



BST 261: Data Science II

Lecture 5

**Convolutional Neural Networks (CNNs):
Data Augmentation
+ Overview so far**

Santiago Romero-Brufau
Harvard T.H. Chan School of Public Health
Spring 2



Administrivia

Fill out the final assignment form!

Office hours for the TAs (in the syllabus): Online after labs on Fridays

Might take attendance in labs (same as in lecture)



“

You can't connect the dots looking forward; you can only connect them looking backwards. So you have to trust that the dots will somehow connect in your future. You have to trust in something - your gut, destiny, life, karma, whatever. This approach has never let me down, and it has made all the difference in my life.

*Steve Jobs (paraphrasing
Kierkegaard)*

<https://youtu.be/UF8uR6Z6KLc>

3

(15 min Stanford Commencement address)

The background of the slide is a complex network diagram. It consists of numerous nodes, represented by small circles, some of which are solid grey and others are hollow with a grey outline. These nodes are interconnected by a web of thin, light-grey lines, creating a dense, interconnected pattern that fills the entire slide area.

Overview so far

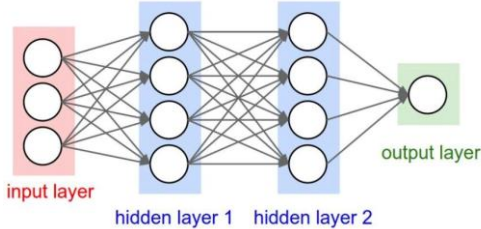
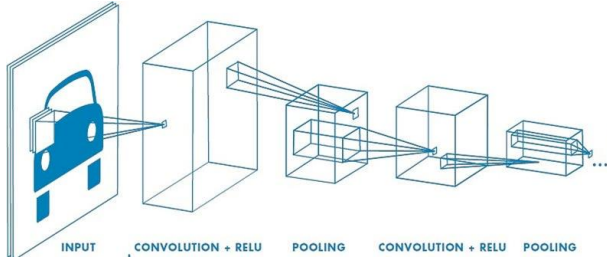
3 Components of a neural network

Network structure: Determines how the inputs are processed


Loss function: How we calculate the “distance” between the prediction and the ground truth (between \hat{Y} and Y). Depends on the output type (classification vs regression)

Optimizer: determines how the parameters are updated. Generally it's a version of Stochastic Gradient Descent.

Neural network structures

Neural Network type	Structure	Best for
MLP (multilayer perceptron)	 <p>input layer hidden layer 1 hidden layer 2 output layer</p>	Structured data (tables)
CNN (convolutional neural network)	 <p>INPUT CONVOLUTION + RELU POOLING CONVOLUTION + RELU POOLING ...</p>	Image data

Forward pass



Key thing to remember: in each layer, the activation function is a **non-linear function**.

Back propagation

Back propagation


We take a loss and calculate the new parameters (W and b)

$$*W_x = W_x - a \left(\frac{\partial \text{Error}}{\partial W_x} \right)$$

Diagram illustrating the backpropagation weight update formula:

- $*W_x$: New weight (indicated by an upward arrow)
- W_x : Old weight (indicated by a downward arrow)
- a : Learning rate (indicated by an upward arrow)
- $\left(\frac{\partial \text{Error}}{\partial W_x} \right)$: Derivative of Error with respect to weight (indicated by a downward arrow)

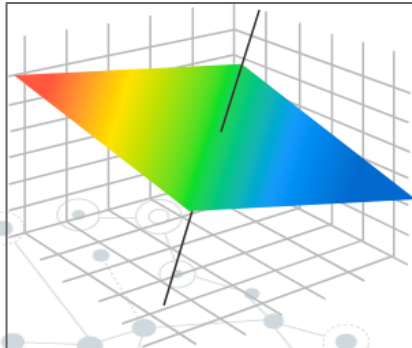

Key thing to remember: in each layer, we move the set of weights in the direction in which the loss function is decreased (downhill).

The background of the slide is a light gray network pattern. It consists of numerous small circles, some of which are solid gray and others are hollow with a gray outline. These circles are interconnected by a web of thin, light gray lines, creating a complex, organic structure that resembles a neural network or a social graph. The pattern is dense and covers the entire area of the slide.


Other quick thoughts

Why non-linearity matters

Without ReLU (only linear functions)



With ReLU (non-linearity!)



$$f(x) = g(x) + g(-x) + g(2x-2) + g(-2x+2)$$

$$\text{where } g(x) = \text{ReLU}(x)$$



Example MLP: multilayer perceptron

Layer (type)	Output Shape	Param #
dense (Dense)	(None, 60)	2700
dense_1 (Dense)	(None, 55)	3355
dense_2 (Dense)	(None, 50)	2800
dense_3 (Dense)	(None, 45)	2295
dense_4 (Dense)	(None, 30)	1380
dense_5 (Dense)	(None, 20)	620
dense_6 (Dense)	(None, 1)	21
Total params: 13,171		
Trainable params: 13,171		
Non-trainable params: 0		




CNNs in Python

Classifying Skin Lesions

- ◎ We'll be using data from the [International Skin Imaging Collaboration: Melanoma Project](#)
- ◎ We'll be classifying images as malignant or benign
- ◎ The overarching goal of the ISIC Melanoma Project is to support efforts to reduce melanoma-related deaths and unnecessary biopsies by improving the accuracy and efficiency of melanoma early detection

Malignant



Malignant



Benign



Benign



Classifying Skin Lesions

- ◎ This archive contains 23k images of classified skin lesions. It contains both malignant and benign examples
 - We'll be using a fraction of this
- ◎ Each example contains the image of the lesion, meta data regarding the lesion (including classification and segmentation) and meta data regarding the patient
- ◎ The data can be viewed in [this link](#) (in the gallery section)
- ◎ It can be downloaded through the site or by using [this repository](#)

Classification Skin Lesions

- ◎ The subsample of the data is available in a [Google Drive](#) folder
- ◎ You can access it with this [code and notebook](#)
- ◎ You can also download the images to your machine if you would like
 - There are zip files available on canvas
- ◎ We'll start with creating a simple CNN


```

1 # Define model
2 model = keras.Sequential([
3     layers.Conv2D(32, (3, 3), activation='relu', input_shape=(150, 150, 3)),
4     layers.MaxPooling2D((2, 2)),
5
6     layers.Conv2D(64, (3, 3), activation='relu'),
7     layers.MaxPooling2D((2, 2)),
8
9     layers.Conv2D(128, (3, 3), activation='relu'),
10    layers.MaxPooling2D((2, 2)),
11
12    layers.Conv2D(128, (3, 3), activation='relu'),
13    layers.MaxPooling2D((2, 2)),
14
15    layers.Flatten(),
16
17    layers.Dense(512, activation='relu'),
18
19    layers.Dense(1, activation='sigmoid')
20 ])

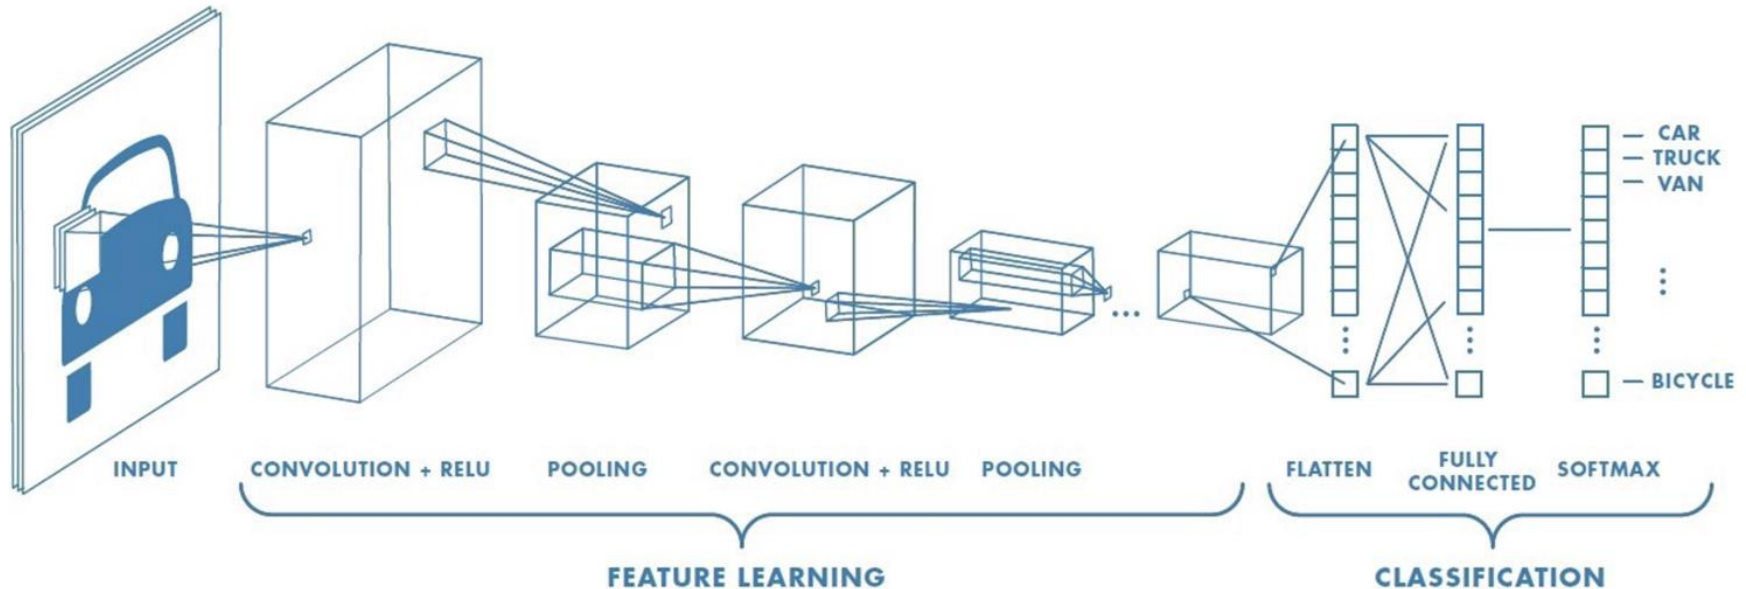
```

Model: "sequential"

Layer (type)	Output Shape	Param #
=====		
conv2d (Conv2D)	(None, 148, 148, 32)	896
max_pooling2d (MaxPooling2D)	(None, 74, 74, 32)	0
conv2d_1 (Conv2D)	(None, 72, 72, 64)	18496
max_pooling2d_1 (MaxPooling2D)	(None, 36, 36, 64)	0
conv2d_2 (Conv2D)	(None, 34, 34, 128)	73856
max_pooling2d_2 (MaxPooling2D)	(None, 17, 17, 128)	0
conv2d_3 (Conv2D)	(None, 15, 15, 128)	147584
max_pooling2d_3 (MaxPooling2D)	(None, 7, 7, 128)	0
flatten (Flatten)	(None, 6272)	0
dense (Dense)	(None, 512)	3211776
dense_1 (Dense)	(None, 1)	513
=====		

Total params: 3,453,121
 Trainable params: 3,453,121
 Non-trainable params: 0

CNN Schematic



Number of parameters

width m, height n, previous layer's filters d and account for all such filters **k in the current layer**. Don't forget the bias term for each of the filter. Number of parameters in a CONV layer would be : **$((m * n * d) + 1) * k$** , added 1 because of the bias term for each filter.

$((\text{shape of width of the filter} * \text{shape of height of the filter} * \text{number of filters in the previous layer} + 1) * \text{number of filters})$

$$((3 * 3 * 32 + 1) * 3)$$

<https://towardsdatascience.com/understanding-and-calculating-the-number-of-parameters-in-convolution-neural-networks-cnns-fc88790d530d>

```
1 # Define model
2 model = keras.Sequential([
3     layers.Conv2D(32, (3, 3), activation='relu', input_shape=(150, 150, 3)),
```

width m, height n, previous layer's filters d and account for all such filters **k in the current layer**. Don't forget the bias term for each of the filter. Number of parameters in a CONV layer would be : $((m * n * d) + 1) * k$, added 1 because of the bias term for each filter.

$$((m * n * d) + 1) * k$$
$$((3 * 3 * 3) + 1) * 32 = 896$$

Model: "sequential"

Layer (type)	Output Shape	Param #
=====		
conv2d (Conv2D)	(None, 148, 148, 32)	896

```

1 # Define model
2 model = keras.Sequential([
3     layers.Conv2D(32, (3, 3), activation='relu', input_shape=(150, 150, 3)),
4     layers.MaxPooling2D((2, 2)),
5
6     layers.Conv2D(64, (3, 3), activation='relu'),
7     layers.MaxPooling2D((2, 2)),
8
9     layers.Conv2D(128, (3, 3), activation='relu'),
10    layers.MaxPooling2D((2, 2)),
11
12    layers.Conv2D(128, (3, 3), activation='relu'),
13    layers.MaxPooling2D((2, 2)),
14
15    layers.Flatten(),
16
17    layers.Dense(512, activation='relu'),
18
19    layers.Dense(1, activation='sigmoid')
20 ])

```

Convolution and pooling layers - notice how the output size decreases with each layer. Remember what applying a filter, padding, and/or strides does to an input.

Model: "sequential"

Layer (type)	Output Shape	Param #
conv2d (Conv2D)	(None, 148, 148, 32)	896
max_pooling2d (MaxPooling2D)	(None, 74, 74, 32)	0
conv2d_1 (Conv2D)	(None, 72, 72, 64)	18496
max_pooling2d_1 (MaxPooling2D)	(None, 36, 36, 64)	0
conv2d_2 (Conv2D)	(None, 34, 34, 128)	73856
max_pooling2d_2 (MaxPooling2D)	(None, 17, 17, 128)	0
conv2d_3 (Conv2D)	(None, 15, 15, 128)	147584
max_pooling2d_3 (MaxPooling2D)	(None, 7, 7, 128)	0
flatten (Flatten)	(None, 6272)	0
dense (Dense)	(None, 512)	3211776
dense_1 (Dense)	(None, 1)	513
=====		
Total params: 3,453,121		
Trainable params: 3,453,121		
Non-trainable params: 0		

```

1 # Define model
2 model = keras.Sequential([
3     layers.Conv2D(32, (3, 3), activation='relu', input_shape=(150, 150, 3)),
4     layers.MaxPooling2D((2, 2)),
5
6     layers.Conv2D(64, (3, 3), activation='relu'),
7     layers.MaxPooling2D((2, 2)),
8
9     layers.Conv2D(128, (3, 3), activation='relu'),
10    layers.MaxPooling2D((2, 2)),
11
12    layers.Conv2D(128, (3, 3), activation='relu'),
13    layers.MaxPooling2D((2, 2)),
14
15    layers.Flatten(),
16
17    layers.Dense(512, activation='relu'),
18
19    layers.Dense(1, activation='sigmoid')
20 ])

```

We finish the network with 1 hidden dense layer and 1 output layer. Note that most of the parameters in the model come from the hidden dense layer and not the convolution or pooling layers.

Model: "sequential"

Layer (type)	Output Shape	Param #
=====		
conv2d (Conv2D)	(None, 148, 148, 32)	896
max_pooling2d (MaxPooling2D)	(None, 74, 74, 32)	0
conv2d_1 (Conv2D)	(None, 72, 72, 64)	18496
max_pooling2d_1 (MaxPooling2D)	(None, 36, 36, 64)	0
conv2d_2 (Conv2D)	(None, 34, 34, 128)	73856
max_pooling2d_2 (MaxPooling2D)	(None, 17, 17, 128)	0
conv2d_3 (Conv2D)	(None, 15, 15, 128)	147584
max_pooling2d_3 (MaxPooling2D)	(None, 7, 7, 128)	0
flatten (Flatten)	(None, 6272)	0
dense (Dense)	(None, 512)	3211776
dense_1 (Dense)	(None, 1)	513

Total params: 3,453,121
Trainable params: 3,453,121
Non-trainable params: 0

```

1 # Define model
2 model = keras.Sequential([
3     layers.Conv2D(32, (3, 3), activation='relu', input_shape=(150, 150, 3)),
4     layers.MaxPooling2D((2, 2)),
5
6     layers.Conv2D(64, (3, 3), activation='relu'),
7     layers.MaxPooling2D((2, 2)),
8
9     layers.Conv2D(128, (3, 3), activation='relu'),
10    layers.MaxPooling2D((2, 2)),
11
12    layers.Conv2D(128, (3, 3), activation='relu'),
13    layers.MaxPooling2D((2, 2)),
14
15    layers.Flatten(),
16
17    layers.Dense(512, activation='relu'),
18
19    layers.Dense(1, activation='sigmoid')
20 ])

```

Model: "sequential"

Layer (type)	Output Shape	Param #
=====		
conv2d (Conv2D)	(None, 148, 148, 32)	896
max_pooling2d (MaxPooling2D)	(None, 74, 74, 32)	0
conv2d_1 (Conv2D)	(None, 72, 72, 64)	18496
max_pooling2d_1 (MaxPooling2	(None, 36, 36, 64)	0
conv2d_2 (Conv2D)	(None, 34, 34, 128)	73856
max_pooling2d_2 (MaxPooling2	(None, 17, 17, 128)	0
conv2d_3 (Conv2D)	(None, 15, 15, 128)	147584
max_pooling2d_3 (MaxPooling2	(None, 7, 7, 128)	0
flatten (Flatten)	(None, 6272)	0
dense (Dense)	(None, 512)	3211776
dense_1 (Dense)	(None, 1)	513

Total params: 3,453,121
 Trainable params: 3,453,121
 Non-trainable params: 0

Total # of parameters that need to be learned by the network

```
1 from keras.preprocessing.image import ImageDataGenerator
2
3 # All images will be rescaled by 1./255
4 train_datagen = ImageDataGenerator(rescale=1./255)
5 test_datagen = ImageDataGenerator(rescale=1./255)
6
7 train_generator = train_datagen.flow_from_directory(
8     # This is the target directory
9     train_dir,
10    # All images will be resized to 150x150
11    target_size = (150, 150),
12    batch_size = 20,
13    # Since we use binary_crossentropy loss, we need binary labels
14    class_mode = 'binary')
15
16 validation_generator = test_datagen.flow_from_directory(
17     validation_dir,
18     target_size = (150, 150),
19     batch_size = 20,
20     class_mode = 'binary')
```

We first scale the data to get values between 0 and 1.

Then we transform the images to be 150x150 pixels in size (this is arbitrary), declare a batch size of 20 (this is also arbitrary), and declare the class mode (i.e. the type of classification we want to do)




```
1 for data_batch, labels_batch in train_generator:
2     print('data batch shape:', data_batch.shape)
3     print('labels batch shape:', labels_batch.shape)
4     break
```

```
data batch shape: (20, 150, 150, 3)
labels batch shape: (20,)
```

We can see the shape of the images: they are in batches of 20, with each image being represented by 3, 150x150 tensors: one for R, one for G and one for B color “channels”.

```
1 history = model.fit(
2     train_generator,
3     steps_per_epoch = 81, # ceil(1609/20)
4     epochs = 30,
5     validation_data = validation_generator,
6     validation_steps = 22) # ceil(426/20)
```

The number of training examples divided by the batch size, i.e. how many batches we need to go through until the model sees all training data. For both the training and validation sets.



Data Augmentation

- ◎ As we have seen, overfitting is caused by having too few training examples to learn from
- ◎ Data augmentation generates more training data from existing training examples by **augmenting** the samples via a number of random transformations
- ◎ These transformations should yield believable images

Data Augmentation

- ◎ Types of augmentation:
 - Rotation
 - Horizontal/vertical flip
 - Random crops/scales
 - ◎ Zoom
 - ◎ Width or height shifts
 - Shearing
 - Brightness, contrast, saturation
 - Lens distortions

Types of data augmentation

1. Rotations



Types of data augmentation

2. Horizontal/Vertical Flips



Types of data augmentation

3. Random crops/scales



Types of data augmentation

4. Shearing



Types of data augmentation

5. Brightness, contrast, saturation



Types of data augmentation

6. Lens distortions



Types of data augmentation

7. Combinations of the above



Data Augmentation

- ◎ If you train a network using data augmentation, it will never see the same input twice, but the inputs will still be heavily correlated
 - You're remixing known information, not producing new information
- ◎ May not completely escape overfitting due to this correlation
- ◎ Adding dropout can also help

Data Augmentation in Keras

```
1 from keras.preprocessing.image import ImageDataGenerator
2 datagen = ImageDataGenerator(
3     rotation_range = 40,
4     width_shift_range = 0.2,
5     height_shift_range = 0.2,
6     shear_range = 0.2,
7     zoom_range = 0.2,
8     horizontal_flip = True,
9     fill_mode = 'nearest')
```

You can create your own data generator with any specifications you'd like. The values chosen here are arbitrary.

You can check out the [Keras documentation](#) to see all of the available options and values each type of augmentation type can take.





Note that only your training data should be augmented - not the test or validation sets. The point of augmentation is to “increase” your training set size.

```
1 from keras.preprocessing.image import ImageDataGenerator
2
3 # All images will be rescaled by 1./255
4 train_datagen = ImageDataGenerator(rescale=1./255)
5 test_datagen = ImageDataGenerator(rescale=1./255)
6
7 train_generator = train_datagen.flow_from_directory(
8     # This is the target directory
9     train_dir,
10    # All images will be resized to 150x150
11    target_size = (150, 150),
12    batch_size = 20,
13    # Since we use binary_crossentropy loss, we need binary labels
14    class_mode = 'binary')
15
16 validation_generator = test_datagen.flow_from_directory(
17     validation_dir,
18     target_size = (150, 150),
19     batch_size = 20,
20     class_mode = 'binary')
```

We first scale the data to get values between 0 and 1.

Then we transform the images to be 150x150 pixels in size (this is arbitrary), declare a batch size of 20 (this is also arbitrary), and declare the class mode (i.e. the type of classification we want to do)

Data Augmentation in Keras

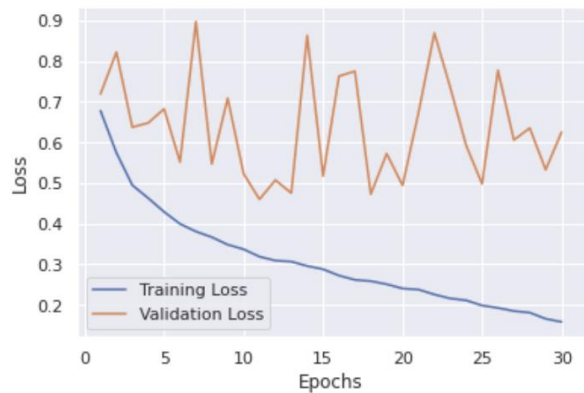


Back to the notebook

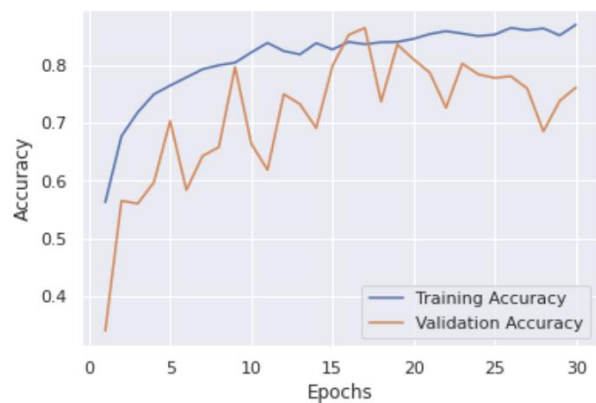

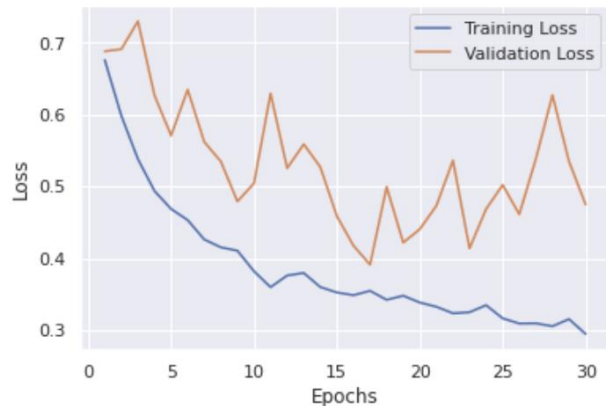
[Skin lesions with data augmentation](#)

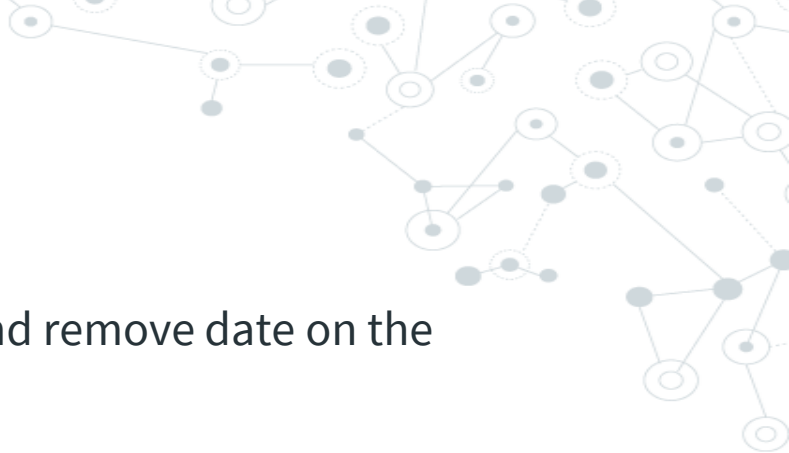
Data Augmentation in Keras

Without augmentation



With augmentation





Change due date for HW1 on Canvas (to Sat), and remove date on the homework itself