column. We take the numeric features and pack them.

(Refer Slide Time: 10:11)

Juc	Disk				anny	9
[19]	NUMERIC_FEATURES = ['age','n_siblings_spouses','parch', 'fare']					
	<pre>packed_train_data = raw_train_data.map(PackNumericFeatures(NUMERIC_FEATURES))</pre>					
	<pre>packed_test_data = raw_test_data.map(PackNumericFeatures(NUMERIC_FEATURES))</pre>					
		\uparrow	\downarrow	Θ	\$ î.	÷
0	show_batch(packed_train_data)	_				_
8	sex : [b'male' b'female' b'male' b'female' b'male'] class : [b'first' b'Second' b'Third' b'First' b'Third'] deck : [b'unknown' b'unknown' b'unknown'] embark_town : [b'Southampton' b'Southampton' b'Southampton' b'Southampton'] alonel : [b'y' b'n' b'n' b'n'] numeric : [[28. 0. 0. 26.55]] [2. 4. 1. 39.688] [45. 1. 1. 164.867] code : [code					

You can see that all the numeric features are packed into a tensor, whereas the other features are kept separate.

(Refer Slide Time: 10:37)



We need to normalize the continuous data.

(Refer Slide Time: 10:43)

						$\uparrow \downarrow$	e û î
0	<pre>import desc = desc</pre>	pandas as pd pd.read_csv(train_file_path)[NUME	RIC_FEATURES]	.describe()		
0		age	n_siblings_spouses	parch	fare		
	count	627.000000	627.000000	627.000000	627.000000		
	mean	29.631308	0.545455	0.379585	34.385399		
	std	12.511818	1.151090	0.792999	54.597730		
	• min	0.750000	0.000000	0.000000	0.000000		
	25%	23.000000	0.000000	0.000000	7.895800		
	50%	28.000000	0.000000	0.000000	15.045800		
	75%	35.000000	1.000000	0.000000	31.387500		
	max	80.000000	8.000000	5.000000	512.329200		

One of the ways of normalizing is by using z-score normalization that we studied earlier. Let us see how to do the z-score normalization on the numerical data or on continuous data. So, if first give the mean and standard deviation of each continuous column and then define a normalized function that centers the data.

(Refer Slide Time: 11:23)



We create a numeric column with *tf.feature_columns.numeric_column* API and pass the normalization function as an argument. The normalization will be run on each batch.

(Refer Slide Time: 11:57)

Bind	the MEAN	and STD to the no	ormalizer	fn using <u>fu</u>	nctools.pa	artial.							
[25]	<pre># See w normali numeric</pre>	hat you just co zer = functool: column = tf.f	r <mark>eated.</mark> s.partia: eature co	l(normalize	e_numeric_c ric column(data, mean: ('numeric'	MEAN, std=S	TD) fn=normal	izer, sha	pe=[ler	NUMERI	: FEATU	RES
	numeric numeric	_columns = [num _column	meric_co	lumn]	_							_	
0	Numeric	Column(key='n	umeric',	shape=(4	,), defaul	lt_value=M	one, dtype	=tf.float3	2, norma	lizer_	fn=func	tools.p	part
Whe	n you trair	n the model, inclu	de this fe	ature colum	nn to select	and center t	his block of i	numeric dat	a:				,
0	example	_batch['numeri	:']							↑ \	6	Î	:
	numeric numeric	_layer = tf.ke _layer(example	ras.layen _batch).n	rs.DenseFea numpy()	atures(nume	eric_colum	is)						
[]													

(Refer Slide Time: 11:59)

oode	+ Text & Copy to Drive V Disk	•	🖍 Editing	NF
		$\uparrow \downarrow$	• 🗘 🖡	:
U	example_batch['humeric']			
0	<pre>ctf.Tensor: id=579, shape=(5, 4), dtype=float32, numpy= array([[28. , 0. , 0. , 26.55], [22. , 1. , 1. , 29.], [2. , 4. , 1. , 39.688], [45. , 1. , 1. , 164.867], [18. , 1. , 0. , 6.496]], dtype=float32)></pre>			
O _c	<pre>numeric_layer = tf.keras.layers.DenseFeatures(numeric_columns) numeric_layer(example_batch).numpy() mean based normalization used here requires knowing the means of each column ahead of time.</pre>			
- Cate	gorical data			
 Cate Som 	egorical data 2 of the columns in the CSV data are categorical columns. That is, the content should be one of a limite	d set of optio	ns.	

Let us put all the numeric columns in the dense features and pass it to the model.

(Refer Slide Time: 12:15)

	with tidata	Teari X @ mays-Salanday X + Here and the second s			•	a (
+ CO	ode	+ Text & Copy to Drive		•	Editi	ng N	TULL
>	[26]	<pre><tf.tensor: 4),="" dtype="float32," id="579," numpy="</pre" shape="(5,"></tf.tensor:></pre>					
	9	array([[28. , 0. , 0. , 26.55], [22. , 1. , 1. , 29.], [2. , 4. , 1. , 39.68], [45. , 1. , 1. , 164.867], [18. , 1. , 0. , 6.496]], dtype=float32)>					
			\uparrow	V€	¢ i		
	0	<pre>numeric_layer = tf.keras.layers.DenseFeatures(numeric_columns) numeric_layer(example_batch).numpy()</pre>				_	
	0	array([[-0.13 , -0.474 , -0.479, -0.144], [-0.61 , 0.395, 0.782, -0.099], [-2.208 , 3.001, 0.782, 0.097], [1.228 , 0.395, 0.782 , 2.39], [-0.93 , 0.395, -0.479, -0.511]], dtype=float32)					
	The r	nean based normalization used here requires knowing the means of each column ahead of time.					
•	Cate	egorical data					
	Some	e of the columns in the CSV data are categorical columns. That is, the content should be one of a limited se	tofo	ptions			
	Unet		h		Loolum		
• *	e .	n 5 🧿 🛛 🖪 🤱			£ ^ J	120 8/10 8/11	09M 12019

So, you can see that the numeric layer has normalized data for each of the numeric features. Let us see how to handle the categorical variable.

(Refer Slide Time: 12:35)



We use tf.feature_column.indicate_column to convert each categorical variable into a one hot encoding. We construct dictionaries of values in each of the categorical variables. For example, sex has male and female in it, class has first second and third in it. There are A to J decks then there are few embarkation towns, and alone has two values yes or no. Then we take each of these categorical variables along with their vocabularies and convert as categorical variables into indicator columns using the vocabulary list.

(Refer Slide Time: 13:39)



We can see that there are indicator columns that are created from this exercise.

(Refer Slide Time: 13:57)



We also create a dense feature layer out of the categorical columns.

(Refer Slide Time: 14:05)



Since the categorical columns are converted into numbers you can combine categorical columns and numerical columns into dense feature layer and pass it to the model.

(Refer Slide Time: 14:27)



(Refer Slide Time: 14:37)

- Bui	ld the model	
Build	atf.keras.Sequential,starting with the preprocessing_layer.	
0	<pre>model = tf.keras.Sequential([preprocessing_layer, tf.keras.layers.Dense(128, activation='relu'), tf.keras.layers.Dense(128, activation='relu'), tf.keras.layers.Dense(1, activation='sigmoid'),]) model.compile(loss='binary_crossentropy', optimizer='adam', metrics=['accuracy'])</pre>	↑ ↓ ∞ ¢ i
[,] Trai	in, evaluate, and predict	

You can see that the preprocessing layer which is a dense feature layer is passed to a model. This is how we handle the dataset creation using CSV and preprocessing. Let us check how do we do shuffling and batching.

(Refer Slide Time: 15:03)



So, we can simply call a shuffle transformation on the training data by specifying the buffer size.

(Refer Slide Time: 15:29)

	ella statu (Tursori X 🧠 cavip buresearch.google.com	aynb - Coluboratory x 🛊 /github/tensorflow/docs/blob/master/site/	/en/12/tutorials/load_data/csx.ipyrb#s	crolTo=sTm_pD90gdl					- (1 ₀× α (10) =
4 Co	de	+ Text	🍐 Copy to Drive				V RAM	•	Editing	NPTEL
-		120/120	[=====	J	- US 3ms/step	- 105S: 0.3184 -	accuracy: 0.	85/9		
>	[37]	Epoch 16	/20							
_	~	126/126	[======================================]	- Øs 3ms/step	- loss: 0.3152 -	accuracy: 0.	8533		
	Θ	Epoch 17	/20							
		126/126	[====================================]	- Øs 3ms/step	- loss: 0.3115 -	accuracy: 0.	8616		
		Enoch 18	/20	,						
		126/126	/ 200 [1	- As 3ms/ston	- loss · 0 3083 -	accuracy: 0	8502		
		120/120	/20	j	- 05 Sills/Step	- 1055. 0.5005 -	accuracy. 0.	5352		
		126/126	/ 20 r		Ac 2mc/ston	locc. 0 2054		0000		
		120/120	(20		- 05 Sms/step	- 1055: 0.5054 -	accuracy: 0.	6004		
		Epoch 20	/20	,						
		126/126	[======================================	j	- 0s 3ms/step	- 1055: 0.3026 -	accuracy: 0.	8633		
		<tensorf< td=""><td>low.python.keras.</td><td>callbacks.Histo</td><td>ry at 0x7fe4678</td><td>c8898></td><td></td><td></td><td></td><td></td></tensorf<>	low.python.keras.	callbacks.Histo	ry at 0x7fe4678	c8898>				
	Once	e the model	is trained, you can ch	ieck its accuracy or	n the test_data s	et.				
	[38]	test los	s, test accuracy =	model.evaluate(t	est data)					
	[50]		·, ····-							
		print('\	n∖nTest Loss {}, T∉	est Accuracy {}'.	format(test_loss	, test_accuracy))				
							1			
	0	53/	Unknown - 1s 11ms	/step - loss: 0	.5138 - accurac	v: 0.7955				
	0									
		Test Los	s 0.5137907432497	673. Test Accur	acv 0.795454561	7103577				
				,	,			A 1	A ==	
1								$\land \lor$	Θ / 🛯	: n -
	Use ·	tf.keras.M	Model.predict to inf	fer labels on a batcl	h or a dataset of b	atches.				
										1343.044
# H	e	J J 🧕	N 🖬 👘						x ² ^ A i	NG 8/17/2019

Then we can evaluate the model on the test data and see the accuracy.

(Refer Slide Time: 15:39)



You can also check out the model predictions and compare that against the actual outcome. We will be doing more and more of model tanning and evaluation in the subsequence secessions. So, we are not giving much stress on that particular part. In the sessions more stress is given on creating input pipelines from multiple sources. Let us look at another source of creating input pipeline, this time with NumPy.

(Refer Slide Time: 16:19)



So, here the data is stored in NumPy array and you want to construct the dataset from this NumPy arrays. We read the NumPy array where we have Fashion-MNIST dataset which is stored as NumPy array.

(Refer Slide Time: 16:59)



We read that we get train_examples, test_examples, train_labels and test_labels corresponding to train examples, test examples, training labels and test labels.

(Refer Slide Time: 17:13)



We use *from_tensor_slices* from the dataset API for constructing datasets from NumPy arrays. Finally, we can do shuffling and batching of the training dataset, whereas on the test dataset we can apply let say batch transformation and then you can build a model and evaluate its accuracy on the test dataset.

(Refer Slide Time: 18:03)

* caberwardt pogle carginality factorie X or support Galaxies X or support Galaxies X + +	e (10) **
reode + Text ▲ Copy to Drive Connect -	Editing NPTEL
> Luau panuas uatamanes with truata	
🏦 View on TensorFlow.org	
This tutorial provides an example of how to load pandas dataframes into a tf.data.Dataset.	
This tutorials uses a small <u>dataset</u> provided by the Cleveland Clinic Foundation for Heart Disease. There are several hundred CSV. Each row describes a patient, and each column describes an attribute. We will use this information to predict whether a heart disease, which in this dataset is a binary classification task.	rows in the patient has
<pre>[] from _future_ import absolute_import, division, print_function, unicode_literals try: # Xtensorflow_version only exists in Colab. %tensorflow_version 2.x except Exception: pass import names as nd</pre>	
import tensorflow as tf	

Let us look, how to construct datasets from panda's data frames. Here we have a dataset provided by Cleveland Clinic Foundation for Heart Disease. There are several hundred rows in the dataset. The dataset is in CSV format. Each row describes a patient and each column describe an attribute of a patient. We will use this information to predict whether a patient has heart disease or not.

(Refer Slide Time: 18:55)



So, let us look at the steps to create a dataset. We use *dataset.from_tensor_slices* to create the dataset from the Pandas DataFrame. So, in NamPy and Pandas whether data is sitting in memory we use from_tensor_slices() function of dataset api to create data sources.

(Refer Slide Time: 19:21)



Let us try to understand how to create a dataset from an image dataset. The image class is provided in the directory name and we have multiple directories of the images. So, let us see how to create input data pipeline for such a data through this example.

(Refer Slide Time: 19:53)



First download the file.

(Refer Slide Time: 20:07)



(Refer Slide Time: 20:15)



You can see that there are directories like dandelion, daisy, tulips which contain images of the respective flowers. There are 3670 images in all.

(Refer Slide Time: 20:45)



And, these are top ten image paths. You can see that there are few images of dandelion flower, some images of sun flowers, some of roses, tulips, daisy and so on. Let us look at couple of images.

(Refer Slide Time: 21:07)



We will first assign the label to each of the images.

(Refer Slide Time: 21:13)



For this what we did is we took the string labels and assign an index to each of the labels. So, we have converted those strings into numbers. (Refer Slide Time: 21:41)



We can read the image using tf.io.read_file() function by providing the image path.

(Refer Slide Time: 21:59)



And, we can decode the image into an image tensor. So, we can see that an image tensor is a 3D tensor which shape $240 \times 380 \times 3$; there are three colour channels and each images of the size 240×180 . We will resize each of the image to 192×192 and we normalize the image by dividing each pixel value by 255.

(Refer Slide Time: 22:37)



So, we wrap all these transformation into a simple function called preprocess_image where we first decode the jpeg file resize image and then normalize the image in the range between 0 to 1. We put load and pre process, we wrap it in load_and_preprocess_image() function.

(Refer Slide Time: 23:15)



Let us look at some of these images after pre-processing.

(Refer Slide Time: 23:23)



Let us see how to build a *tf.dataset* from this images. So, you first construct a dataset of paths. All the paths are in memory.

(Refer Slide Time: 23:37)



We use *from_tensor_slices*, as you are doing earlier in case of NumPy and Pandas datasets. Now, we create a new dataset that loads and formats image on the fly by mapping *preprocess_image* over the dataset of path.

(Refer Slide Time: 24:17)



So, we use a map transformation and the and apply this function on each and every image in the dataset of paths. Let us look at some of the images from the dataset. We look at first four images. There are images along with their descriptions.

(Refer Slide Time: 24:43)



We build dataset of labels next using from tensor slices function.

(Refer Slide Time: 24:55)



We use dataset.zip to assign label to each of the image. And, you can see that each images of the size $192 \times 192 \times 3$ and each label is a scalar.

(Refer Slide Time: 25:27)



In order to train a model with this dataset, we want to first shuffle the data and then we want to batch it and repeat forever and we want to make sure that batches are available as soon as possible.

So, we use shuffle() transformation for shuffling the data by giving sufficient buffer size, then we use repeat() for repeating the dataset, epochs after epochs; specify the batch and we use a prefetch() function that clears the dataset page batches in the background while model is training.

(Refer Slide Time: 26:15)



There are few points you note here; the order is important. A shuffle after repeat would shuffle items app across epoch boundaries. Some items will be seen twice before others are seen at all. A shuffle after batch would shuffle the order of the batches, but not shuffle the items across the batches. We use *buffer_size* of the same size as a dataset for a full shuffle up to the dataset size large values provide better randomization, but they need more memory. Shuffle buffer is filled before any elements they are pulled from it.

Large buffer size may cause a delay when your dataset is starting. The shuffle dataset does not report the end of a dataset until the shuffle buffer is completely empty a dataset is started by repeat causing another wait for shuffle buffer to be filled. We use the *dataset.apply()* method to address the last point.

(Refer Slide Time: 27:31)



And, we fuse it with shuffle and repeat() function as seen here.

(Refer Slide Time: 27:43)



Finally, we pipe the dataset to the model and model and train the model. The MobileNet expects the input to be normalized to -1 to +1 range. So, we convert our data from 0 to 1 range to -1 to 1 range.

(Refer Slide Time: 28:19)



We define a change_range() function and apply the tern each element with map function. The MobileNet returns 6 x 6 special grade of features for each image. We pass it a batch of images to see the output of the MobileNet.

(Refer Slide Time: 28:43)



You can see that it returns 6 x 6 special grids. Finally, we built a model wrapped around MobileNet and GlobalAveragePooling2D to average over those special dimensions and then use a dense layer with softmax activation as output layer.

(Refer Slide Time: 29:33)



We can see that it produces output of the shape 32×5 .

(Refer Slide Time: 29:45)

Coue	1	- Text	🙆 Co	py to Driv	e							~	Disk 🔳		•	/	Editing	
Co	omp	ile the n	nodel to c	escribe th	ie training	g proce	edure:											
[3	9]	model.	compile(optimizer loss=' <mark>spa</mark> netrics=[=tf.kera rse_cate "accurac	as.opt egoric cy"])	timizers. cal_cross	.Adam(), sentropy'	,									
Th	nere	are 2 tra	ainable va	riables - t	he Dense	eweigh	hts and b	bias:										
[4	0]	len(mo	del.trai	nable_var	iables)													
6	3	2													1			
	5	model.	summary()										T	¥	e ,	•	-
Yc	You are ready to train the model.																	
No de	Note that for demonstration purposes you will only run 3 steps per epoch, but normally you would specify the real number of steps, as defined below, before passing it to model.fit():																	

We compile the model to describe the training procedure and train the model.

(Refer Slide Time: 29:55)

oue	model.summarv()			Disk Disk	Eultilig			
0	model. Summary ()							
0	Model: "sequential"							
	Layer (type)	Output Shape	Param #					
	mobilenetv2_1.00_192 (Model)	(None, 6, 6, 1280)	2257984					
	global_average_pooling2d (GI	(None, 1280)	0					
	dense (Dense)	(None, 5)	6405					
	Total params: 2,264,389 Trainable params: 6,405 Non-trainable params: 2,257,	984						
You	eready to train the model.							
Note defi	e that for demonstration purposes yo ned below, before passing it to mode:	ou will only run 3 steps per ep l.fit():	poch, but normally you v	vould specify the real numbe	r of steps, as			
r 1	stons non onoch_tf moth coil/	len(all image naths)/BATC	"H ST7E) numpy()					

(Refer Slide Time: 30:01)



For the demonstration purpose, we will only run three steps per epoch.

(Refer Slide Time: 30:11)

	htás (bu x i i nejvist-talason x i i ferundastr-talason x i i nejvist-talason x i i nejvist-talason x i + k otámetrágojá zenýtháhrodhulóztáki-mánya vritánicka kaj karnya synchrotine, fisatoly	a () *
and the second s	e + Text 🕼 Copy to Drive 🗸 RAM 🛄 👻 🌶	Editing NPTEL
>	 W8817 07:00:47.466316 140085485397888 deprecation.py:323] From /tensorflow-2.0.0b1/python3.6/tens Instructions for updating: Use tf.where in 2.0, which has the same broadcast rule as np.where 3/3 [===================================	orflow/pyt
Ŧ	Performance Iote: This section just shows a couple of easy tricks that may help performance. For an in depth quide see <u>Input Pipeline Perfo</u>	ormance.
	he simple pipeline used above reads each file individually, on each epoch. This is fine for local training on CPU, but may not be r GPU training and is totally inappropriate for any sort of distributed training.	sufficient
	o investigate, first build a simple function to check the performance of our datasets:	
	<pre>import time default_timeit_steps = 2*steps_per_epoch+1 def timeit(ds, steps=default_timeit_steps): overall_start = time.time() # Fetch a single batch to prime the pipeline (fill the shuffle buffer), # before starting the timer the direction to the timer</pre>	
•	8 🗷 🧕 🚺 🛄 🐮	^ J ⁴ ENG ^{1222,PM} □