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

Lecture – 15 Regression

[FL]. In this session, we will build a deep learning model for Regression. In the last session, we built a deep learning model for classifying fashion accessories into one of the ten categories. So, in this particular session we will build a regression model for predicting fuel efficiency of vehicles.

(Refer Slide Time: 00:45)

 Initias/ magnessaring/organite group we are investigated in the set of the		ж
This notebook is open with private outputs. Outputs will not be saved. You can disable this in <u>Notebook settings</u> .		
O basic_regression.ipynb B		GÐ SHARE
	RAM I	/ EDITIN
Regression: Predict fuel efficiency		
🍸 New on Tennorflow.org 🚥 Rum in Goorde Colan 💭 View source on Githus 🛓 Download notebook		
is a represented problem we aim to product the output of a continuous value, like a price or a probability. Contract this with a classification problem, where we aim to		
In a regression problem, we aim to predict the output of a continuous value, like a price or a probability. Contrast this with a classification problem, where we aim to select a class from a list of classes (for example, where a picture contains an apple or an orange, recognizing which fruit is in the picture).		
In a regression problem; we amit predict the output of a continuous value, like a proce or a probability. Contrast this with a classification problem; where we amit to select a class from a list of classes (for example, where a picture contains an apple or an orange, recognizing which fruit is in the picture). This notebook uses the classic <u>alum MPS</u> Dataset and builds a model to predict the fuel efficiency of late 1975 and early 1980 automobiles. To do this, will		
In a regression problem, we am to predict the output of a continuous value, like a proce a probability. Contrast this with a classification problem, where we am to select a class from a list of classes (for example, where a picture contains an apple or an orange, recognizing which fruit is in the picture). This notebook uses the tho <u>ods (Amb MBC</u>) bases are tailed as a model to predict the fuell efficiency of the automation problem, the automation of the set of the picture). This notebook uses the tho <u>ods (Amb MBC</u>) bases are tailed as a model to predict the fuell efficiency of the set of this automation is a subscription of many automabiles. To do this, well provide the model with a description of many automabiles from that time period. This description includes attributes like: cylinders, displacement, horsepower, and		
In a negression problem, we am to predict the output of a continuous value, like a proce as probability. Contrast this with an assistantiation problem, where we am to select causes from a list classes (for earning-kine a price contrains and open or an orange, recogning which this is in the product.) This notebook uses the classic <u>Appa MPG</u> basest and ballists amodel to predict the fuel efficiency of late 1970s and early 1960s assorables. To do this, will provide the model with a description of many automobiles from that time period. This description includes attrabutes like: cylinders, displacement, horsepower, and weight.		
In a regression problem, we am to predict the output of a continuous value, like a proce as probability. Contrast this with an a classification problem, where we am to select class form and its classific (dataset) for sample, where a price contrains and good or an orange, recogning mich that is in that of puture). This notebook uses the classific <u>dataset MPG</u> bataset we build to predict the fuel efficiency of late 1970s and early 1980s automobiles. To do his, will provide the model with a description of many automobiles from that time period. This description includes attributes like cylinders, displacement, honespower, and weight. This example uses the tF. Averas API, use <u>this guide</u> for details.		
In a regression problem, we am to predict the output of a continuous value, like a proce as probability. Contrast this with an adstratication problem, where we am to select calcus form and its collasses (the sample, where a price contrains and good on a ronger, recogning which this is in the product). This notebook uses the classic <u>Addro MPG</u> Dataset and builds a model to predict the fuel efficiency of late 1970s and early 1980s automobiles. To do this, we'll provide the model with a decorption of many automobiles from that time period. This description includes attributes like cylinder, displacement, horsepower, and weyf. This readers the f. Kersa X-RI see <u>this guide</u> for details.		
In anymetering problem, we am to product the output of a continuous value, like a proce or a probability. Contrast the winh an cassing control where we am to select cassiss than its additionate (the output of the optical control and the output of the ou		
In a regression problem, we am to predict the output of a continuous value, like a proce or a probability. Contrast the winh an classification problem, where we am to select classifion fail and classific descending where a price contrains and operation of the selection of the problem. This description of many automobiles from that time period. This description includes attributes like cylinder, dipalacement, horspower, and weight. This extenders due to the selection of many automobiles from that time period. This description includes attributes like cylinder, dipalacement, horspower, and weight. This extenders due to the selection of many automobiles for details.		
In a negression problem, we am to predict the output of a continuous value, like a proce as probability. Contrast the wint and castralization problem, where we am to set act casts from an like act casts from an like we proceed the output of a contrast the switch probability. Contrast the wint and castralization problem, where we am to set act casts from an like act casts for an like accession problem (with the example of the output		
in a negression problem, we am to predict the output of a continuous value, like a proce as probability. Contrast the wint and castralization problem, where we am to select castra form and in a classic data. DBMC classes is a classic according on a ronzer, ecception problem (in the classic the classic data. DBMC classes (be earning where explance ontrains and problem). This description includes ante-base share on the selection of the selectio		
In anymetering modeling was anto bysech the output of a continuous walks also give a process a probability. Contrast this with an assistantion prodem, where was not observed catasts from a list of classes (the sample was explane contrast and prode or a normer, ecosympt which fur is in the prode.) This notebook uses the classic (here have a produce of the same produce or a normer, ecosympt which fur is in the product. To do this, we'll produce the norm (with any associated to produce the norm (with a description of many automobiles from that time period. This description includes attributes like cylinder, diplacement, horseposer, and weyld. This example uses the ft.kers a APs we flug give for details.		
h a ngresson problem, we am to predict the output of a continuous value, like a proce or a producing, probability, Contrast the wind and castralization problem, where we am to setted casts from an like we protoce or norms, encopied primode in the output contrast the wind we have been accounted on the set of the output contrast. The value denity 1900 a submobile, To do this, we'll provide the model with a description of many automobiles from that time period. This description includes attributes like opfinder, diplacement, horsepower, and weight. The easember sets the f.k.erss API, use this going for details.		
in a ngreson problem, we am to predict the output of a continuous value, like a proce as probability. Contrast the wint and costantication problem, where we am to set act casts from a like disest (bases proke varies) where we am to set act casts from a like disest (bases proke varies) where we am to set act casts from a like disest (bases proke varies) where we am to prove the model to predict the fuel efficiency of like 1970s and entry 1960s automobiles. To do this, well provide the model with a description of many automobiles from that time period. This description includes attributes like cylinder, displacement, horspoore, and weight. The campion uses the f1, kerns APL tee this guide for details. Like the model of the predict of the set of the		
h a ngresson problem, we am to predict the output of a continuous when the a proc or a problem (potential the wind and castralic tor) poder, where we am to set exist casts that and its description of themse of the output of a continuous when the approx or a more, necessary minh that is in the priore. This description is description of many automobiles from that there peiced the fund effective of the set of themse. We prove the more description of many automobiles from that there peiced. This description includes attributes like optimizes, and wenty. This example uses the tf.kers APR see this guide for details.		
h a ngresson problem, we am to predict the output of a continuous when the a proce or problemity. Contrast the winh a costantication problem, where we am to set exist casts from a large explane contrast and by even of a contrast the set of the cost of the set of the set of the cost of the set of the set of the set of the cost of the set of the cost of the set of the s		

In regression problem, the output is a real number which is like a price or fuel efficiency measured by miles per gallon. Contrasting this with classification problem where we aim to select a class from list of classes, regression problem strives to predict a real number.

In this particular exercise we will use a classic Auto MPG Dataset and build a model to predict the fuel efficiency of 1970s and 80s automobiles. To do this, we will provide a model with description of many automobiles from that time period. These features include cylinders, displacement, horsepower and weight of the automobile. As usual before starting the

notebook let us connect to a collab runtime and install the useful softwares which are not present. First we will install a software called seaborn for plotting a pair plot.

(Refer Slide Time: 01:59)



So, what is the seaborn is installed we will import some of the plotting libraries like matplotlib.pyplot. We will import pandas for manipulating data. Then we will we will import seaborn for pair plot and we will also install tensorflow 2.0 and import keras and layers library from tensorflow.

(Refer Slide Time: 02:29)



Let us run this particular cell. At the end of the cell we make sure that we have the right version of tensor flow present in our colab runtime. We ensure that by printing the tf version which is 2.0.0 beta 1 which is the desired version for this particular exercise.

(Refer Slide Time: 02:49)



The auto MPG dataset is available from UCI machine learning repository. Our first job is to get this particular data. In the last exercise, we the dataset was present in the tensor flow data sets and it was easier for us to import and load the data in tensorflow. In the case of auto

MPG, this data set is not available in tensorflow so, we will have to first download the file and then load the data using pandas and we will import the data from pandas data frame into tensorflow. Let us first understand that particular process.

So, at the first step we will download the data. We will use keras utils.get_file() function for that the first argument is where we want to store the data and in the second argument here is the url of the file containing the auto mpg data. So, let us run this which will download this particular cell will download the data and the data is now present in auto MPG dot data as printed by the dataset path.

(Refer Slide Time: 03:55)

		mups//coldo.	researcingo	giecontryiana	u) tersoritor	Nuocs/Diou/III	astel/swereivitelu	DIOHAIS/ NEIG	phrase lei	тезаютаруно	rsuunio-unz	nagern				н
					This n	otebook is oper	h with private out	puts. Output	s will not b	e saved. You o	an disable this	in Notebook setti	ngs.			
0) basia	c_regressio	on.ipynb	B												GD SHAR
Fil	le Edit	View Inser	rt Runtime	Tools Help												
	ODE 🖬	TEXT 🛊	e CELL 🐐	CELL & CO	OPY TO DRIVE	1								V RAM I	3.	/ EDITIN
Gett	ine dat	a														
First o	downloa	ad the datase	я.													
[4]	datase datase	et_path = ke et_path	aras.utils.	get_file("auto	o-mpg.data	", "http://ar	chive.ics.uci.e	du/ml/mach:	ine-learni	ng-databases	/auto-mpg/aut	co-mpg.data")				
θ	st download the da il dataset_path dataset_path Downloading 32768/30286 ['/root/.kera: nport it using panda colume_names raw_detaset = dataset = raw dataset = raw	oading data /30286 [=== t/.keras/da	from <u>http</u> stasets/auf	o://archive.i	cs.uci.ed	<mark>u/ml/machine</mark> - Os 3us/ste	<u>-learning-data</u> p	ibases/auto	o-mpg/aut	<u>>-mpg.data</u>						
Impor	rt it usir	ng pandas														
0	<pre>colmm_pass = [1907.'()[isder', '0[islosest', '0erspoor', '0eight',</pre>															
	datase datase	et = raw_dat et.tail()	aset.copy()												
0		MPG Cylin	nders Disp	placement Ho	rsepower	Weight Acc	eleration Mod	el Year O	urigin							
	393	27.0	4	140.0	86.0	2790.0	15.6	82	ŕ.							
	394	44.0	4	97.0	52.0	2130.0	24.6	82	2							
	395	32.0	4	135.0	84.0	2295.0	11.6	82	1							
	306	28.0	4	120.0	79.0	2625.0	18.6	82	1							
	197	1.0	4	119.0	82.0	2720.0	19.4	82	1							
and a state		-														

Let us import the data using pandas dataframe. We will first list out the column names. We have column names like Miles Per Gallon, Cylinders, Displacement, Horsepower, Weight, Acceleration, Model Year and the Origin. We will use pandas.read_csv() function because this data is present in a CSV format.

This particular function takes the data set, path column names and we tell how to handle the NA values. So, we want to replace NA values by question mark. We want to ignore anything after the tap and that is specified using the comment argument. The separator is a space and we want to skip initial spaces in the file. After specifying all these arguments we are able to

read the content of auto MPG dataset into the data frame raw_dataset. We will make a copy of raw_dataset as dataset. Let us run this code cell and examine what is there in the dataset.

So, we printed last five rows of dataset using dataset.tail() command and you can see here there are features like MPGs, Cylinders, Displacement, Horsepower and so on. Most of the features are numeric features in nature whereas, origin is a discrete feature it just has value 1 or 2. So, this particular part takes care of downloading the data and loading that into pandas dataframe. Once the data is loaded our next job is to clean the data and remove null values and perform normalization on the dataset. Let us take steps to preprocess the data.

(Refer Slide Time: 06:01)

C https://colab.research.google.com/github/tensorflow/docs/blob/master/site/en/r2/tutorials/keras/basic_regression.jpynb#scrollTo=4ZUDosChC1UN	\$
This notebook is open with private outputs. Outputs will not be saved. You can disable this in Notebook settings.	
C basic_regression.jpyhb BL File Edt Wew Insert Runtime Tods Help	GD SHAF
CODE TEXT 🔶 CELL 🕹 COPY TO DRIVE	V RAM I V EDITI
lean the data	
he dataset contains a few unknown values.	
<pre>6] dataset.isna().sum()</pre>	
PFC 0 Ultiplexent 0 Restspacet 0 alight 0 type: 1154	
to keep this initial tutorial simple drup those rows.	
he "origin" column is really categorical, not numeric. So convert that to a one-hot:	
Final (Dr-frat) Grand (Dr-frat) (art bit or the manual + the manual + 23.4, a (art bit or the manual + 1.4, a) (art bit or the manual + 1.4, a)	

So, let us first find out if dataset has got any null values. So, we essentially use isna() function and then followed by a some function to identify columns where null values are present. As you can see on the screen, horsepower column has few null values. In order to keep this tutorial or this exercise simple we will simply drop these rows. So, we will drop all the rows containing null values using function dropna().

Next we will convert some of the categorical attributes into one-hot encoding. So, we just saw that origin is a categorical attribute and we will use one-hot encoding to convert that into something that is usable by the machine learning algorithm.

(Refer Slide Time: 06:57)

					This n	olebook I:	s open with priva	le outputs. O	utput	IS WILL	ot be s	aved. You	I can disable this in <u>Notebook s</u>	ettings.			
C B) basi	c_reg	ression.ip	ynb 🖻	in .											GÐ SHARE	
	ODE E	TEXT	+ CELL	CELL &	COPY TO DRIV										RAM -	/ EDITING	
[6]	Accel Model Origi dtype	Year Year n : inte	0 0 54												USX		
To ke	ep this	initial t	tutorial simple	e drop those row:	s.												
[7]	datas	et = d	ataset.drop	na()													
The "	Origi	n" colu	mn is really c	ategorical, not nu	umeric. So conv	ert that to	a one-hot:										
[8]	origi	n = da	taset.pop('	Origin')													
0	datas datas datas datas	et['US et['Eu et['2a et.tai	A'] = (orig rope'] = (or pan'] = (or 1()	in == 1)*1.0 rigin == 2)*1.0 igin == 3)*1.0													
θ		MPG	Cylinders	Displacement	Horsepower	Weight	Acceleration	Model Yea	r L	JSA EI	arope	Japan					
	393	27.0	4	140.0	86.0	2790.0	15.6	8	2	1.0	0.0	0.0					
	394	44.0	4	97.0	52.0	2130.0	24.6	8	2 1	0.0	1.0	0.0					
	395	32.0	4	135.0	84.0	2295.0	11.6	8	2	1.0	0.0	0.0					
	396	28.0	4	120.0	79.0	2625.0	18.6	8	2	1.0	0.0	0.0					
-	-397	31.0	4	119.0	82.0	2720.0	19.4	8	2	1.0	0.0	0.0					
_	- Mr	8						_	-								-

So, let us pop the; let us pop the origin column from the data set and replace the origin column with three values because origin originally has three values USA, Europe and Japan. So, for USA we use the value of 1, for Europe we will use the value of 2 and for Japan we will use the value of 3.

So, let us do this particular step of converting the categorical attribute into numeric attribute and you can see that whenever USA is present you will see the value 1 appearing; whenever USA is present Europe and Japan are obviously, not present so, the values are 0; whenever Europe is present you will see a value of 1 corresponding to a column Europe; whenever value is 1 that column corresponds to whenever there is a value of one in the column of Japan that corresponds to the third value in the original data set. So, this is how we converted our numerical attribute into a categorical attribute and we have a transform data set.

(Refer Slide Time: 08:07)

dat [9] dat dat	taset["E taset["]	[urope'] = (o lapan'] = (or ail()	rigin == 2)*1.0 igin == 3)*1.0									Disk		
θ	MPG	i Cylinders	Displacement	Horsepower	Weight	Acceleration	Model Year	USA	Europe	Japan				
39	3 27.0	4	140.0	86.0	2790.0	15.6	82	1.0	0.0	0.0				
39	4 44.0	4	97.0	52.0	2130.0	24.6	82	0.0	1.0	0.0				
39	5 32.0	4	135.0	84.0	2295.0	11.6	82	1.0	0.0	0.0				
39	6 28.0	4	120.0	79.0	2625.0	18.6	82	1.0	0.0	0.0				
39	7 31.0	4	119.0	82.0	2720.0	19.4	82	1.0	0.0	0.0				
Now split We will us	the data se the te	aset into a train st set in the fir aset = dataset set = dataset	ing set and a tes al evaluation of c t.sample(frac=0 .drop(train_dat	t set. ur model. .8, nandom_sta aset.index)	te=0)									

Now that we have pre processed and cleaned dataset, the next job is to split the data into test and training set. So, we will simply use dataset.sample() and we set random seed so that every time we run this particular collab we get essentially the same training set. So, let us run this particular command this will select 80 percent of the examples in the training and anything that is in the training we drop and the remaining examples are copied into test dataset.

So, to running this colab we have two data sets - train_dataset and test_dataset containing training and test examples respectively. Let us have a quick look at the joint distribution of a few pairs of columns from the training set. We will use sns.pairplot() for plotting the joint distribution of a few pairs of columns and we do that for columns like MPG miles per gallon which is our target column or the column that we want to predict Cylinders, Displacement and Weights. Let us run this particular code cell and look at the pairplot.

(Refer Slide Time: 09:23)



So, in this particular pairplot on diagonal you have you can see the distribution of the values in each of the features. The off-diagonal elements essentially show the relationship between the feature on the row and the feature on the column.

As you can see that there is hardly any relationship between cylinder and MPG. Cylinder and MPG seems to be uncorrelated whereas, displacement and weight have some correlation with MPG because as weight goes up MPG seems to be coming down the same thing is happening with displacement as displacement increases the MPG seems to be coming down. You can see that MPG and weight are correlated features because as the weight increases displacement also tends to increase. So, we can get some useful insights into the features by looking at this plot and how these features affect the outcome can also be seen in this pair plot.

(Refer Slide Time: 11:01)



We first use describe() command on the dataset and obtain training stats, we remove the field for miles per gallon because there is a field that we want to predict and we do the transpose of the stats so that we can display it in a nice tabular fashion. So, you can see all the features of the model on the row side and on the column we have various statistics for each of the features. So, you can see that there are 314 rows in the data set that is why we have 314 as the count for each of the feature.

Then we can look at the mean of every feature, then the standard deviation and pretty much the min, max and few quartiles in between. As we can see that different features are on different scales. For example, minimum number of cylinders is 3 and the maximum number of cylinders is 8; minimum weight is 1649 maximum weight is 5114. So, you can see that the features are on a different scale and in order to bring them on the same scale later we will perform normalization operation.

(Refer Slide Time: 12:21)



Before doing normalization we will split features from the label. So, right now in our train_dataset we have features as well as labels. So, we will use pop command and we will remove the column corresponding to the label the label column. So, we essentially remove miles per gallon column and that gives us the training label and using the pop command on the test data we get the test label. Let us normalize the data leaving aside the labeling column.

Now, this normalization is a very important process. So, we want to make sure that we do we perform exactly the same normalization on test data as a training data or any other future data set that is coming to the model for prediction we apply the same normalization. So, in order to do that we will we have already calculated a train statistics. We define a normalization function that does the z score normalization which is defined as the value minus mean divided by the standard deviation. So, we calculate we perform normalization on the training data as well as test data.

And, we will also store this normalization parameters which are present in train underscore test and we will apply them for the prediction as well as any test on any other test instances. Now, that we have explored the data and pre-processed it, the next step is to build a model. Here we will try to build a neural network model to perform the regression task. Let us look at the architecture of neural network model that we will be building.

(Refer Slide Time: 14:19)

egression: Predict fuel efficiens: 🗴 🥨 basic_regression: pynb - Colabo:: x 🛉 therasuntisget, file Terocoffic:: x 4/ pandas.read, csv — pandas.024: x 4		- 0	
C https://colab.research.google.com/github/tensorflow/docs/blob/master/site/en/r2/tutorials/keras/basic_regression.jpynb#scrollTo=r26juK7Z68j-		\$	0
This notebook is open with private outputs. Outputs will not be saved. You can disable this in Notebook settings.			
O basic_regression.ipynb		GD SHARE	1
File Edit View Insert Runtime Tools Help			
CODE TEXT I CELL COPY TO DRIVE	Disk -	/ EDITING	
The model			
Build the model Build the model Here, will use a sequential model with two densely connected hidden layers, and an output layer that returns a single, continuous value. The mode building steps are snapped in a function, build_model, since we'll create a second model later on.			
<pre> out</pre>			
<pre>[] mode] = build model()</pre>			-
Inspect the model Use the .sumary method to print a simple description of the model			
Nate which is model. Take a batch of 10 examples from the training data and call wodel.predict on it.			

Before proceeding further I would like you to pay attention that we have 9 features in our model. We will use a keras sequential model or a feed forward neural network with a couple of hidden layers.

The first hidden layer is a dense layer with 64 unit, the second layer is also another dense layer is 64 units. In both the layers we use activation as relu and we have finally, a dense layer with exactly 1 unit. Let us look at this particular architecture on a board.



So, again coming back to neural network you have to fix the architecture of the neural network. We need to specify the number of hidden layers and number of units in each hidden layer. You also specify the input layer or what is the input looking like and the output layer. So, here we are using two hidden layers it is using 64 units say you have 9 features and the output layer we just have 1 unit.

(Refer Slide Time: 16:25)



We are using dense layers and in dense layers, we connect each node from the previous layer to all the 64 units in the current layer. There is an additional bias unit on each of the hidden units which I am not showing explicitly. We have a total of 10 parameters per unit in the first layer.

In the second layer, here we have 64 inputs from the previous layer plus one bias parameter. So, we have 65 parameters per hidden unit. The second thing you have to care about is what kind of activation we will be using. We will use relu as an activation function in the hidden layers and linear activation in the output layer.

So, you can see here we wrote a special function called build_model() and this particular part of the code is building the model layers.dense defines the first layer where we use dense units we are going to use 64 such kind of units using activation of relu and the input shape is exactly equal to the number of features that we are using in the training data set.

Next we also stack another dense layer with 64 units which uses relu activation function and finally, we have an output layer containing one dense unit. After specifying the model let us also specify the optimizer and the loss function. So, in the case of regression we use mean squared loss as loss function which is a standard loss function for regression and we use RMSprop as an optimizer in using RMSprop with learning rate of 0.001 and we track metrics like mean absolute error and mean squared error. Let us look at mean squared error and mean absolute error and let us look at mathematically what do they mean.

(Refer Slide Time: 21:57)



So in the regression framework what happens is that we learn; so, this is our training data for example and let us say this is the model that we learnt. Now, we make end up making some error on each of the points, but you can see that the point on the line is the predicted point and the point denoted by cross is the actual point. So, there is a difference over here there is some difference, here there is some difference here and so on.

So, if we essentially look at every value of y which is y i or let us use slightly different notation. So, we take i-th value of y and we look at the predicted value of i-th example and squared it. So, this defined as squared error for example, if you sum it across all the example we get sum squared error and if we divide it by 2n we get mean squared error.

Instead of doing the square instead if we just use look at the difference or mod of difference between the true value and the predicted value and average it out this is called as mean absolute error and this is called as mean squared error. Let us go back to collab and now train the model. So, let us run this particular code cell and let us build a model. So, now we have a model. Let us see how model looks like.

(Refer Slide Time: 24:15)

Insp	ect the model	imple description of the mor	el	Dok me	, contra
0	model.summary()				
0	Model: "sequential_2"	Outout Shane	Daram B		
	dense_6 (Dense)	(None, 64)	640		
	dense_7 (Dense)	(None, 64)	4160		
	dense_8 (Dense)	(None, 1)	65		
	Total params: 4,865 Trainable params: 4,865 Non-trainable params: 0				
Now [44]	try out the model. Take a batch o example_batch = normed_trai example_result = model.pred	of 10 examples from the train n_data[:10] lict(example_batch)	ing data and call model.p	i k	

We can see from the model summary that we have a sequential model or a feed forward neural network. We are using essentially three dense layer; two of them are hidden layer and the last layer is an output laye with exactly one output.

Now, that we have built a model let us try the model on a few examples. Let us take a batch of 10 examples from their training data and call model.predict() on it. Mind you, I mean we have not yet trained the model, but we are using randomly initialize rates and seeing whether model.predict() works.

(Refer Slide Time: 25:11)

	This investow is open will private outputs, outputs will for a safety. For call usable lins in <u>sociation acturity</u> s.			-
C)) basic_regression.ipynb 🛍		GÐ SHARE	(
00	COLI VYEW INDELL NUMBER TOOD THP DOE ■ TEXT	RAM I	/ EDITING	
lt see	is to be working, and it produces a result of the expected shape and type.			
Train	the model			
Train	the model for 1000 epochs, and record the training and validation accuracy in the history object.			
0	8 Study exclude programs by perturbing a single dot for each completed spech (completed spech ed(saf, spech, log)); if speck [3] = or prior(1); prior(1); end(1);			
Visua	lize the model's training progress using the stats stored in the history object.			
461	hit = nd DataFrama(hithory.hithory)			
1	hist["epoch] = history.epoch			
5	*			

Yes, it seems to be working and it produces a result of expected shape and type. So, we use this in order to make sure that the model has been set up properly and all shapes of the tensors are set as expected. Let us train the model for 1000 epochs and record training and validation accuracy in the history object.

We give the normalized training data and corresponding labels as input to mode.fit function and we record all the history in the history object that is returned by model.fit function. Let us train the model. So, model is training as you can see from the progress and once model is trained we will visualize a history data whatever is told in the history data it seems model is training is complete.

(Refer Slide Time: 26:23)



So, let us look at the last five rows of history data. History will have 1000 rows; one for every epoch and you can see that how the loss, so the loss is actually it is very interesting to note here loss is quite fluctuating loss is going up than coming down and again going up and again up. So, let us look at what is happening to the validation loss. Validation loss seem to be going up then coming down and here it has gone up. So, let us plot the history and see how it looks like.

(Refer Slide Time: 27:07)



Interesting, so, you can see that training data mean in terms of mean absolute error the training error is going down, but validation error is going up. The same thing we observe in the mean squared error that while training error is going down validation error does not seem to be improving.

(Refer Slide Time: 27:33)

This notebook is open with private outputs. Outputs will not be saved. You can disable this in <u>Notebook settings</u> .	
O basic_regression.ipynb B	
File Edit View Insert Runtime Tools Help	GD SHARE
CODE D TEXT + CELL + CELL COPY TO DRIVE	Disk - PEDITING
Comparison of the second	
This graph shows little improvement, or even degradation in the wildation error after about 100 epochs. Let's update the wolel. Let's call to automatically stop training when the vialition score doesn'improve. We'll use an ExityStopping caliback that tests a training condition for every epoch. If a set amount of epochs elapses without showing improvement, then automatically stop the testing.	
This graph hows filte impovement, or even degradation in the validation error after about 100 epoch, Lefs update the wold. Fit call to automatically stop training when the validation score devent impours (the use an anaforgispic) caliback that tests a training condition for every epoch. If a set amount of epochs elapses without showing improvement, then automatically stop the training. You can learn more about this caliback <u>here</u> .	
This graph hows liftle impovement, or even degradation in the validation ever after about 100 epoch, Lefs update the wold. If it call to automatically stop training when the validation score devent impours (Wie use an anaforgissing callback that tests a training condition for every epoch. If a set amount of epochs elapses without showing improvement, then automatically stop the training. You can learn more about this callback <u>then</u> : [ag] morel = the stall_model()	
This graph hows lifte improvement, or even degradation in the validation ere after about 100 epoch, Lefs update the wold. Fit call to automatically stop training when the validation score deem in more will use an analyticipic addack that tests a varineg condition for every epoch. If a set amount of epochs elapses without showing improvement, then automatically stop the training. You can learn more about this callback <u>here</u> [28] model = build_model() * The aptience parameter is the amount of epochs to check for japrovement. envy_tate = set-callback_here.	
This graph thow lifts improvement, or even degradation in the validation error after about 100 epoch, Let's update the work. If it call to automatically stop training whet the validation is not even the maximum call the use an adjoingous adjust that tests a training condition for every epoch. If a set amount of epochs edgeses which about about approximate the automatically stop the training. Your can learn more about this callback <u>base</u> [63] model = built_model() [#The patience presents in it. In the provide the training. #The patience presents in the interpret of epochs to your in the present end of the presentence of the present end your interpret of the patience of the present end of the present end of the present end your interpret end of the present end of the present end of the present end of the present fixture = makel.fittionend.presin_effect, your_lefted() [https://www.fittionend.	
This graph thow lifts improvement, or even degradation in the validation ere after about 100 spoch, Lefs update the work. If it call to automatically stop training whet the validation is not even the use an adjoingous allack that tests a training condition for every epoch. If a set amount of epochs edgeses which about about a properties the automatically stop the training. Nov can learn more about this callback base: [23] model = build_model() [23] model = build_model() [24] model = build_model() [26] model() [26] model = build_model() [26] model = build_	

So, this points to some kind of a; some kind of an over fitting problem in this particular data set. So, we will in order to overcome this particular over fitting problem, we will try to use early stopping as a means of correcting this issue. So we will stop the model before it starts over fitting. So, we will look at we will keep an eye on validation error and use early stopping callback for stopping the model before it starts over fitting.

So. callback early stopping is set up in this particular way; we use keras.callbacks.EarlyStopping where we use validation loss as a monitoring mechanism and we wait for 10 epochs. If in 10 epochs validation loss does not improve then we will then we decide to stop the model. And, we use and let us train the model again with the early stopping callback and we will plot the history to see how model does this time.

(Refer Slide Time: 28:41)



The model looks much better compared to the earlier model. The earlier model's training error seems to improve while validation error became worse or at least stayed the same. So that was an overfitting problem, but after doing early stopping regularization; regularization is one way of addressing the over fitting problem. So, we applied early stopping regularization here and after applying early stopping regularization we see that training loss and validation loss or validation error both seems to be improving as we keep training for more epochs.

(Refer Slide Time: 29:45)



So, let us evaluate the model on the test data and obtain numbers like mean absolute error and mean squared error. So, we have mean absolute error was of 1.85 miles per gallon. So, now, we have a model and we will use the model to predict the miles per gallon for the test data.

So, here we will give normalized test data has an input and ask the model to predict miles per gallon. Let us plot the predicted value of miles per gallon. So, we have true values of miles per gallon.

(Refer Slide Time: 30:23)

Regression Predict hel efficien: x 🤨 basic, regression, pynb - Colabo: x 📍 therasuralis.get, file TemorFic: x 🌵 pandas.read, cov — pandas.024. x +		- 0)
O a https://colab.research.google.com/github/tensorflow/docs/blob/master/site/en/i2/tutorials/keras/basic_regression.ipynb#scrolTo=Xe7RXH3N3CWU		\$)
This notebook is open with private outputs. Outputs will not be saved. You can disable this in Notebook settings.			
Chasic_regression.ipymb Eat View Insert Rumme Tools Help		GÐ SHARE	A
CODE DI TEXT 🔶 CELL 🗳 COPY TO DRIVE	V RAM	/ EDITING	^
r Make predictions			
Finally, predict MPG values using data in the testing set:			
<pre>tractices: = woil.media(cone.tst_data).fatted() plt.satter(un_ide), set_predictions plt.satter(un_ide), rest_predictions plt.satter(un_ide) p</pre>		1	
Contraction of the former of t			
The second predicts reasonably well. Lefs take a look at the error distribution.			
		1606	

And, the predicted values of miles per gallon and you can see that in most of the cases predicted values the true value and the predicted values are actually close by and are model seem to be working reasonably well.

(Refer Slide Time: 30:43)



Let us also plot the error distribution the mean absolute error distribution; mean absolute error distribution is not really Gaussian. But, we expect we might expect that because here we

have a very small number of samples. We just had 314 rows out of which 80 percent where use for training and only 20 percent for use for test.

(Refer Slide Time: 31:05)

So, this brings us to the end of the regression exercise. In this exercise we learned how to use deep neural network model for the task of regression; we also studied how to how to remedy or how to prevent the overfitting using early stopping as regularization mechanism. In the coming session we will look at how to use this deep neural network model to make predictions on the structured data. We will also study how to store and retrieve the model for the deployment purpose and also study underfitting and overfitting through some practical examples.