
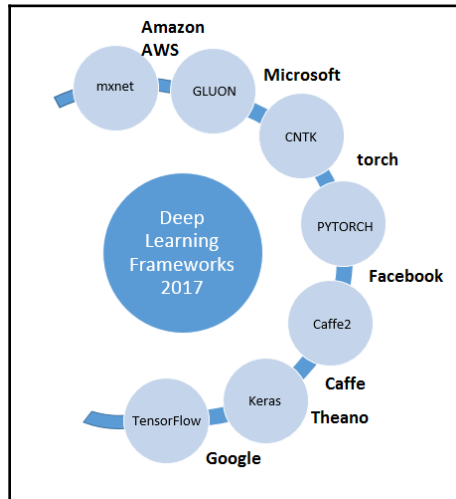
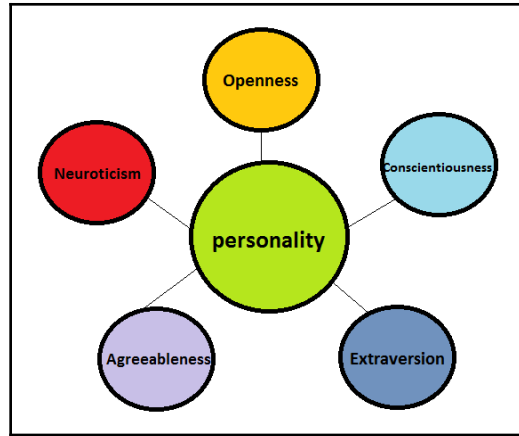
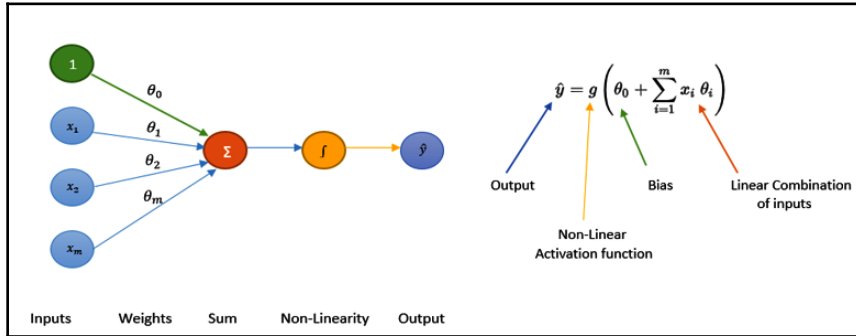


Chapter 01: Overview of Neural Networks




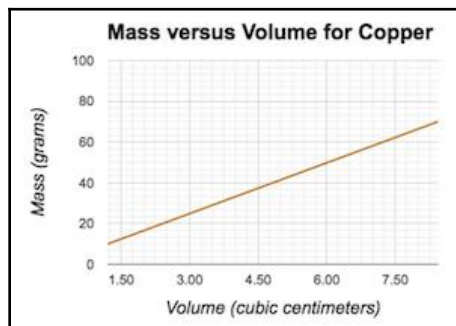
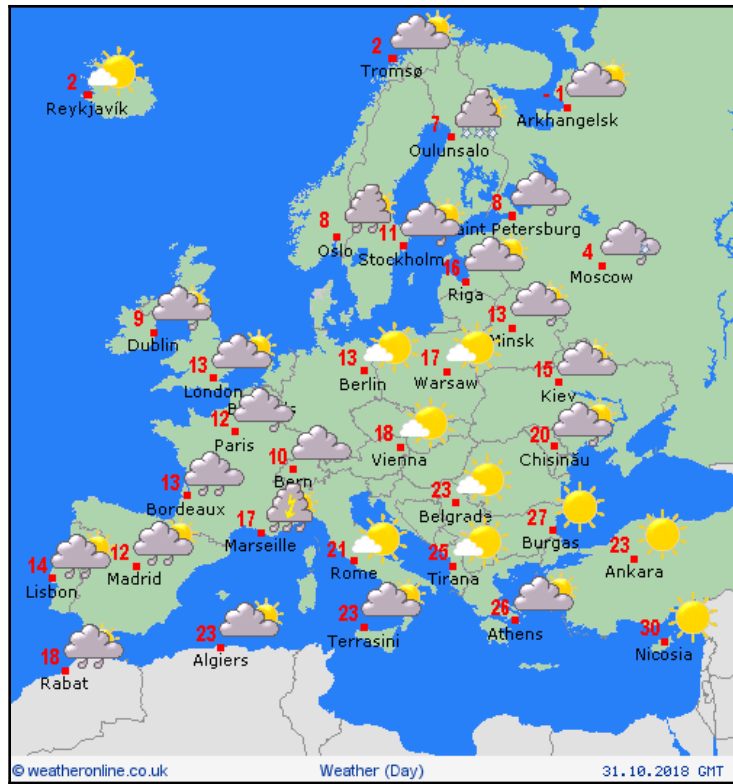



Chapter 02: A Deeper Dive into Neural Networks




$$\mathbf{x} = \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ \vdots \\ \vdots \\ x_n \end{bmatrix} \quad \mathbf{x}^T = [x_1 \ x_2 \ x_3 \ \dots \ x_n]$$









The purpose of activation functions is to **introduce non-linearities** into the network



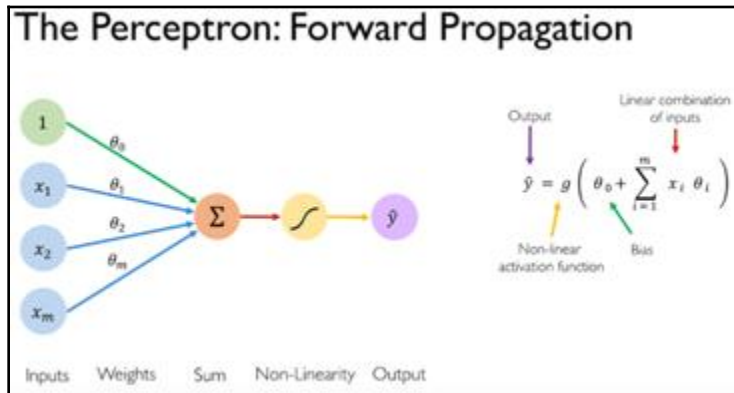
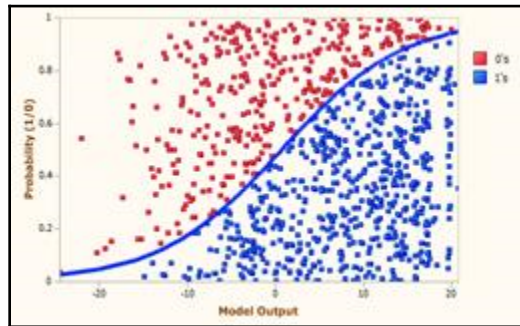

Linear Activation functions produce linear decisions no matter the network size





Non-linearities allow us to approximate arbitrarily complex functions



$$\text{output} = \begin{cases} 0 & \text{if } \sum_j w_j x_j \leq \text{threshold} \\ 1 & \text{if } \sum_j w_j x_j > \text{threshold} \end{cases}$$







$$\hat{y} = g \left(\theta_0 + \sum_{i=1}^m x_i \theta_i \right)$$


$$\hat{y} = g \left(\theta_0 + \mathbf{X}^T \boldsymbol{\theta} \right)$$

where: $\mathbf{X} = \begin{bmatrix} x_1 \\ \vdots \\ x_m \end{bmatrix}$ and $\boldsymbol{\theta} = \begin{bmatrix} \theta_1 \\ \vdots \\ \theta_m \end{bmatrix}$

The Perceptron: Example



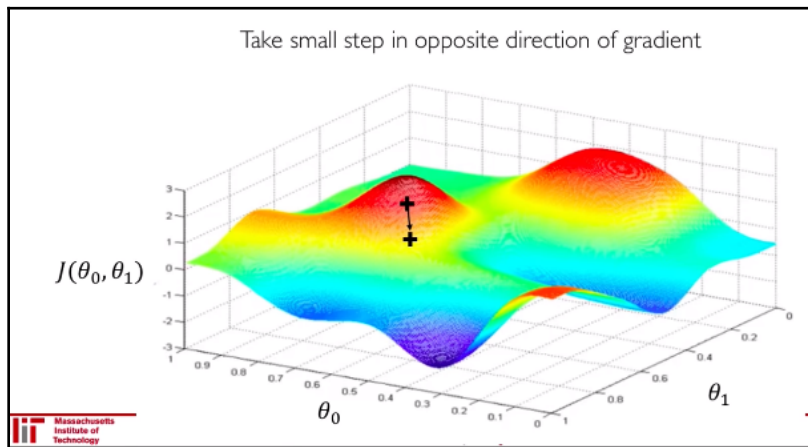
Loss Optimization



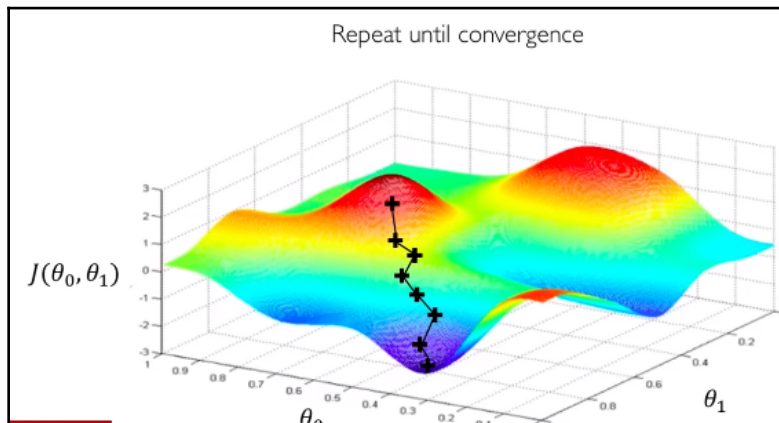
Algorithm





1. Initialize weights randomly $\sim \mathcal{N}(0, \sigma^2)$
2. Loop until convergence:
3. Compute gradient $\frac{\partial J(\theta)}{\partial \theta}$
4. Update weights, $\theta \leftarrow \theta - \eta \frac{\partial J(\theta)}{\partial \theta}$
5. Return weights

Take small step in opposite direction of gradient



Repeat until convergence






$$\theta \leftarrow \theta - \eta \frac{\partial J(\theta)}{\partial \theta}$$


How can we set the learning rate?

Multi Output Perceptron




$$z_i = \theta_{0,i} + \sum_{j=1}^m x_j \theta_{j,i}$$

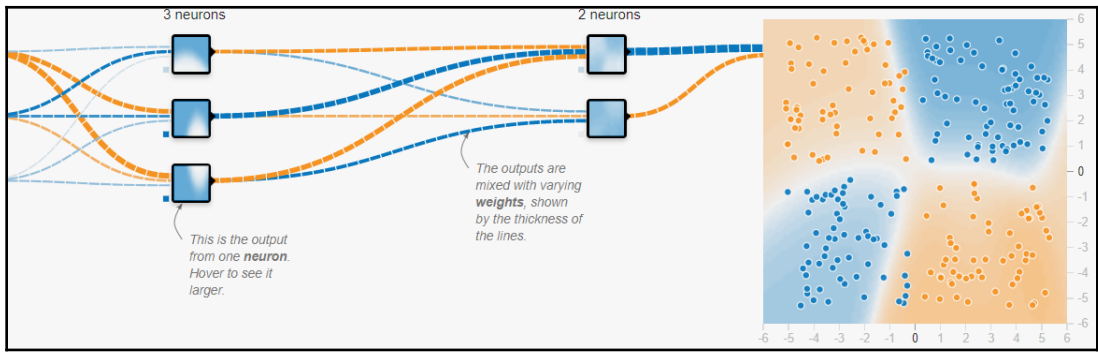


Single Layer Neural Network

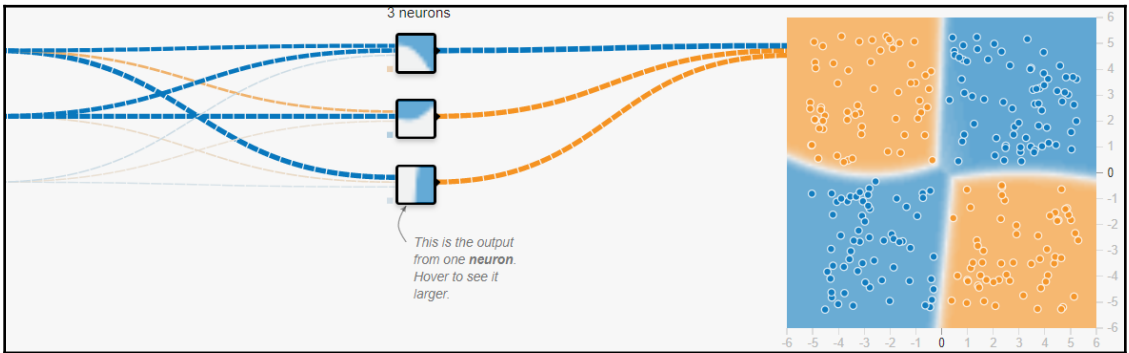
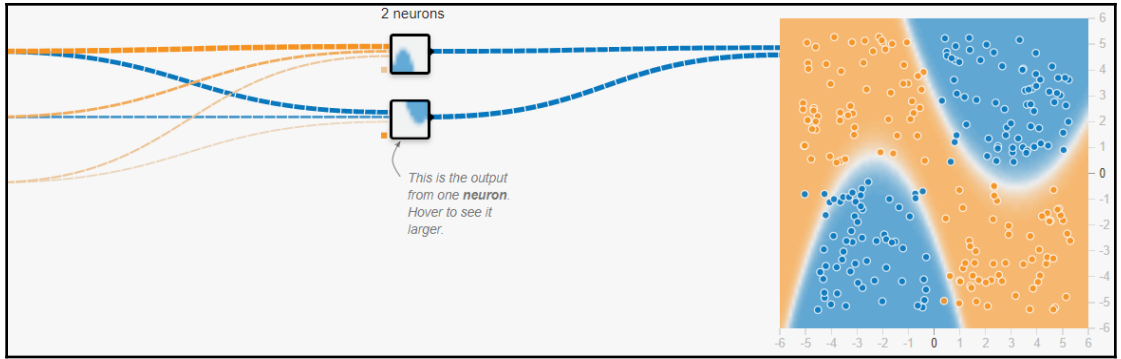


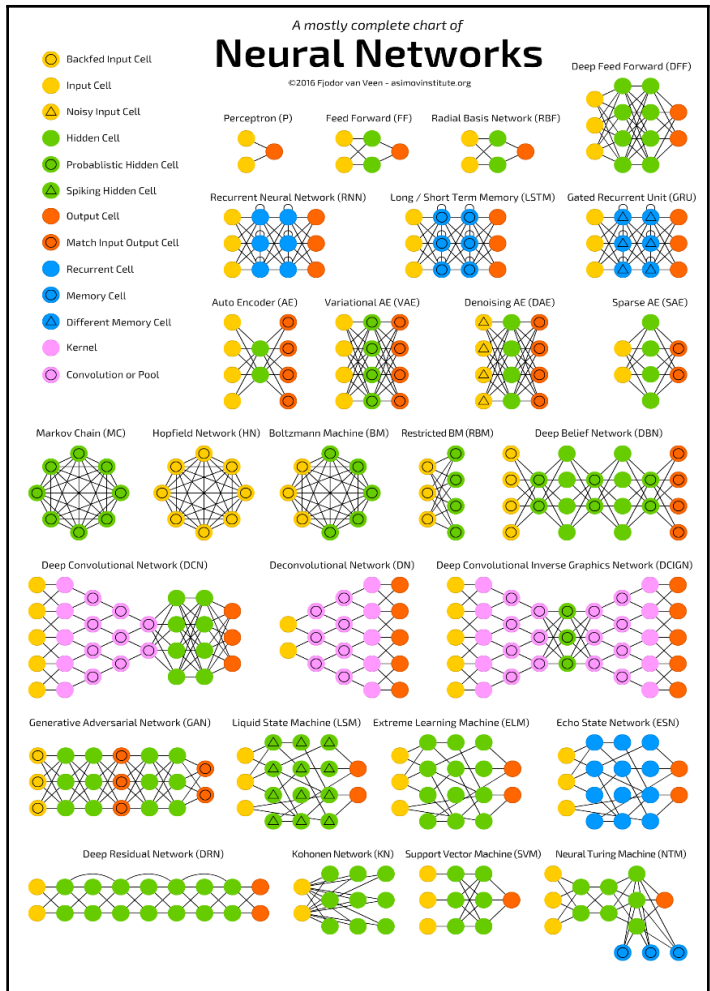
Inputs Hidden Final Output

$$z_i = \theta_{0,i}^{(1)} + \sum_{j=1}^m x_j \theta_{j,i}^{(1)} \quad \hat{y}_i = \theta_{0,i}^{(2)} + \sum_{j=1}^{n_{d1}} z_j \theta_{j,i}^{(2)}$$

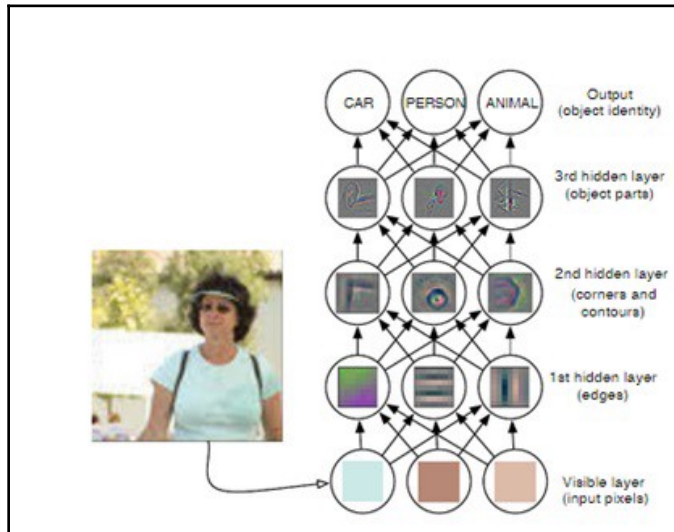











Chapter 03: Signal Processing - Data Analysis with Neural Networks






Activation Functions


Sigmoid

$$\sigma(x) = \frac{1}{1+e^{-x}}$$




tanh

$$\tanh(x)$$




ReLU

$$\max(0, x)$$



Leaky ReLU

$$\max(0.1x, x)$$








Maxout



$$\max(w_1^T x + b_1, w_2^T x + b_2)$$



ELU


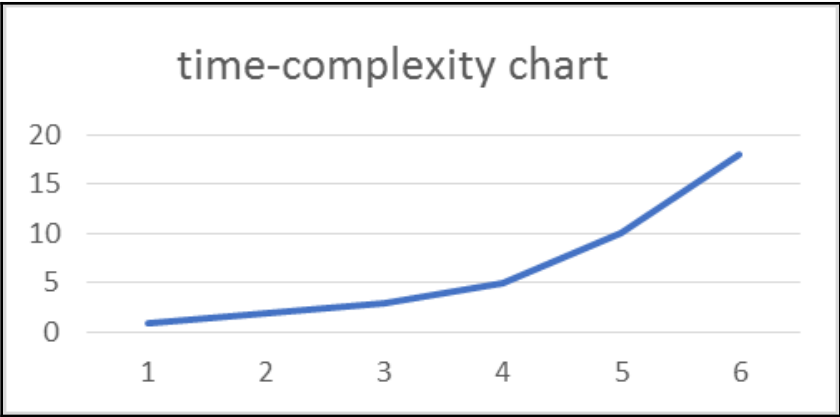
$$\begin{cases} x & x \geq 0 \\ \alpha(e^x - 1) & x < 0 \end{cases}$$
















```

1 import matplotlib.pyplot as plt
2
3 import numpy as np
4
5 def plot_history(history):
6     plt.figure()
7     plt.xlabel('Epoch')
8     plt.ylabel('Mean Abs Error [1000$]')
9     plt.plot(history.epoch, np.array(history.history['mean_absolute_error']),
10             label='Train Loss')
11     plt.plot(history.epoch, np.array(history.history['val_mean_absolute_error']),
12             label = 'Val loss')
13     plt.legend()
14     plt.ylim([0, 15])
15
16 plot_history(network_metadata)

```



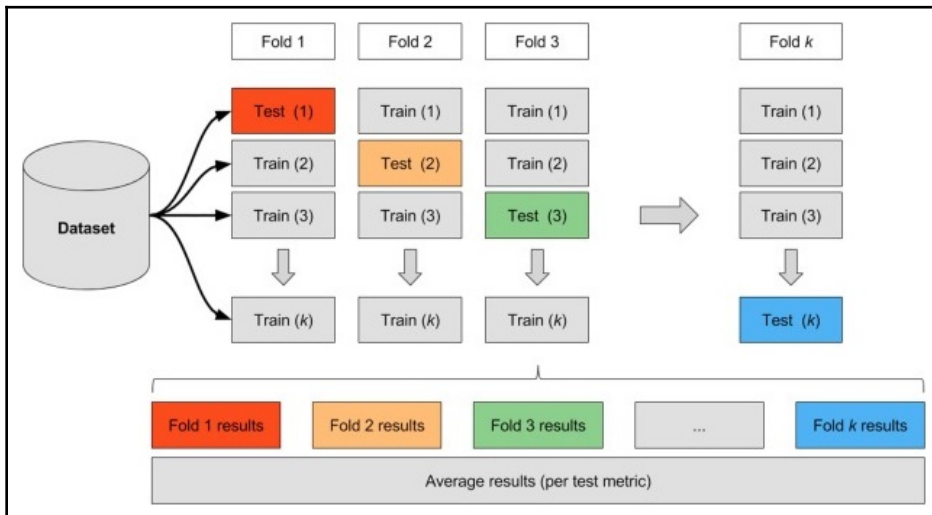

```
1 test_predictions = model.predict(x_test).flatten()
2
3 plt.scatter(y_test, test_predictions)
4 plt.xlabel('True Values [1000$]')
5 plt.ylabel('Predictions [1000$]')
6 plt.axis('equal')
7 plt.xlim(plt.xlim())
8 plt.ylim(plt.ylim())
9 _ = plt.plot([-100, 100], [-100, 100])
```




```



1 error = test_predictions - y_test
2 plt.hist(error, bins = 50)
3 plt.xlabel("Prediction Error [1000$]")
4 _ = plt.ylabel("Count")

```



Chapter 04: Convolutional Neural Networks






FACIAL RECOGNITION


Deep-learning neural networks use layers of increasingly complex rules to categorize complicated shapes such as faces



Layer 1: The computer identifies pixels of light and dark.




Layer 2: The computer learns to identify edges and simple shapes.

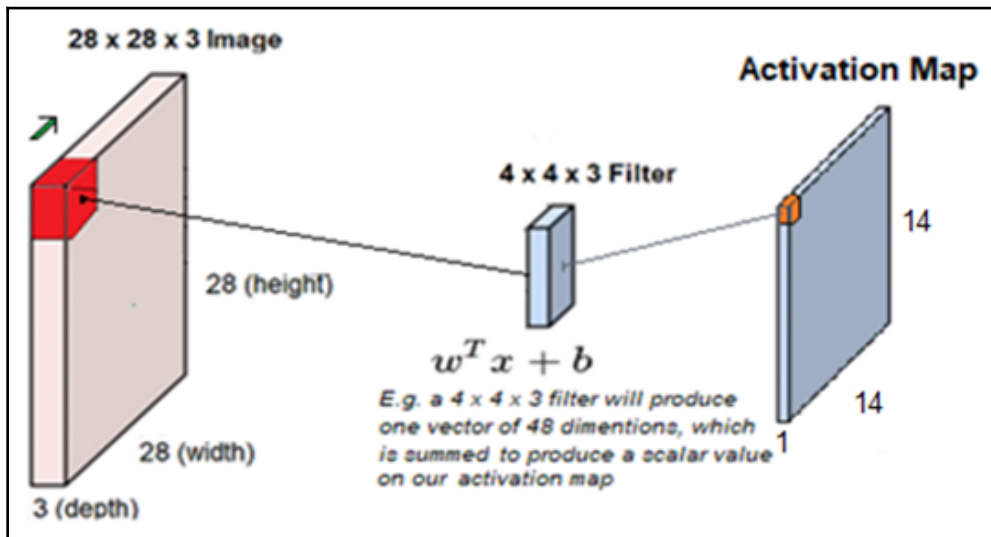








Layer 3: The computer learns to identify more complex shapes and objects.

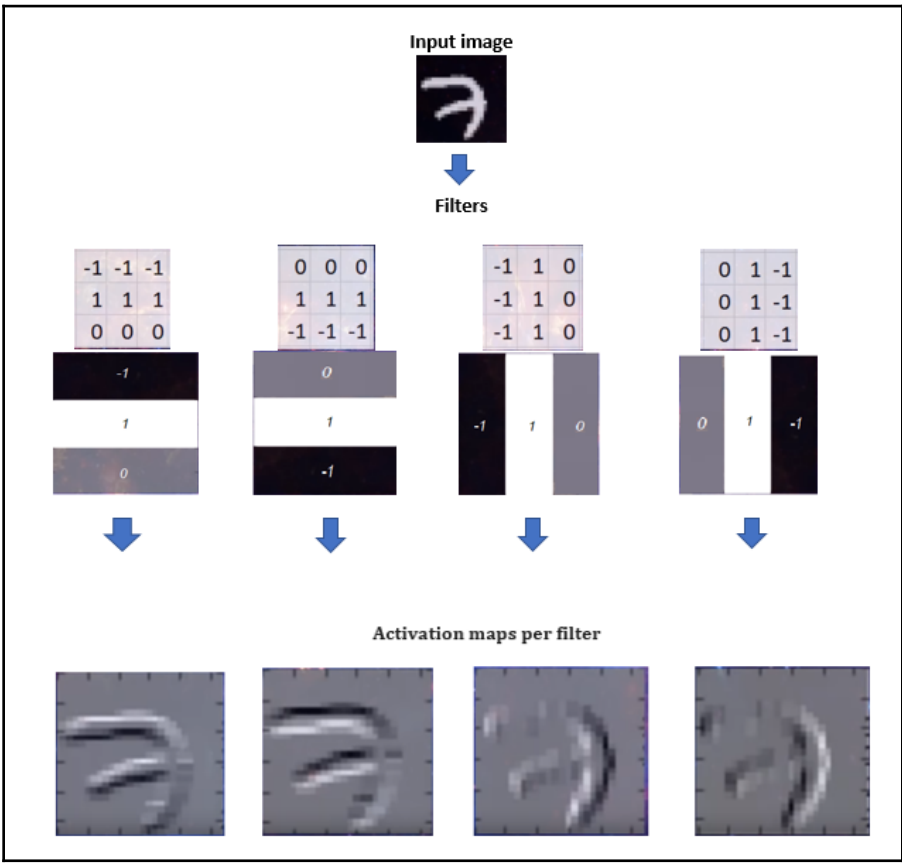



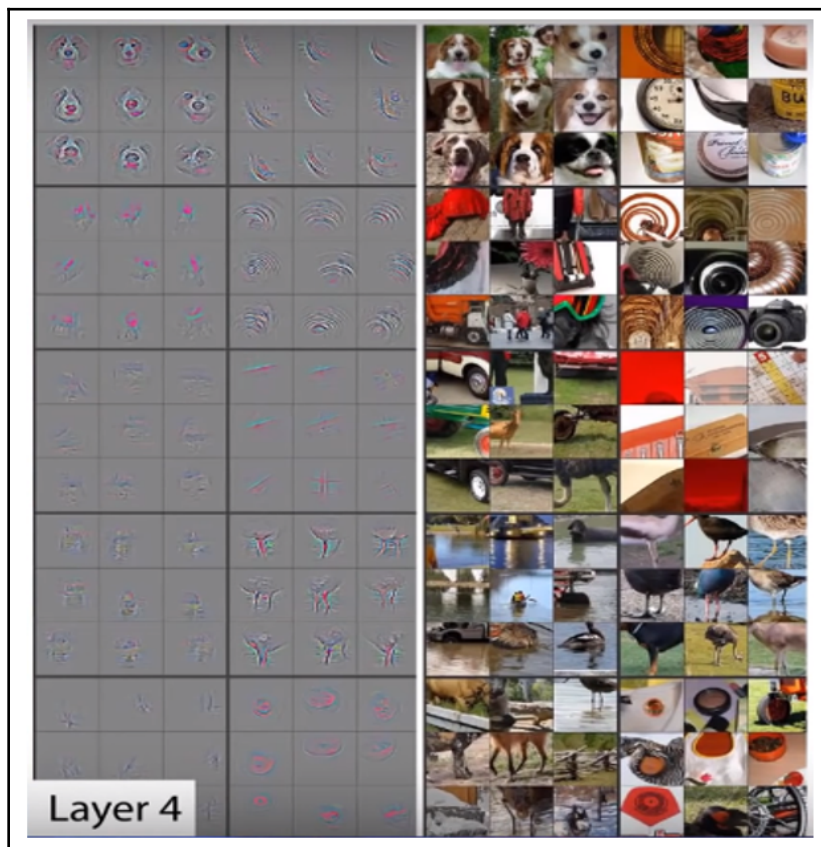
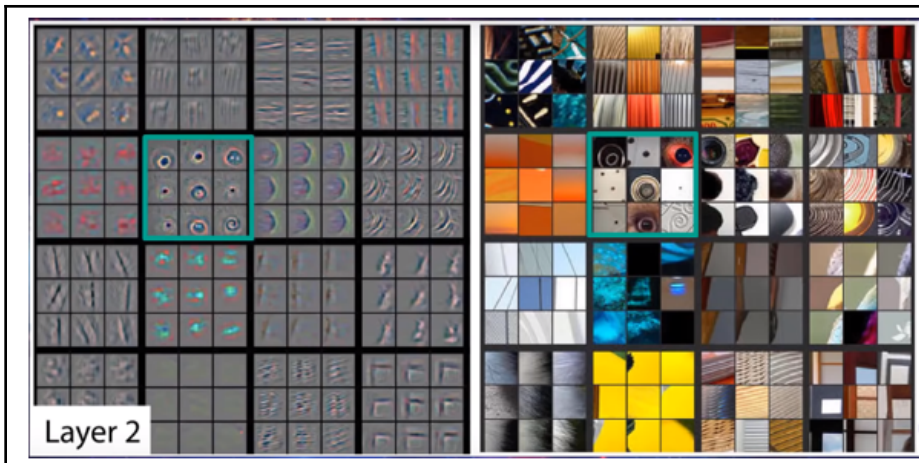
Layer 4: The computer learns which shapes and objects can be used to define a human face.



















```

1 model.summary()

```

Layer (type)	Output Shape	Param #
conv2d_29 (Conv2D)	(None, 64, 64, 16)	1216
batch_normalization_29 (Batch Normalization)	(None, 64, 64, 16)	64
max_pooling2d_29 (MaxPooling2D)	(None, 32, 32, 16)	0
dropout_29 (Dropout)	(None, 32, 32, 16)	0
conv2d_30 (Conv2D)	(None, 32, 32, 32)	12832
batch_normalization_30 (Batch Normalization)	(None, 32, 32, 32)	128
max_pooling2d_30 (MaxPooling2D)	(None, 16, 16, 32)	0
dropout_30 (Dropout)	(None, 16, 16, 32)	0
flatten_15 (Flatten)	(None, 8192)	0
dense_29 (Dense)	(None, 128)	1048704
dense_30 (Dense)	(None, 1)	129

Total params: 1,063,073
 Trainable params: 1,062,977
 Non-trainable params: 96

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

```
1 from keras.callbacks import EarlyStopping
2
3 early_stopping = keras.callbacks.EarlyStopping(monitor='val_loss')
```


```
1 model.fit(X_train, Y_train, validation_data=(X_test, Y_test),
2           epochs=20,
3           batch_size=50,
4           callbacks=[early_stopping])
```

```
Train on 600 samples, validate on 150 samples
Epoch 1/20
600/600 [=====] - 24s 40ms/step - loss: 1.1793 - acc: 0.5583 - val_loss: 0.5761 - val_acc: 0.6267
Epoch 2/20
600/600 [=====] - 21s 35ms/step - loss: 0.5251 - acc: 0.7850 - val_loss: 0.4851 - val_acc: 0.7600
Epoch 3/20
600/600 [=====] - 20s 34ms/step - loss: 0.3586 - acc: 0.8400 - val_loss: 0.3339 - val_acc: 0.8667
Epoch 4/20
600/600 [=====] - 20s 34ms/step - loss: 0.2379 - acc: 0.9117 - val_loss: 0.2766 - val_acc: 0.8800
Epoch 5/20
600/600 [=====] - 20s 34ms/step - loss: 0.1905 - acc: 0.9267 - val_loss: 0.1928 - val_acc: 0.9400
Epoch 6/20
600/600 [=====] - 20s 34ms/step - loss: 0.1688 - acc: 0.9250 - val_loss: 0.1650 - val_acc: 0.9600
Epoch 7/20
600/600 [=====] - 20s 34ms/step - loss: 0.1287 - acc: 0.9533 - val_loss: 0.1554 - val_acc: 0.9400
Epoch 8/20
600/600 [=====] - 21s 35ms/step - loss: 0.0968 - acc: 0.9600 - val_loss: 0.2002 - val_acc: 0.8800
<keras.callbacks.History at 0x2baa51ba7b8>
```

```
1 # Predict the test set results
2 Y_pred = model.predict_classes(X_test)
```

```
1 from sklearn.metrics import accuracy_score, confusion_matrix, recall_score, precision_score, f1_score
2
3
4 print ("test accuracy: %s" %accuracy_score(Y_test, Y_pred))
5 print ("precision: %s" %precision_score(Y_test, Y_pred))
6 print ("recall: %s" %recall_score(Y_test, Y_pred))
7 print ("f1 score: %s" %f1_score(Y_test, Y_pred))
8
9
```

```
test accuracy: 0.88
precision: 0.9852941176470589
recall: 0.7976190476190477
f1 score: 0.8815789473684211
```



```

1 model.save('C:/Users/npurk/Desktop/Chapter_3_CNN/smile_detector.h5py')
2 model = keras.models.load_model('C:/Users/npurk/Desktop/Chapter_3_CNN/smile_detector.h5py')

```



```

1 input_image = X_test[8]
2 input_image = np.expand_dims(input_image, axis=0)
3 print(input_image.shape)

```

(1, 64, 64, 3)

```
1 #retrieve a blank multi-output model from the functional API
2 from keras.models import Model
3
4 #retrieve layer outputs for layers in previously trained sequential model
5 layer_outputs_smile_detector = [layer.output for layer in model.layers[:]]
6
7 #define multi-output model that takes input image tensors and outputs intermediate layer activations
8 multioutput_model = Model(inputs=model.input, outputs=layer_outputs)
9
10 #Generate activation tensors of intermediate layers
11 activations = multioutput_model.predict(input_image)
```

```
1 print('number of layers:', len(activations))
2 print('Datatype:', type(activations[0]))
3 print('1st layer output shape:', activations[0].shape)
4 print('4th layer output shape:', activations[3].shape)
5 print('8th layer output shape:', activations[7].shape)
```


```
number of layers: 8
Datatype: <class 'numpy.ndarray'>
1st layer output shape: (1, 64, 64, 16)
4th layer output shape: (1, 32, 32, 16)
8th layer output shape: (1, 16, 16, 32)
```



```
1 model.summary()
```

Layer (type)	Output Shape	Param #
conv2d_29 (Conv2D)	(None, 64, 64, 16)	1216
batch_normalization_29 (Batch Normalization)	(None, 64, 64, 16)	64
max_pooling2d_29 (MaxPooling2D)	(None, 32, 32, 16)	0
dropout_29 (Dropout)	(None, 32, 32, 16)	0
conv2d_30 (Conv2D)	(None, 32, 32, 32)	12832
batch_normalization_30 (Batch Normalization)	(None, 32, 32, 32)	128
max_pooling2d_30 (MaxPooling2D)	(None, 16, 16, 32)	0
dropout_30 (Dropout)	(None, 16, 16, 32)	0
flatten_15 (Flatten)	(None, 8192)	0
dense_29 (Dense)	(None, 128)	1048704
dense_30 (Dense)	(None, 1)	129

=====
Total params: 1,063,073
Trainable params: 1,062,977
Non-trainable params: 96
=====



$$\frac{\partial \text{output}}{\partial \text{input}}$$

```

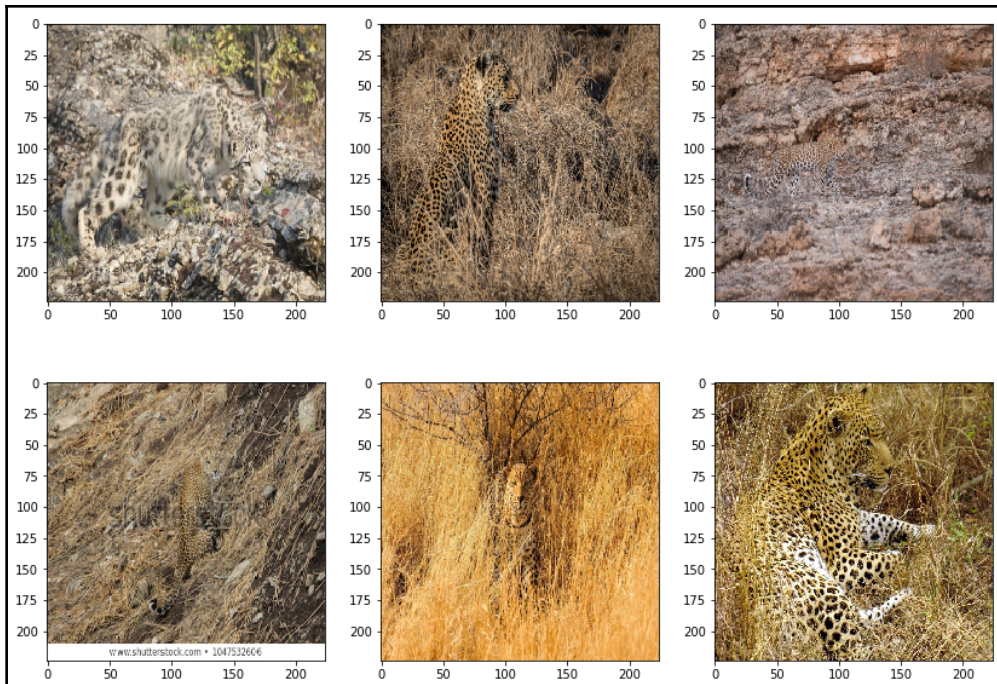
1 from keras.applications import ResNet50
2 from vis.utils import utils
3 from keras import activations
4
5 # Hide warnings on Jupyter Notebook
6 import warnings
7 warnings.filterwarnings('ignore')
8
9 # Build the ResNet50 network with ImageNet weights
10 model = ResNet50(weights='imagenet', include_top=True)
11
12 # Utility to search for layer index by name.
13 # Alternatively we can specify this as -1 since it corresponds to the last layer.
14 layer_idx = utils.find_layer_idx(model, 'fc1000')
15
16 # Swap softmax with linear
17 model.layers[layer_idx].activation = activations.linear
18 model = utils.apply_modifications(model)

```

```

1 from vis.utils import utils
2 from matplotlib import pyplot as plt
3 %matplotlib inline
4 plt.rcParams['figure.figsize'] = (14, 10)
5
6 img1 = utils.load_img('C:/Users/npurk/Pictures/Saved Pictures/leo_1.jpg', target_size=(224, 224))
7 img2 = utils.load_img('C:/Users/npurk/Pictures/Saved Pictures/leo_2.jpg', target_size=(224, 224))
8 img3 = utils.load_img('C:/Users/npurk/Pictures/Saved Pictures/leo_3.jpg', target_size=(224, 224))
9 img4 = utils.load_img('C:/Users/npurk/Pictures/Saved Pictures/leo_4.jpg', target_size=(224, 224))
10 img5 = utils.load_img('C:/Users/npurk/Pictures/Saved Pictures/leo_5.jpg', target_size=(224, 224))
11 img6 = utils.load_img('C:/Users/npurk/Pictures/Saved Pictures/leo_6.jpg', target_size=(224, 224))
12
13 f, ax = plt.subplots(nrows=2, ncols=3)
14 ax[0, 0].imshow(img1)
15 ax[0, 1].imshow(img2)
16 ax[0, 2].imshow(img3)
17 ax[1, 0].imshow(img4)
18 ax[1, 1].imshow(img5)
19 ax[1, 2].imshow(img6)

```



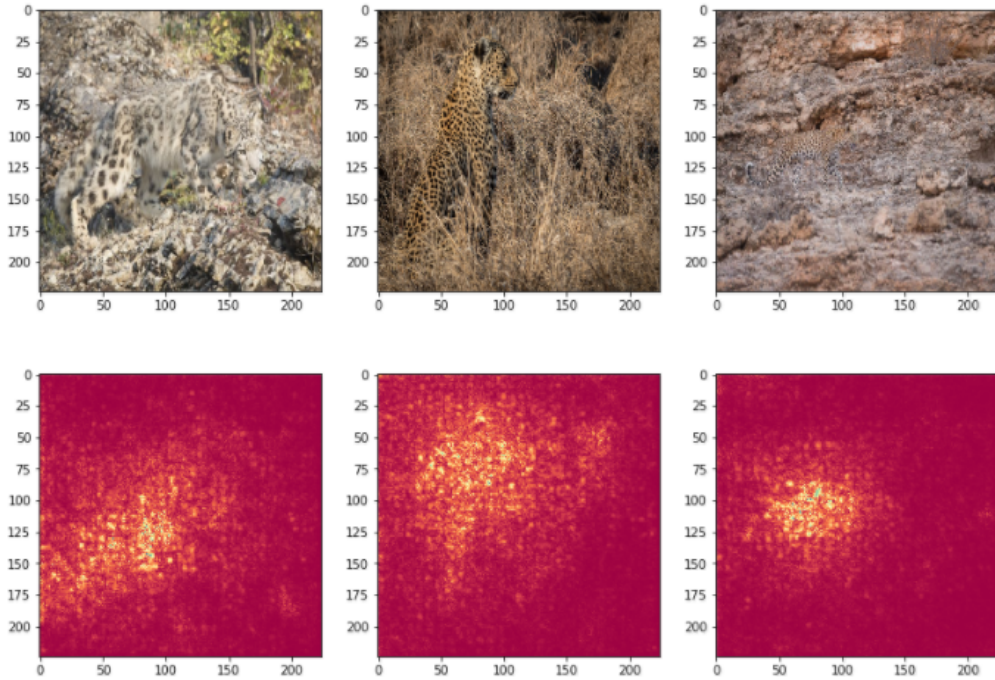
```

1  from vis.visualization import visualize_saliency
2  from vis.utils import utils
3
4
5  leopards = [img1,img2,img3,img4,img5,img6]
6  # Utility to search for Layer index by name.
7  # Alternatively we can specify this as -1 since it corresponds to the Last Layer.
8  layer_idx = utils.find_layer_idx(model, 'fc1000')
9
10 gradients=[]
11 for i, img in enumerate(leopards):
12     # 288 is the imagenet index corresponding to `leopard, panthera pardus`
13     grads = visualize_saliency(model, layer_idx, filter_indices=288, seed_input=cougars[i])
14     gradients.append(grads)

```

```
1 f, ax = plt.subplots(2, 3)
2 ax[0,0].imshow(leopards[0])
3 ax[0,1].imshow(leopards[1])
4 ax[0,2].imshow(leopards[2])
5 ax[1,0].imshow(gradients[0], cmap='Spectral')
6 ax[1,1].imshow(gradients[1], cmap='Spectral')
7 ax[1,2].imshow(gradients[2], cmap='Spectral')
```

<matplotlib.image.AxesImage at 0x1df4ba81208>




```

1 import matplotlib.cm as cm
2 from vis.visualization import visualize_cam, overlay
3 from keras import activations
4
5 # Find the fully connected output Layer
6 layer_idx = utils.find_layer_idx(model, 'fc1000')
7 #Find the penultimate convolutional Layer
8 final_conv_layer = utils.find_layer_idx(model, 'res5c_branch2c')
9
10
11 plt.figure()
12 f, ax = plt.subplots(1, 2)
13
14 for i, img in enumerate([img3, img6]):
15
16     grads = visualize_cam(model, #ResNet50 model on ImageNet weights
17                           seed_input=img, #Image with Leopard
18                           filter_indices=288, #Filter Index for Leopard in imagenet dataset
19                           layer_idx=layer_idx, #Last fully connected Layer
20                           penultimate_layer_idx=final_conv_layer) #Penultimate convolutional Layer
21
22 # overlay the heatmap on top of the original image.
23 jet_heatmap = np.uint8(cm.jet(grads)[..., :3] * 255)
24 ax[i].imshow(overlay(jet_heatmap, img))

```

<Figure size 1008x720 with 0 Axes>



```
1
2 from keras.preprocessing.image import img_to_array
3 # convert the image pixels to a 4-D numpy array
4 img_predict = img_to_array(img1)
5 img_predict = np.expand_dims(img_predict, axis=0)
6 img_predict.shape
7
```

(1, 224, 224, 3)

```
1 yhat = model.predict(img_predict)
2
3 from keras.applications.resnet50 import decode_predictions
4 # convert the probabilities to class labels
5 predictions = decode_predictions(yhat, top=1000)
6 # retrieve the most likely result, e.g. highest probability
7 label = predictions[0][0]
8 labels = predictions[0][:]
9 # print the classification
10 print('%s (%.2f%%)' % (label[1], label[2]*100))
```

snow_leopard (1080.04%)

```
1 labels[:5]
```


```
[('n02112875', 'snow_leopard', 10.800396),
 ('n02130308', 'cheetah', 9.888457),
 ('n02100735', 'English_setter', 8.046139),
 ('n02128385', 'leopard', 7.3484097),
 ('n02110341', 'dalmatian', 6.5762362)]
```

```
1 from keras.applications import VGG16
2 from vis.utils import utils
3 from keras import activations
4
5 # Build the VGG16 network with ImageNet weights
6 model = VGG16(weights='imagenet', include_top=True)
7
8 # Utility to search for layer index by name.
9 # Alternatively we can specify this as -1 since it corresponds to the last layer.
10 layer_idx = utils.find_layer_idx(model, 'predictions')
11
12 # Swap softmax with linear
13 model.layers[layer_idx].activation = activations.linear
14 model = utils.apply_modifications(model)
```





```
1 from vis.visualization import visualize_activation
2
3 from matplotlib import pyplot as plt
4 %matplotlib inline
5 plt.rcParams['figure.figsize'] = (18, 6)
6
7 # 20 is the imagenet category for 'ouzel'
8 img = visualize_activation(model, layer_idx, filter_indices=[288])
9 plt.imshow(img)
```

<matplotlib.image.AxesImage at 0x1c381974f60>

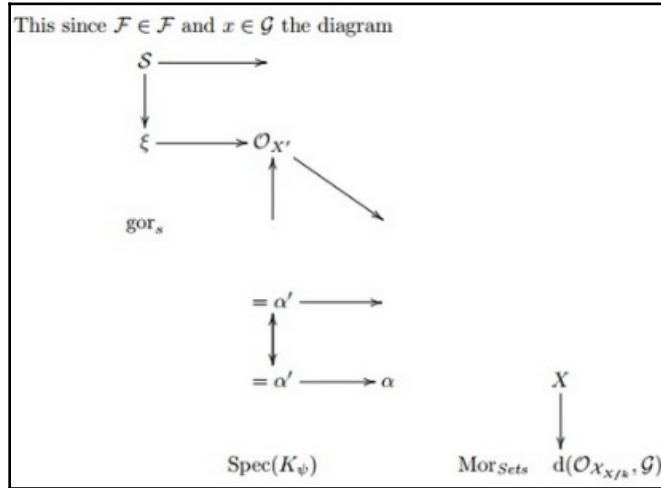


```
1 img = visualize_activation(model, layer_idx, filter_indices=288, max_iter=500, verbose=True)
```





Chapter 05: Recurrent Neural Networks



is a limit. Then \mathcal{G} is a finite type and assume S is a flat and \mathcal{F} and \mathcal{G} is a finite type f_* . This is of finite type diagrams, and

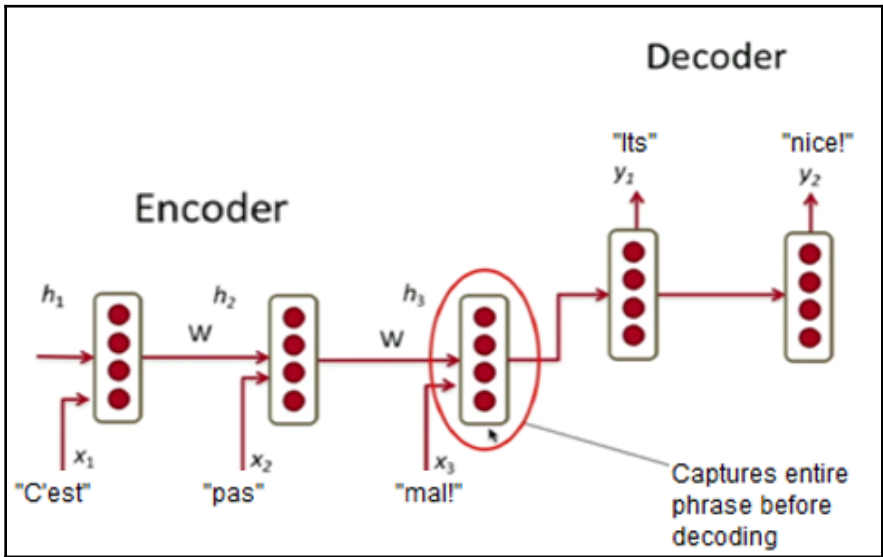

- the composition of \mathcal{G} is a regular sequence,
- $\mathcal{O}_{X'}$ is a sheaf of rings.


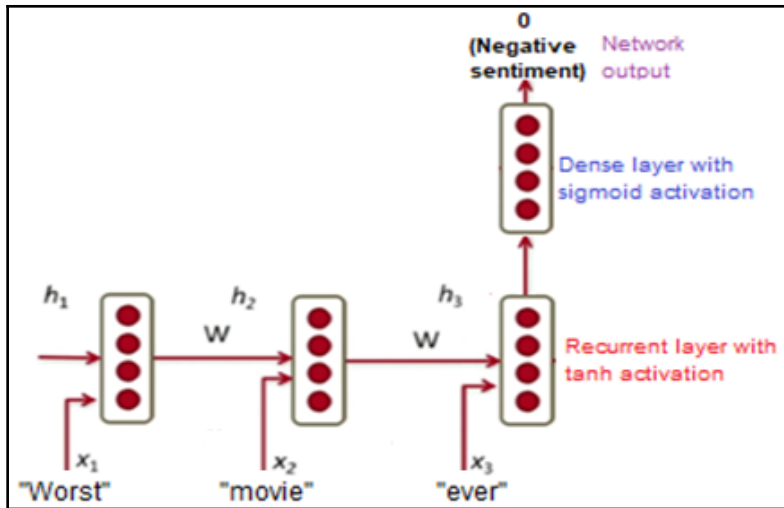
```

1 | verify_predictions(22)
Network falsely predicts that review 22 is negative
(None,
 array([0.16930683], dtype=float32),
 "? how managed to avoid attention remains a mystery a potent mix of comedy and crime this one takes chances where tarantino pl
ays it safe with the hollywood formula the risks don't always pay off one character in one sequence comes off ? silly and falls
flat in the lead role thomas jane gives a wonderful and complex performance and two brief appearances by mickey rourke hint at
the high potential of this much under and mis used actor here's a director one should keep one's eye on")

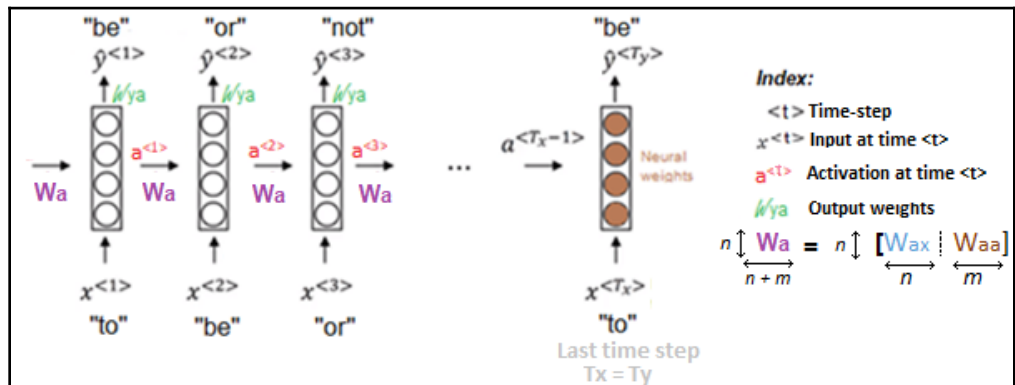
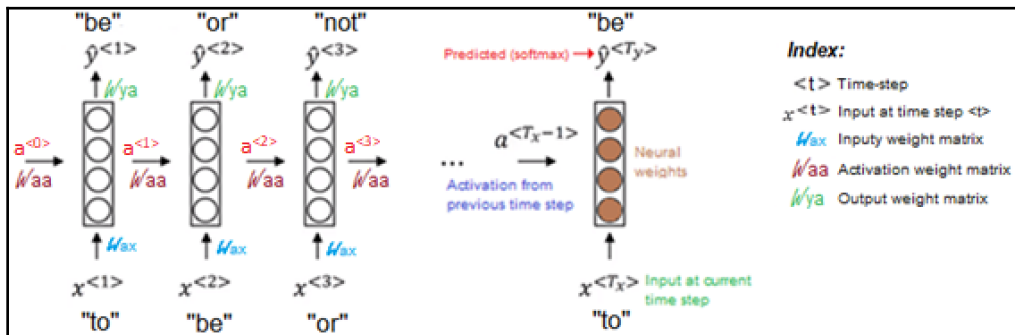
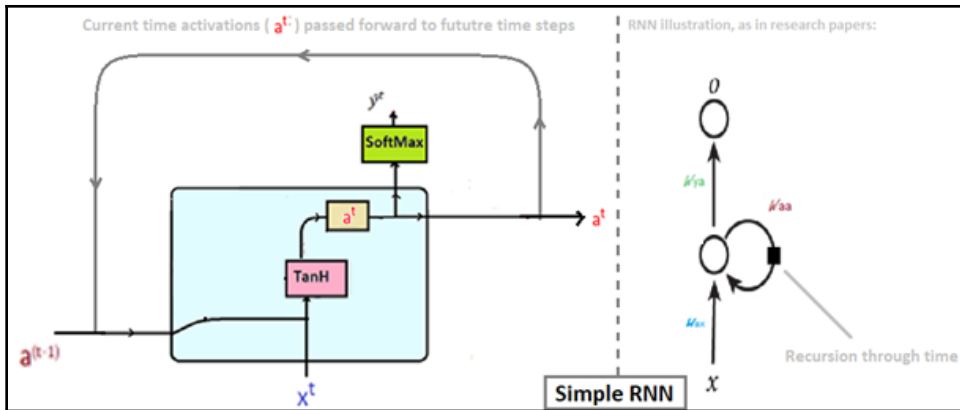
```

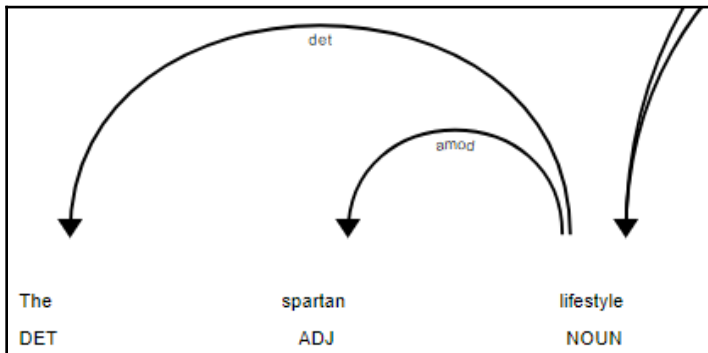


Advantages	Drawbacks
<ul style="list-style-type: none"> • Possibility of processing input of any length • Model size not increasing with size of input • Computation takes into account historical information • Weights are shared across time 	<ul style="list-style-type: none"> • Computation being slow • Difficulty of accessing information from a long time ago • Cannot consider any future input for the current state

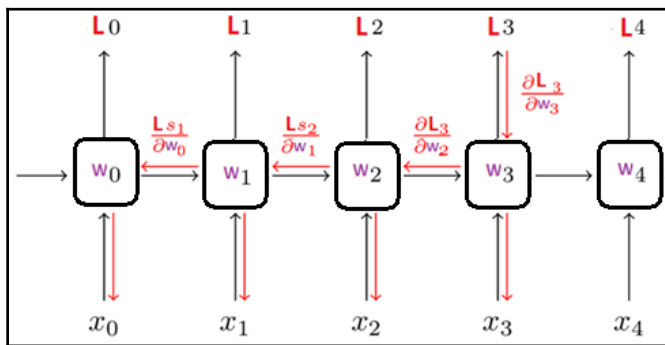





	Misc.	Words describing image	Sentiment of review	Translated version of input	Words matching sound
Output type:	Misc.	Words describing image	Sentiment of review	Translated version of input	Words matching sound
Layer Schema:					
Input type:	Misc.	Image	review of words	Foreign language input	Sequence of sounds
Use case:	Basic Neural Network	RNN for image captioning	RNN for sentiment analysis	RNN for machine language translation	RNN for parking car or performing speech recognition
Type:	One-to-One	One-to-Many	Many-to-One	Many-to-Many (semi-sync)	Many-to-many(full-sync)







$$\begin{aligned} \frac{\partial L_3}{\partial W} &= \frac{\partial L_3}{\partial y_3} \cdot \frac{\partial y_3}{\partial h_2} \cdot \frac{\partial h_2}{\partial W} \\ &= \sum_{t=0}^2 \frac{\partial L_3}{\partial y_3} \cdot \frac{\partial y_3}{\partial h_2} \cdot \frac{\partial h_2}{\partial h_t} \cdot \frac{\partial h_t}{\partial W} \\ &= \sum_{t=0}^2 \frac{\partial L_3}{\partial y_3} \cdot \frac{\partial y_3}{\partial h_2} \cdot \left(\prod_{j=t+1}^2 \frac{\partial h_j}{\partial h_{j-1}} \right) \cdot \frac{\partial h_t}{\partial W} \end{aligned}$$

```



from keras import optimizers


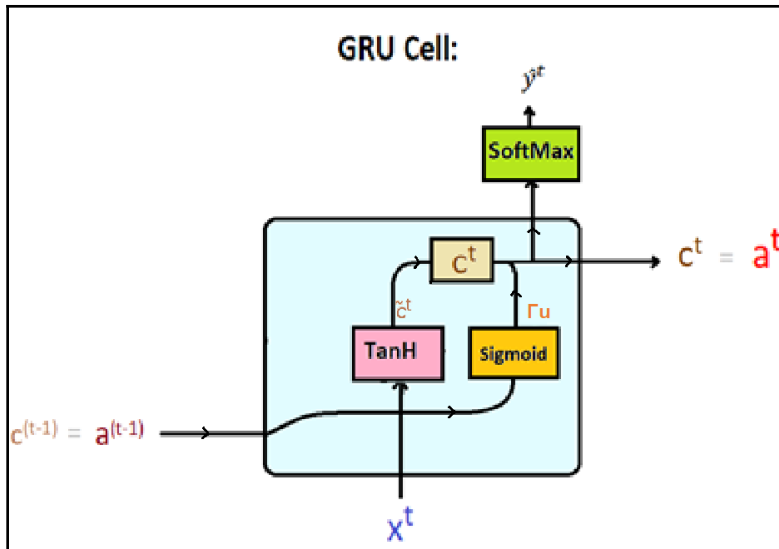
# All parameter gradients will be clipped to
# a maximum norm of 1.
sgd = optimizers.SGD(lr=0.01, clipnorm=1.)

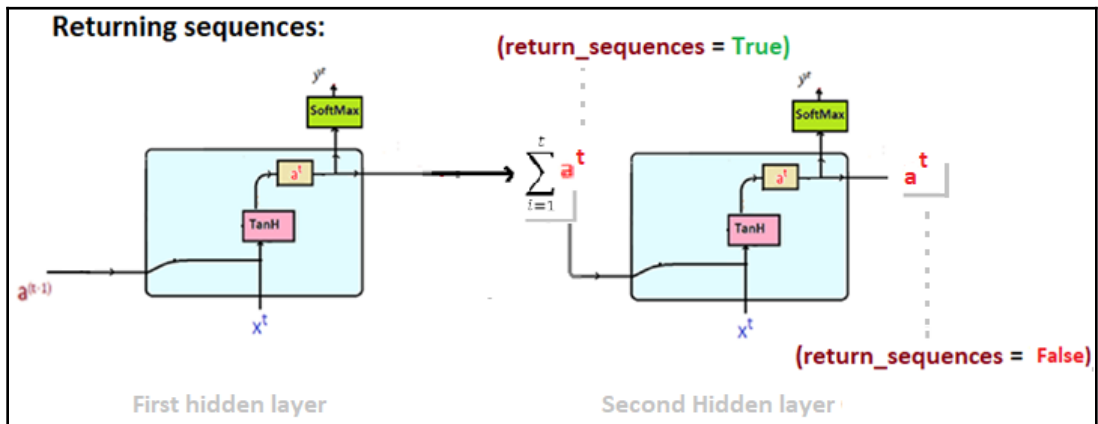
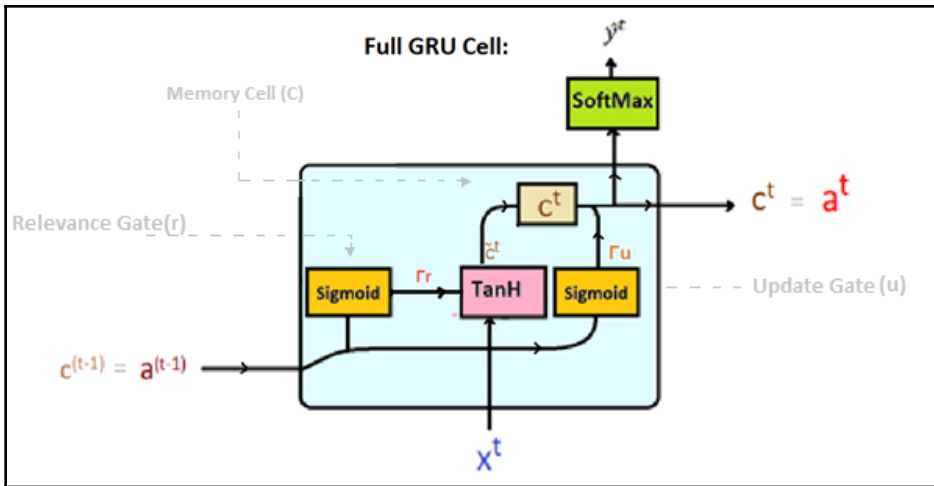
from keras import optimizers

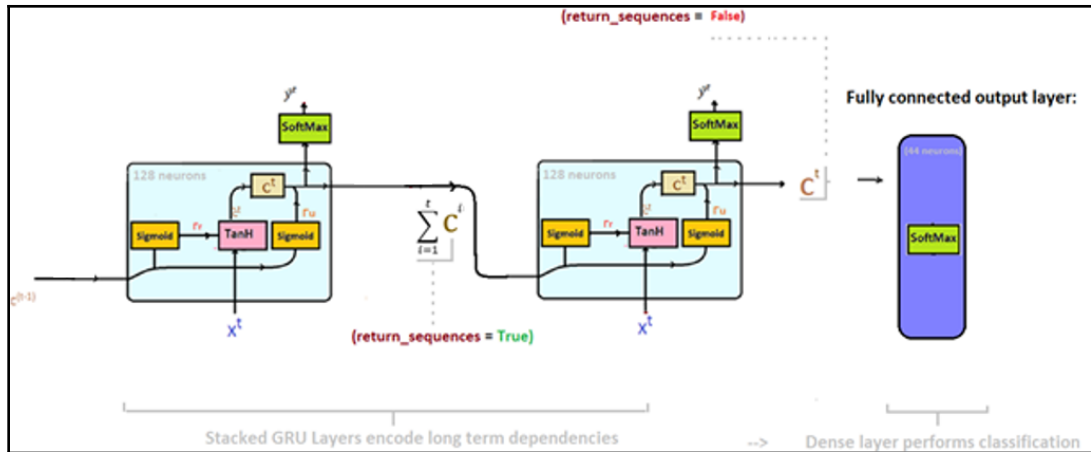
# All parameter gradients will be clipped to
# a maximum value of 0.5 and
# a minimum value of -0.5.
sgd = optimizers.SGD(lr=0.01, clipvalue=0.5)

```








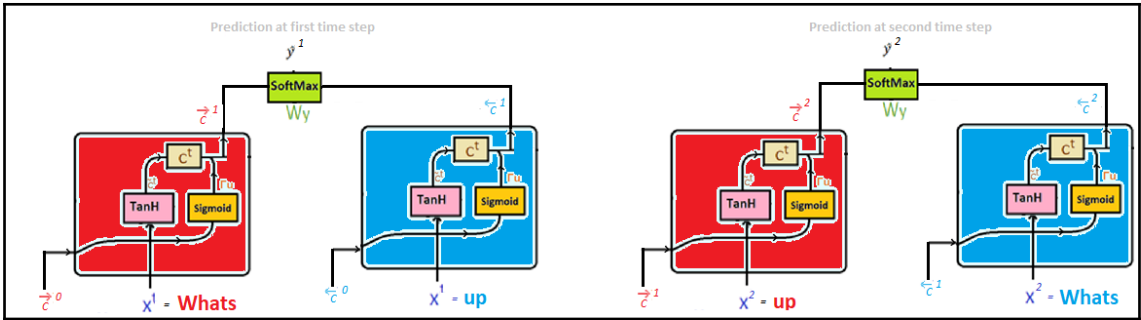


Normal order: \rightarrow
 The spartan marched forward despite the obstacles thrown at him

Reverse Order: \leftarrow
 him at throw obstacles the despite forward marched spartan The



Time-step->	1	2
Ordered sequence->	Whats	up
Reverse sequence->	up	Whats



```

1 test_models(all_models, epochs=1)

Initiating compilation...
Compiled: SimpleRNN_model
Epoch 1/1
166730/166730 [=====] - 33s 199us/step - loss: 2.2307

----- Generating text after Epoch: 0
----- *Sampling Threshold* : 0.3
----- *Input sequence to generate from* : "nd mine more of the same beauty tha"
nd mine more of the same beauty that in the ther and he wert will , and , th me the me the of in the fere , the port ou hame th
e pance , and he fore if in the wing the late , and whe ham . ham . and with me the heme , and , and lith , and me the wort and
my the for . in the peate , and ntte fore the fere , and list , and lord ham . in the ponge , the ny lord , and lore , and with
ma that har . of the fere , and , and and , and will

----- *Sampling Threshold* : 0.5
----- *Input sequence to generate from* : "nd mine more of the same beauty tha"
nd mine more of the same beauty thas and in the hinen . wor . in were more , in the lith , gham . not the sering of of her . in
g the galle , and here pall the peuae , ard hat . lowe , and ham . in the senge henat , and hame the harresen . and mente ing t
he wenthin the foll the thend in this ruen the helleere , for . me that ding ous andue the inge in in whe ham . gore the , and
you houd and and , the perte in s and lyour will the o

----- *Sampling Threshold* : 0.7
----- *Input sequence to generate from* : "nd mine more of the same beauty tha"
nd mine more of the same beauty that in the lour inghe ding mist on wire ke by . whe is aun that ine s of hiant hat sheraund qu
einel of s abe shower , thouthes ind . than ' nne so deuenger for mose so fer ancede the mande not mower ham . or saplee tha ki
ng the e tham . io h mest is phis ne came to the becheno , thit with mene hepere foll the forishis nor wing in shin with that .
ineu . not me hor , and , in thre tone ming , thet i

----- *Sampling Threshold* : 1.0
----- *Input sequence to generate from* : "nd mine more of the same beauty tha"
nd mine more of the same beauty than and audy myring , fouge ' ses hi kere thit rnse andes nuthieers ched erit my yous afco tor
n tis ine myourille , py lowd stbenen apowee prayet , all , tf feryeakestring ot hey for allwingor this at ganse . (et is myof
ltaeda dod . god iphon so h ctay fither ffothis lenne , nsowhetr share it ser the weeles dane , orabughis szall , knsdge llot ?
hea , so fother hamll : th then ther phat , and 'it b

----- *Sampling Threshold* : 1.2
----- *Input sequence to generate from* : "nd mine more of the same beauty tha"
nd mine more of the same beauty thangrdi on welt thees cnoune hel nSawit) hon serusl te hent ay h nereriolr thes ghe : eresdoin
kfe houchie ht sdour sanngrraune ? and bede : ti malrr .ee fred fare wond oy sabcki harereus , s aukurd oy vert yhe yestineckink
a , and hourd fay - on ' mtseed : at hiubge , eu) . yountomrend , nfares ou hase mowes of thet aptham v nimhte winh' wholes ye
at th gue , onze , avse nownd goar no' cere pusin the

Layer (type)                Output Shape                Param #
-----
simple_rnn_11 (SimpleRNN)    (None, 128)                 22016
-----
dense_20 (Dense)           (None, 43)                  5547
-----
Total params: 27,563
Trainable params: 27,563
Non-trainable params: 0
-----

```

```

----- Generating text after Epoch: 1

----- *Sampling Threshold* : 0.3

----- *Input sequence to generate from* : "n , like sweet bells iangled out of "

n , like sweet bells iangled out of the restion , my lord , and the pard , and the more of the forth , and the parse hamlet , and
the beare , the pray his seene the prease , the sonce , the does , and the mant , and the seene the bodne , i am the man his hau
e his beaue , and dis a passe , the brought , and the growne , and the read , and the seece , and for the father , but the fathe
r , and the does the but his hause so , i part haml

----- *Sampling Threshold* : 0.5

----- *Input sequence to generate from* : "n , like sweet bells iangled out of "

n , like sweet bells iangled out of he purpe , and loce , his most pol . but is the soule , and but the ringed and my frace , or
the clowes , and the forle , and the riplest hamlet : i shall bor ' s to the for his beare , the drowne of resonance , that i pol
. he well he dous it of the clayer . my lord , and the in passion of will his pare his as ham . i makes of deare , and reane for
the fire ic horrow of hing , how she shall heare ,

----- *Sampling Threshold* : 0.7

----- *Input sequence to generate from* : "n , like sweet bells iangled out of "

n , like sweet bells iangled out of , and we or ists to deuitly doe , that most none all is at this at it enter of your ham . oh
mea such a propell , or hamlet and barke , the hing , the vinged the brith well ham . it not of beare play , and soule good meto
' s thing take , and expeare , like the king , the king , the seat your : but i ampees heart of him loue , no more ham . i hau
e his benner : as the grace , i am bet vndon ' d w

----- *Sampling Threshold* : 1.0

----- *Input sequence to generate from* : "n , like sweet bells iangled out of "

n , like sweet bells iangled out of who ? other his begne ? leouer grooke madge , the dow , dor heare indeed hum tudre oones fro
m it denstrack to screcent , but woud light ? buards to i hor as ' l : yet they to , haue so will to cleames , that is , and hi
s not as off the ringes trutinight vndition of thing does of heare , may it he ray fither of are intme till aest it onelles art
ham . i warr there ? rospinisbconouor ? it wernyir

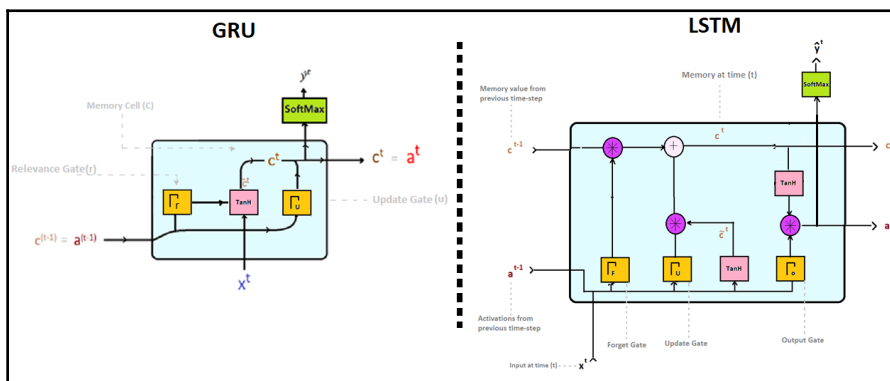
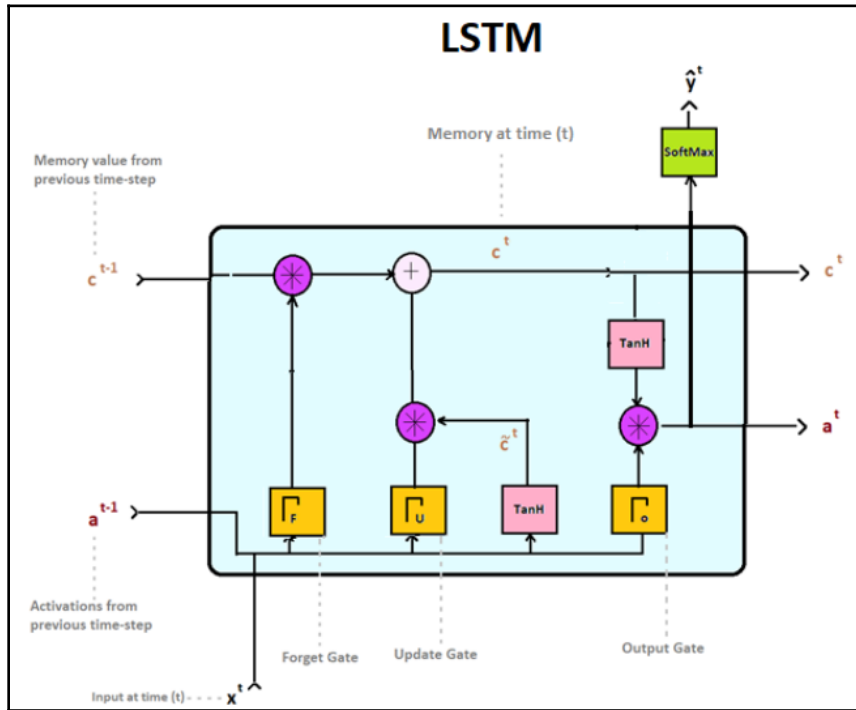
----- *Sampling Threshold* : 1.2

----- *Input sequence to generate from* : "n , like sweet bells iangled out of "

n , like sweet bells iangled out of & ineperues louth me ? yrule , you ; gucch fick hinge of endirg of quiettade mad , the won hi
s put deatter it sonder did naturte ; aglioredkes stachet oyce haflef off rone this liked keeke whes must to - mannom : 'f he to
colturalboustlement ingenoyge into is ammeare it nothers treabersed but hea . buboken itgins . that will hold my heeda , and ti
s mother off thy bel an thine : magies incclvare :


```

Chapter 06: Long Short-Term Memory Networks



$$c^t = (\Gamma_u * e^t) + [(1 - \Gamma_u) * c^{t-1}]$$

Update Gate



$$c^t = (\Gamma_u * e^t) + (\Gamma_f * c^{t-1})$$

Forget Gate

Update Gate

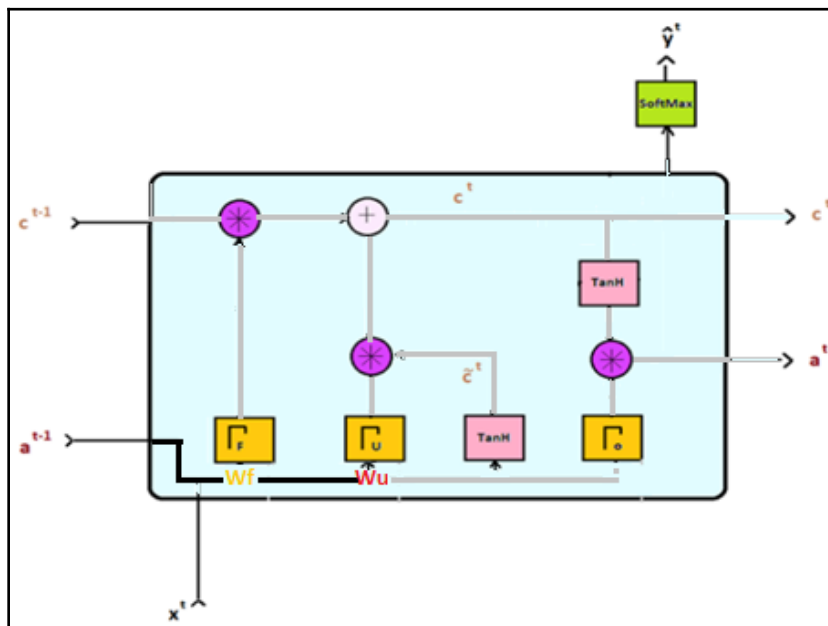
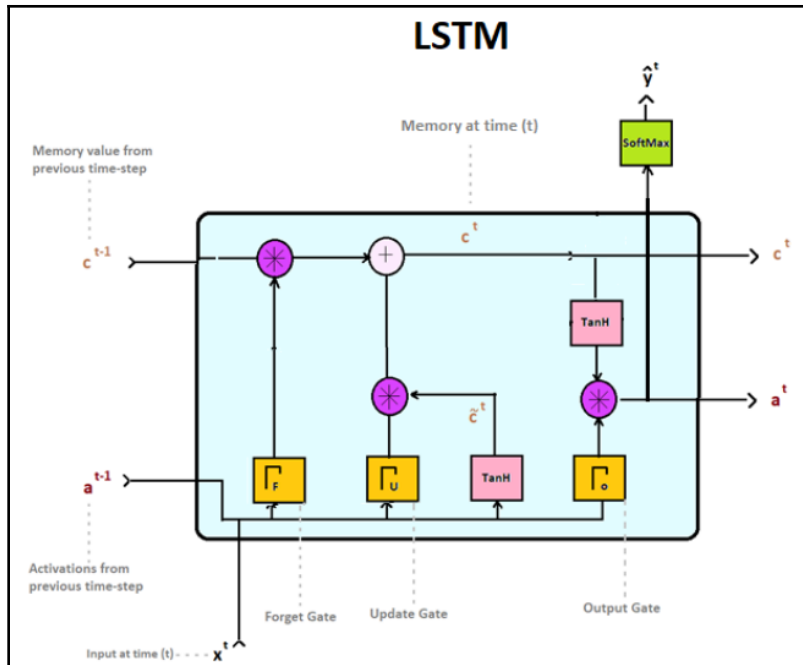
$$c^t = (\Gamma_u * e^t) + [(1 - \Gamma_u) * c^{t-1}]$$

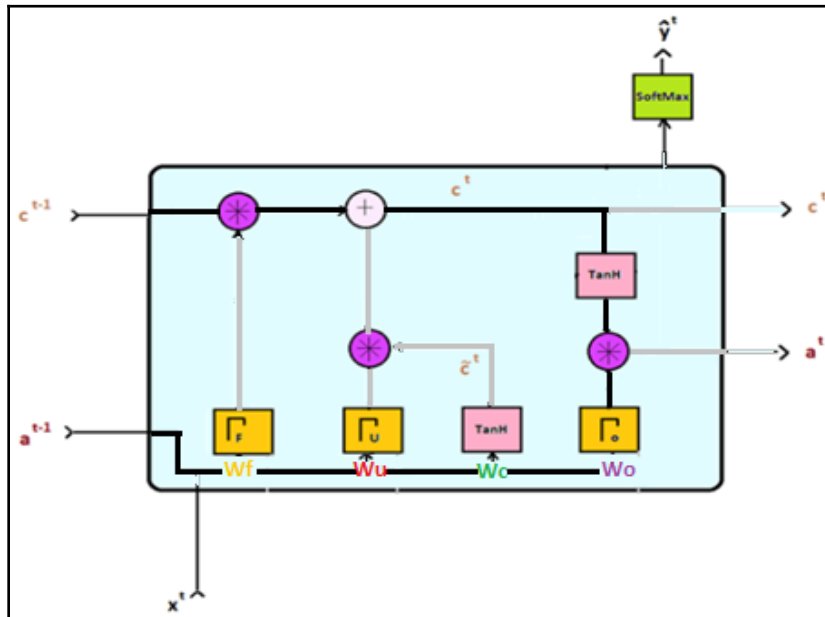

Update Gate

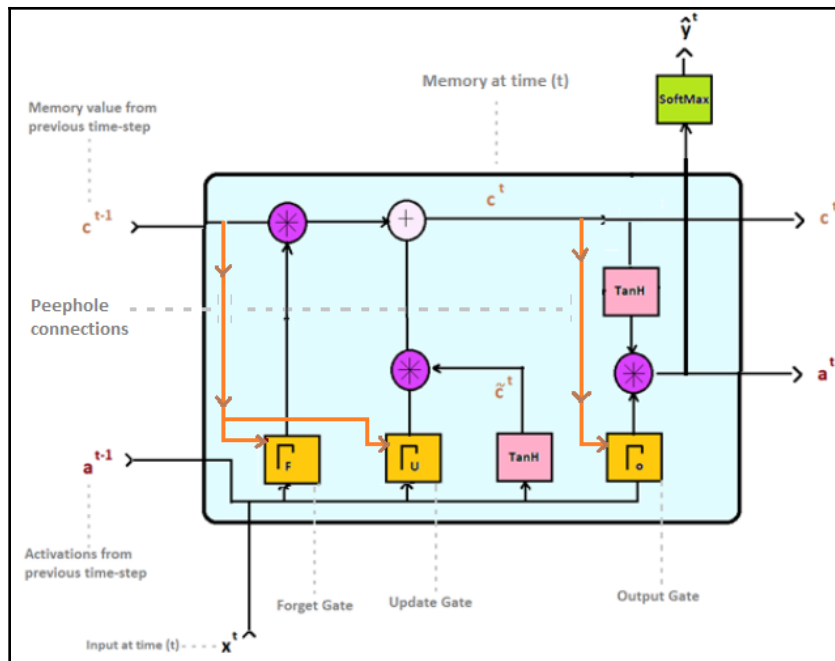
$$c^t = (\Gamma_u * e^t) + (\Gamma_f * c^{t-1})$$

Forget Gate

Update Gate









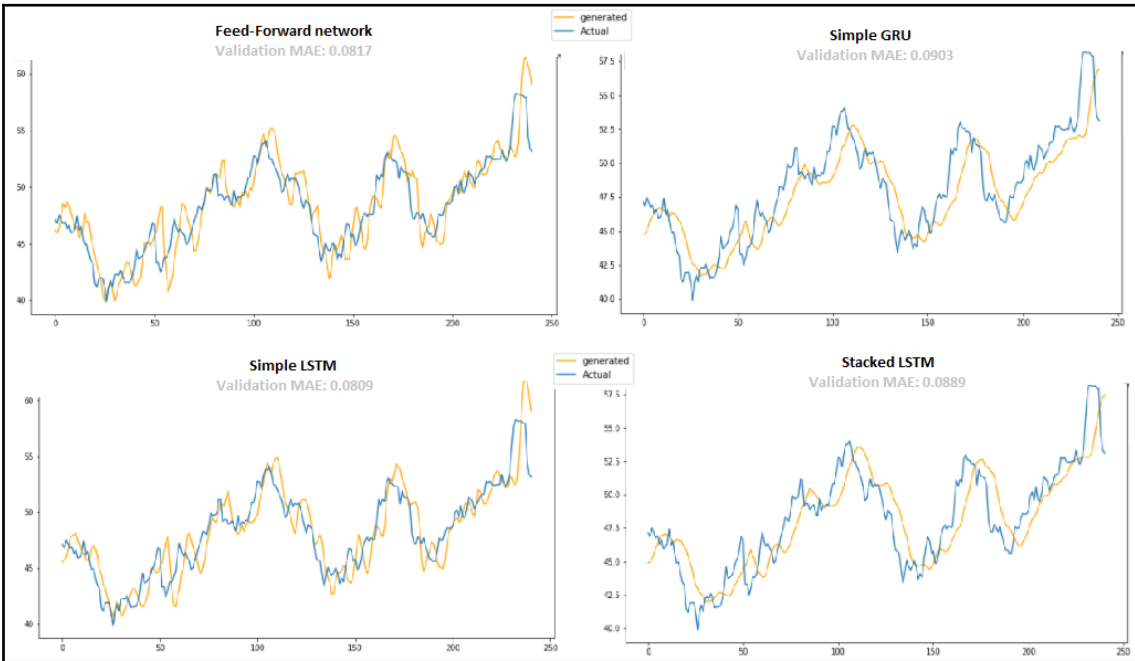



Out[2]:

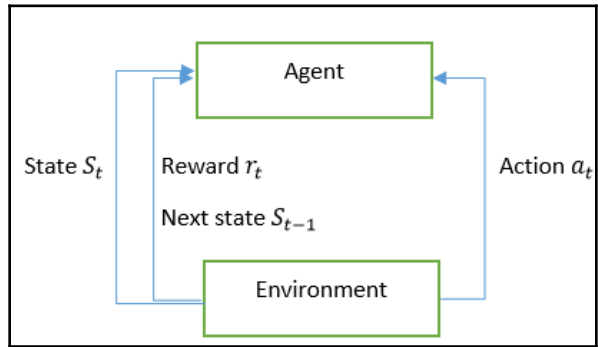
	date	open	high	low	close	volume	Name
0	2013-02-08	15.07	15.12	14.63	14.75	8407500	AAL
1	2013-02-11	14.89	15.01	14.26	14.46	8882000	AAL
2	2013-02-12	14.45	14.51	14.10	14.27	8126000	AAL
3	2013-02-13	14.30	14.94	14.25	14.66	10259500	AAL
4	2013-02-14	14.94	14.96	13.16	13.99	31879900	AAL







Chapter 07: Reinforcement Learning with Deep Q-Networks



Initial state of environment: 204

```

+-----+
|R: | : :G|
| : : : : |
| : : : : |
| | : | : |
|Y| : |B: |
+-----+

```

```

+-----+
|R: | : :G|
| : : : : |
| : : : : |
| | : | : |
|Y| : |B: |
+-----+
(South)

```

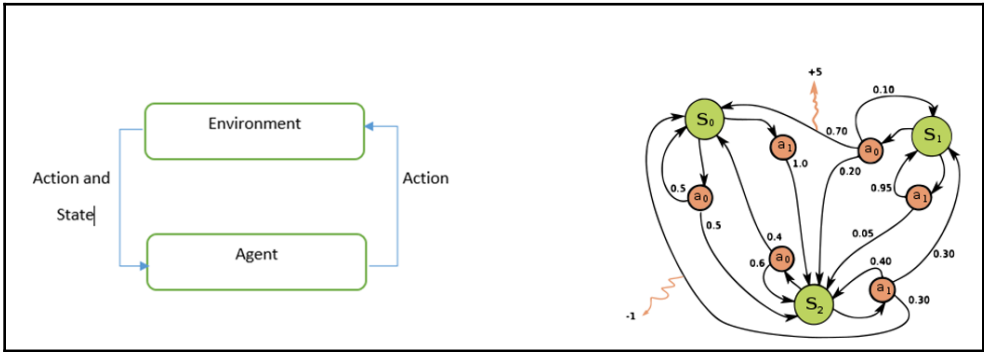
```

2143
+-----+
|R: | : :G|
| : : : : |
| : : : : |
| | : | : |
|Y| : |B: |
+-----+
      (South)
2144
+-----+
|R: | : :G|
| : : : : |
| : : : : |
| | : | : |
|Y| : |B: |
+-----+
      (West)
2145
+-----+
|R: | : :G|
| : : : : |
| : : : : |
| | : | : |
|Y| : |B: |
+-----+
      (West)
2146
+-----+
|R: | : :G|
| : : : : |
| : : : : |
| | : | : |
|Y| : |B: |
+-----+
      (Pickup)

```


Defined by: $(\mathcal{S}, \mathcal{A}, \mathcal{R}, \mathbb{P}, \gamma)$

- \mathcal{S} : set of possible states
- \mathcal{A} : set of possible actions
- \mathcal{R} : distribution of reward given (state, action) pair
- \mathbb{P} : transition probability i.e. distribution over next state given (state, action) pair
- γ : discount factor




Loss, as function of model weights at given time (θ_t):

$$L_t(\theta_t) = \mathbb{E}_{s, a \sim \rho(\cdot)} [(y_t - Q(s, a; \theta_t))^2]$$

Target
Prediction

where $y_t = \mathbb{E}_{s' \sim \mathcal{E}} \left[r + \gamma \max_{a'} Q(s', a'; \theta_{t-1}) \mid s, a \right]$


Target value at time (t)
Current Reward
Max reward at next state
Model weights at (t-1)



```

1 INPUT_SHAPE = (84, 84)
2 WINDOW_LENGTH = 4

```



Loss, as function of model weights at given time (θ_t)

$$L_t(\theta_t) = \mathbb{E}_{s, a \sim \rho(\cdot)} [(y_t - Q(s, a; \theta_t))^2]$$

Target Prediction

where $y_t = \mathbb{E}_{s' \sim \mathcal{E}} [r + \gamma \max_{a'} Q(s', a'; \theta_{t-1}) | s, a]$



Target value at time (t) Current Reward Max reward at next state Model weights at (t-1)

$$L_i(\theta_i) = \mathbb{E}_{s, a \sim \rho(\cdot)} [(y_i - Q(s, a; \theta_i))^2]$$


Target network computes Q-value for state-action pair


$$y_i^{DDQN} = r + \gamma Q(s', \arg \max_{a'} Q(s', a'; \theta_i); \theta^-)$$


Target Value DDQN network selects action for next state



Chapter 08: Autoencoders







Layer (type)	Output Shape	Param #
input_1 (InputLayer)	(None, 784)	0
dense_1 (Dense)	(None, 32)	25120
dense_2 (Dense)	(None, 784)	25872
=====		
Total params: 50,992		
Trainable params: 50,992		
Non-trainable params: 0		

Layer (type)	Output Shape	Param #
input_1 (InputLayer)	(None, 784)	0
dense_1 (Dense)	(None, 32)	25120
=====		
Total params: 25,120		
Trainable params: 25,120		
Non-trainable params: 0		

Layer (type)	Output Shape	Param #
input_2 (InputLayer)	(None, 32)	0
dense_2 (Dense)	(None, 784)	25872
Total params: 25,872		
Trainable params: 25,872		
Non-trainable params: 0		

```

1 autoencoder.fit(x_train, x_train,
2               epochs=50,
3               batch_size=256,
4               shuffle=True,
5               validation_data=(x_test, x_test))

```

Train on 60000 samples, validate on 10000 samples


Epoch 1/50
60000/60000 [=====] - 4s 61us/step - loss: 0.5303 - val_loss: 0.4575

Epoch 2/50
60000/60000 [=====] - 2s 32us/step - loss: 0.4240 - val_loss: 0.4048

Epoch 3/50
60000/60000 [=====] - 2s 38us/step - loss: 0.3940 - val_loss: 0.3865

Epoch 4/50
60000/60000 [=====] - 2s 41us/step - loss: 0.3768 - val_loss: 0.3709

Epoch 5/50
60000/60000 [=====] - 2s 40us/step - loss: 0.3629 - val_loss: 0.3588



```

1 cols = ['Label', 'Latin Name', 'Common Name', 'Train Images', 'Validation Images']
2 labels = pd.read_csv("C:/Users/npurk/Desktop/VAE/monkey_labels.txt", names=cols, skiprows=1)
3 labels

```


Label	Latin Name	Common Name	Train Images	Validation Images	
0	n0	alouatta_palliata	mantled_howler	131	28
1	n1	erythrocebus_patas	patas_monkey	139	28
2	n2	cacajao_calvus	bald_ukari	137	27
3	n3	macaca_fuscata	japanese_macaque	152	30
4	n4	oebuella_pygmaea	pygmy_marmoset	131	28
5	n5	oebus_capucinus	white_headed_capuchin	141	28
6	n6	mico_argentatus	silvery_marmoset	132	28
7	n7	saimiri_sciureus	common_squirrel_monkey	142	28
8	n8	aotus_nigriceps	black_headed_night_monkey	133	27
9	n9	trachypithecus_johnii	nilgiri_langur	132	28






```
1 from sklearn.model_selection import train_test_split
2 x_train, x_test, y_train, y_test = train_test_split(all_monkeys, all_monkeys, test_size=0.2, random_state=42)

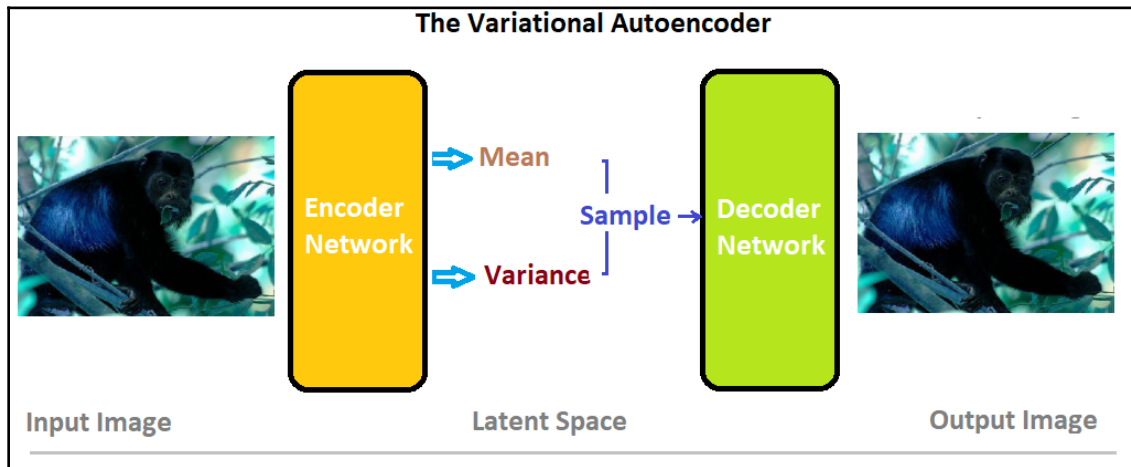
1 plt.imshow(x_train[6].reshape(64, 64, 3))
<matplotlib.image.AxesImage at 0x22b03ff6198>
0
10
20
30
40
50
60
0 10 20 30 40 50 60

1 x_train.shape, x_test.shape
((875, 12288), (219, 12288))
```





Chapter 09: Generative Networks



```
1 encoder_module= Model(input_layer, (z_mean, z_log_var))
2
3 encoder_module.summary()
```

Layer (type)	Output Shape	Param #	Connected to
input_4 (InputLayer)	(None, 784)	0	
Intermediate_layer (Dense)	(None, 256)	200960	input_4[0][0]
z_mean (Dense)	(None, 2)	514	Intermediate_layer[0][0]
z_log_var (Dense)	(None, 2)	514	Intermediate_layer[0][0]

Total params: 201,988
Trainable params: 201,988
Non-trainable params: 0

```
1 def sampling(args):
2     z_mean, z_log_var = args
3     epsilon = K.random_normal(shape=(K.shape(z_mean)[0], latent_dim), mean=0.,
4                               stddev=epsilon_std)
5     return z_mean + K.exp(z_log_var / 2) * epsilon
```

```
1 z = Lambda(sampling, output_shape=(latent_dim,))([z_mean, z_log_var])
```

```

1 # Custom Loss Layer
2 class CustomVariationalLayer(Layer):
3
4     def __init__(self, **kwargs):
5         self.is_placeholder = True
6         super(CustomVariationalLayer, self).__init__(**kwargs)
7
8     def vae_loss(self, x, x_decoded_mean):
9         xent_loss = original_dim * metrics.binary_crossentropy(x, x_decoded_mean)
10        kl_loss = - 0.5 * K.sum(1 + z_log_var - K.square(z_mean) - K.exp(z_log_var), axis=-1)
11        return K.mean(xent_loss + kl_loss)
12
13    def call(self, inputs):
14        x = inputs[0]
15        x_decoded_mean = inputs[1]
16        loss = self.vae_loss(x, x_decoded_mean)
17        self.add_loss(loss, inputs=inputs)
18        return x

```

```

1 y = CustomVariationalLayer()(input_layer, x_decoded_mean)

```

```

1 vae = Model(input_layer, y)
2 vae.compile(optimizer='rmsprop', loss=None)

```

```

1 vae.summary()

```


Layer (type)	Output Shape	Param #	Connected to
input_4 (InputLayer)	(None, 784)	0	
Intermediate_layer (Dense)	(None, 256)	200960	input_4[0][0]
z_mean (Dense)	(None, 2)	514	Intermediate_layer[0][0]
z_log_var (Dense)	(None, 2)	514	Intermediate_layer[0][0]
lambda_1 (Lambda)	(None, 2)	0	z_mean[0][0] z_log_var[0][0]
dense_5 (Dense)	(None, 256)	768	lambda_1[0][0]
dense_6 (Dense)	(None, 784)	201488	dense_5[0][0]
custom_variational_layer_3 (Cus	[(None, 784), (None, 0		input_4[0][0] dense_6[0][0]


=====
 Total params: 404,244
 Trainable params: 404,244
 Non-trainable params: 0

```
1 vae.fit(x_train,  
2       shuffle=True,  
3       epochs=epochs,  
4       batch_size=100,  
5       validation_data=(x_test, None))
```

Train on 60000 samples, validate on 10000 samples

```
Epoch 1/50  
60000/60000 [=====] - 15s 246us/step - loss: 190.8604 - val_loss: 172.9880  
Epoch 2/50  
60000/60000 [=====] - 3s 43us/step - loss: 171.1713 - val_loss: 168.7661  
Epoch 3/50  
60000/60000 [=====] - 3s 43us/step - loss: 167.3232 - val_loss: 165.9071  
Epoch 4/50  
60000/60000 [=====] - 3s 43us/step - loss: 164.7716 - val_loss: 164.0052  
Epoch 5/50  
60000/60000 [=====] - 3s 42us/step - loss: 162.9735 - val_loss: 162.5149
```






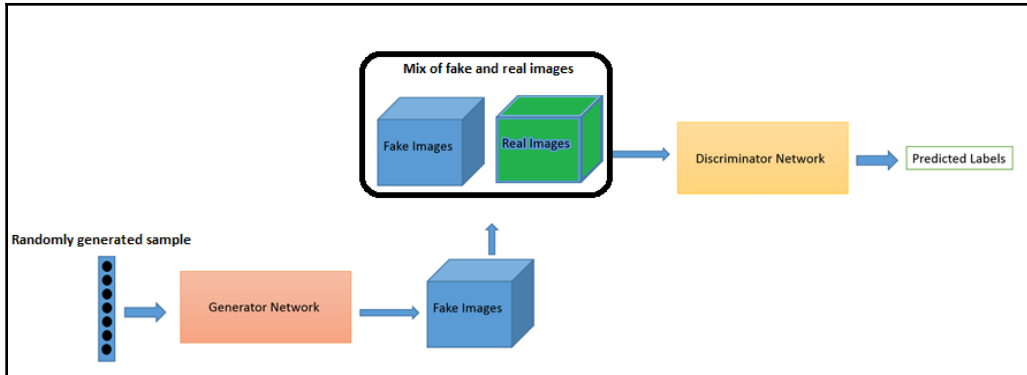
```
1 # build a generator network to sample from the learned distribution
2
3
4 decoder_input = Input(shape=(latent_dim,))
5 _h_decoded = decoder_h(decoder_input)
6 _x_decoded_mean = decoder_mean(_h_decoded)
7 generator = Model(decoder_input, _x_decoded_mean)
```

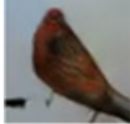

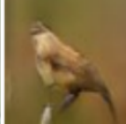

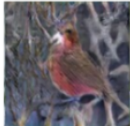



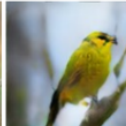

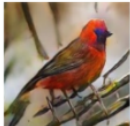
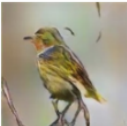

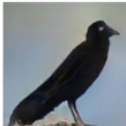

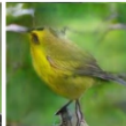

```

1
2
3 # display a 2D manifold of the digits
4 n = 15 # figure with 15x15 digits
5 digit_size = 28
6 figure = np.zeros((digit_size * n, digit_size * n))
7 # linearly spaced coordinates on the unit square were transformed through the inverse CDF (ppf) of the Gaussian
8 # to produce values of the latent variables z, since the prior of the latent space is Gaussian
9 grid_x = norm.ppf(np.linspace(0.05, 0.95, n))
10 grid_y = norm.ppf(np.linspace(0.05, 0.95, n))
11
12 for i, yi in enumerate(grid_x):
13     for j, xi in enumerate(grid_y):
14         z_sample = np.array([[xi, yi]])
15         x_decoded = generator.predict(z_sample)
16         digit = x_decoded[0].reshape(digit_size, digit_size)
17         figure[i * digit_size: (i + 1) * digit_size,
18             j * digit_size: (j + 1) * digit_size] = digit
19
20 plt.figure(figsize=(10, 10))
21 plt.imshow(figure, cmap='binary')
22 plt.show()

```





Text description	This bird is red and brown in color, with a stubby beak	The bird is short and stubby with yellow on its body	A bird with a medium orange bill white body gray wings and webbed feet	This small black bird has a short, slightly curved bill and long legs	A small bird with varying shades of brown with white under the eyes	A small yellow bird with a black crown and a short black pointed beak	This small bird has a white breast, light grey head, and black wings and tail
64x64 GAN-INT-CLS [22]							
128x128 GAWWN [20]							
256x256 StackGAN							

```

1 import numpy as np
2 from tqdm import tqdm
3 from pathlib import Path
4
5 import keras
6 import keras.backend as K
7 from keras.datasets import cifar10
8
9 from keras.models import Sequential, Model
10 from keras.layers import Input, Dense, Activation, LeakyReLU, BatchNormalization, Dropout
11 from keras.layers import Conv2D, Conv2DTranspose, Reshape, Flatten
12 from keras.initializers import RandomNormal
13 from keras.optimizers import Adam
14
15 from sklearn.model_selection import train_test_split
16
17 import matplotlib.pyplot as plt
18 %matplotlib inline

```




```

1 # Load data
2 (x_train, y_train), (x_test, y_test) = cifar10.load_data()
3
4 # Pick the car category
5 x_train = x_train[y_train.flatten() == 1]
6
7 # Check shape
8 x_train.shape, x_test.shape

```


((5000, 32, 32, 3), (1000, 32, 32, 3))



```

1 def preprocess(x):
2     return (x/255)*2-1
3
4 def deprocess(x):
5     return np.uint8((x+1)/2*255)
6
7 X_train_real = preprocess(x_train)
8 X_test_real = preprocess(x_test)

```

<p>At epoch: 1/200, Discriminator Loss: 1.5402 Generator Loss: 4.5400</p>	<p>At epoch: 14/200, Discriminator Loss: 1.1352 Generator Loss: 2.5135</p>	<p>At epoch: 31/200, Discriminator Loss: 1.5590 Generator Loss: 3.1474</p>	<p>At epoch: 55/200, Discriminator Loss: 2.1255 Generator Loss: 3.4757</p>
<p>At epoch: 83/200, Discriminator Loss: 2.2190 Generator Loss: 3.8314</p>	<p>At epoch: 109/200, Discriminator Loss: 2.7688 Generator Loss: 3.9201</p>	<p>At epoch: 176/200, Discriminator Loss: 3.4202 Generator Loss: 4.2350</p>	<p>At epoch: 186/200, Discriminator Loss: 3.2218 Generator Loss: 3.7737</p>

Chapter 10: Contemplating Present and Future Developments

Layer (type)	Output Shape	Param #
input_1 (InputLayer)	(None, 32, 32, 3)	0
block1_conv1 (Conv2D)	(None, 32, 32, 64)	1792
block1_conv2 (Conv2D)	(None, 32, 32, 64)	36928
block1_pool (MaxPooling2D)	(None, 16, 16, 64)	0
block2_conv1 (Conv2D)	(None, 16, 16, 128)	73856
block2_conv2 (Conv2D)	(None, 16, 16, 128)	147584
block2_pool (MaxPooling2D)	(None, 8, 8, 128)	0
block3_conv1 (Conv2D)	(None, 8, 8, 256)	295168
block3_conv2 (Conv2D)	(None, 8, 8, 256)	590080
block3_conv3 (Conv2D)	(None, 8, 8, 256)	590080
block3_conv4 (Conv2D)	(None, 8, 8, 256)	590080
block3_pool (MaxPooling2D)	(None, 4, 4, 256)	0
block4_conv1 (Conv2D)	(None, 4, 4, 512)	1180160
block4_conv2 (Conv2D)	(None, 4, 4, 512)	2359808
block4_conv3 (Conv2D)	(None, 4, 4, 512)	2359808
block4_conv4 (Conv2D)	(None, 4, 4, 512)	2359808
block4_pool (MaxPooling2D)	(None, 2, 2, 512)	0
block5_conv1 (Conv2D)	(None, 2, 2, 512)	2359808
block5_conv2 (Conv2D)	(None, 2, 2, 512)	2359808
block5_conv3 (Conv2D)	(None, 2, 2, 512)	2359808
block5_conv4 (Conv2D)	(None, 2, 2, 512)	2359808
block5_pool (MaxPooling2D)	(None, 1, 1, 512)	0
Total params: 20,024,384		
Trainable params: 20,024,384		

Layer (type)	Output Shape	Param #
input_1 (InputLayer)	(None, 32, 32, 3)	0
block1_conv1 (Conv2D)	(None, 32, 32, 64)	1792
block1_conv2 (Conv2D)	(None, 32, 32, 64)	36928
block1_pool (MaxPooling2D)	(None, 16, 16, 64)	0
block2_conv1 (Conv2D)	(None, 16, 16, 128)	73856
block2_conv2 (Conv2D)	(None, 16, 16, 128)	147584
block2_pool (MaxPooling2D)	(None, 8, 8, 128)	0
block3_conv1 (Conv2D)	(None, 8, 8, 256)	295168
block3_conv2 (Conv2D)	(None, 8, 8, 256)	590080
block3_conv3 (Conv2D)	(None, 8, 8, 256)	590080
block3_conv4 (Conv2D)	(None, 8, 8, 256)	590080
block3_pool (MaxPooling2D)	(None, 4, 4, 256)	0
Total params: 2,325,568		
Trainable params: 2,325,568		
Non-trainable params: 0		

Layer (type)	Output Shape	Param #
input_1 (InputLayer)	(None, 32, 32, 3)	0
block1_conv1 (Conv2D)	(None, 32, 32, 64)	1792
block1_conv2 (Conv2D)	(None, 32, 32, 64)	36928
block1_pool (MaxPooling2D)	(None, 16, 16, 64)	0
block2_conv1 (Conv2D)	(None, 16, 16, 128)	73856
block2_conv2 (Conv2D)	(None, 16, 16, 128)	147584
block2_pool (MaxPooling2D)	(None, 8, 8, 128)	0
block3_conv1 (Conv2D)	(None, 8, 8, 256)	295168
block3_conv2 (Conv2D)	(None, 8, 8, 256)	590080
block3_conv3 (Conv2D)	(None, 8, 8, 256)	590080
block3_conv4 (Conv2D)	(None, 8, 8, 256)	590080
block3_pool (MaxPooling2D)	(None, 4, 4, 256)	0
flatten_1 (Flatten)	(None, 4096)	0
dense_1 (Dense)	(None, 1024)	4195328
dropout_1 (Dropout)	(None, 1024)	0
dense_2 (Dense)	(None, 1024)	1049600
dense_3 (Dense)	(None, 10)	10250
Total params: 7,580,746		
Trainable params: 7,580,746		
Non-trainable params: 0		

```

Train on 50000 samples, validate on 10000 samples
Epoch 1/10
50000/50000 [=====] - 33s 668us/step - loss: 2.3378 - acc: 0.5036 - val_loss: 0.7967 - val_acc: 0.7214
Epoch 2/10
50000/50000 [=====] - 30s 607us/step - loss: 0.7708 - acc: 0.7343 - val_loss: 0.6386 - val_acc: 0.7780
Epoch 3/10
50000/50000 [=====] - 30s 601us/step - loss: 0.6087 - acc: 0.7893 - val_loss: 0.5872 - val_acc: 0.7979
Epoch 4/10
50000/50000 [=====] - 32s 643us/step - loss: 0.4998 - acc: 0.8269 - val_loss: 0.5339 - val_acc: 0.8184

```