

AI/ML BOOTCAMP

William Phelps

Christopher Newport University/Jefferson Lab

Course Topics*

I Semester DS/AI Course in 4 lectures
45 hours in 8 hours

- **Lectures 1 & 2**

- **Lecture 1**

- Introduction to workflow and tools (git/anaconda/pycharm)
- Introduction to Data Science
- Python Review with jupyter notebooks

- **Lecture 2**

- Data Visualization
- Principles of data visualization from Edward Tufte
- Pandas introduction

- **Lectures 3 & 4**

- **Lecture 3**

- Machine Learning Introduction
- Regression/Curvefit
- Overfitting/Underfitting
- kNN classifier
- MNIST/Handwritten digit classifier

- **Lecture 4**

- Overfitting/Underfitting
- Early Stopping
- Hyperparameter optimization
- Convolutional Neural Networks
- CIFAR 10 datasets

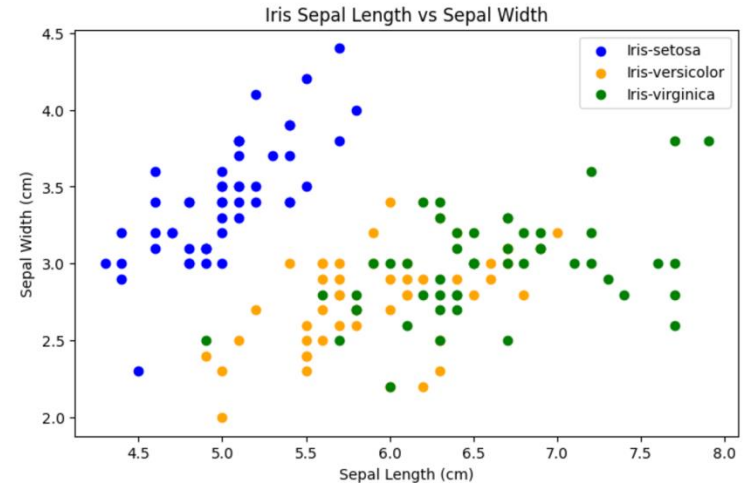

*Tentative

RECAP/MNIST

Recap – Regression and kNN

```
# Let's fit with curve fit and extract parameters
from scipy.optimize import curve_fit
def func(x, a, b):
    return a*np.sin(x) + b
popt, pcov = curve_fit(func, x, y)

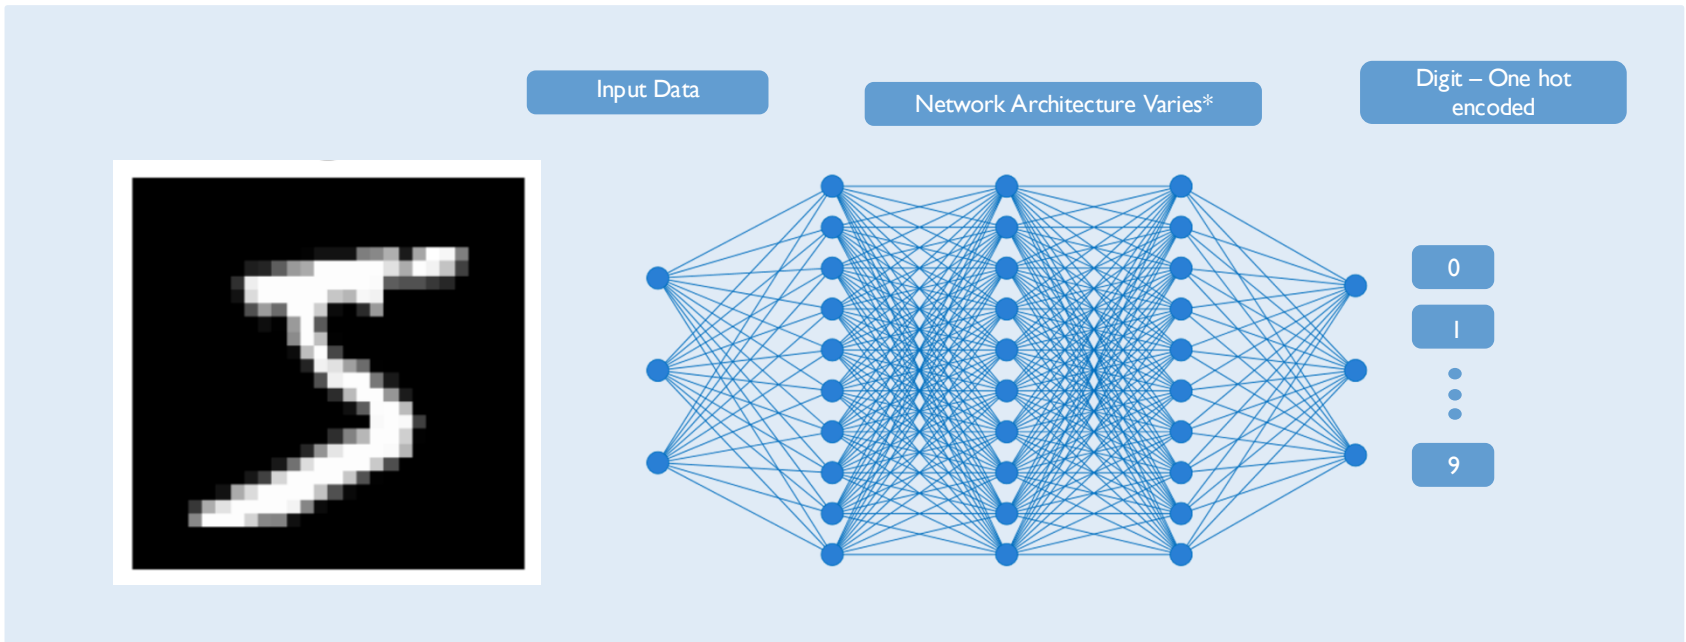
# Plot function on top of dataset
plt.plot(x, y, 'o')
plt.plot(x, func(x, *popt), 'r-');
```



```
from sklearn.neighbors import KNeighborsClassifier
knn = KNeighborsClassifier(n_neighbors=3)
knn.fit(X_train, y_train)
```


MLP Example – MNIST

- Read in 28x28 images of handwritten digits
- Model Architecture: MLP, relu activation function




MNIST MLP

```
# Build a MLP
model = keras.Sequential()
model.add(layers.Flatten(input_shape=(28, 28)))
model.add(layers.Dense(512, activation='relu'))
model.add(layers.Dense(512, activation='relu'))
model.add(layers.Dense(512, activation='relu'))
model.add(layers.Dense(10, activation='softmax'))
model.summary();

# Compile model
model.compile(optimizer='adam',
              loss='sparse_categorical_crossentropy',
              metrics=['accuracy'])

# Train the model
history = model.fit(train_images,
                   train_labels,
                   epochs=150,
                   batch_size=128,
                   validation_split=0.2)
```

Is this the best we
can do with a MLP?



How many epochs?

How well did we do?


MNIST MLP

```
# Build a MLP
model = keras.Sequential()
model.add(layers.Flatten(input_shape=(28, 28)))
model.add(layers.Dense(512, activation='relu'))
model.add(layers.Dense(512, activation='relu'))
model.add(layers.Dense(512, activation='relu'))
model.add(layers.Dense(10, activation='softmax'))
model.summary();

# Compile model
model.compile(optimizer='adam',
              loss='sparse_categorical_crossentropy',
              metrics=['accuracy'])

# Train the model
history = model.fit(train_images,
                   train_labels,
                   epochs=150,
                   batch_size=128,
                   validation_split=0.2)
```

Is this the best we
can do with a MLP?



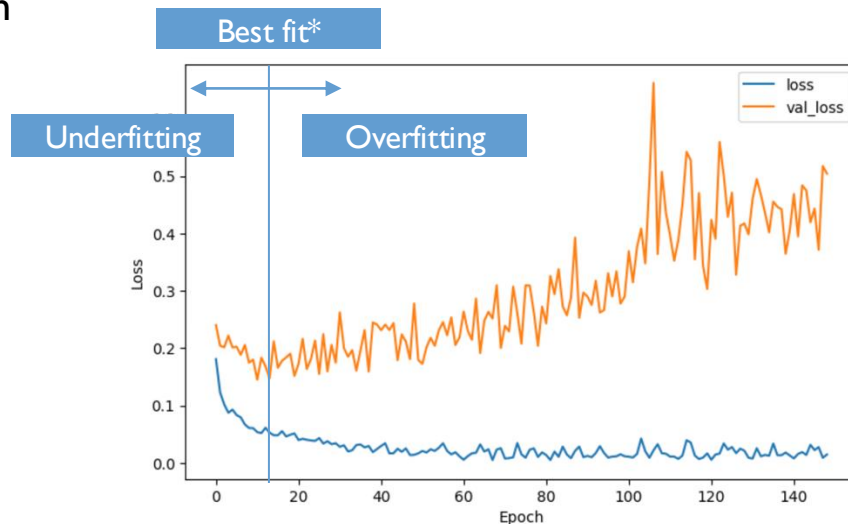
How many epochs?

How well did we do?

Early Stopping with Callbacks

- Callbacks monitor training and trigger actions during `model.fit()`.
- Early stopping stops training when a chosen metric stops improving.
- Here, `monitor='val_loss'` tracks validation loss.
- `patience=3` allows 3 epochs without improvement before stopping.
- This helps reduce overfitting and saves training time.

```
# Example of early stopping with callbacks
from tensorflow.keras.callbacks import EarlyStopping
early_stopping = EarlyStopping(monitor='val_loss', patience=3)
history = model.fit(train_images,
                    train_labels,
                    epochs=50,
                    validation_split=0.2,
                    callbacks=[early_stopping])
```



* Somewhat debatable in this example


MNIST MLP

```
# Build a MLP
model = keras.Sequential()
model.add(layers.Flatten(input_shape=(28, 28)))
model.add(layers.Dense(512, activation='relu'))
model.add(layers.Dense(512, activation='relu'))
model.add(layers.Dense(512, activation='relu'))
model.add(layers.Dense(10, activation='softmax'))
model.summary();

# Compile model
model.compile(optimizer='adam',
              loss='sparse_categorical_crossentropy',
              metrics=['accuracy'])

# Train the model
history = model.fit(train_images,
                    train_labels,
                    epochs=150,
                    batch_size=128,
                    validation_split=0.2)
```

Is this the best we
can do with a MLP?



How many epochs?


How well did we do?

Accuracy and Confusion Matrix

- `model.evaluate(test_images, test_labels)` measures performance on unseen test data.
- It returns the test loss and any metrics specified during `model.compile()`, such as accuracy.
- The test data should not be used in training!
- Next we will look at the confusion matrix


```
# Calculate accuracy on the test set
test_loss, test_acc = model.evaluate(test_images, test_labels)

313/313 [=====] - 1s 4ms/step - loss: 0.3390 - accuracy: 0.9720
```



Confusion Matrix Cont'd

- `model.predict(test_images)` generates class probabilities for each test image.
- `argmax(axis=1)` converts probabilities into predicted class labels.
- The confusion matrix compares true vs. predicted labels
- strong performance appears as large values along the diagonal.



```
# Plot confusion matrix
from sklearn.metrics import confusion_matrix
import seaborn as sns
predictions = model.predict(test_images)
predicted_labels = predictions.argmax(axis=1)
cm = confusion_matrix(test_labels, predicted_labels)
plt.figure(figsize=(10, 8))
sns.heatmap(cm,
            annot=True,
            fmt='d',
            cmap='Blues',
            xticklabels=range(10),
            yticklabels=range(10))
plt.title('Confusion Matrix')
plt.xlabel('Predicted Label')
plt.ylabel('True Label')
plt.show();
```


MNIST MLP

```
# Build a MLP
model = keras.Sequential()
model.add(layers.Flatten(input_shape=(28, 28)))
model.add(layers.Dense(512, activation='relu'))
model.add(layers.Dense(512, activation='relu'))
model.add(layers.Dense(512, activation='relu'))
model.add(layers.Dense(10, activation='softmax'))
model.summary();

# Compile model
model.compile(optimizer='adam',
              loss='sparse_categorical_crossentropy',
              metrics=['accuracy'])

# Train the model
history = model.fit(train_images,
                   train_labels,
                   epochs=150,
                   batch_size=128,
                   validation_split=0.2)
```

Is this the best we
can do with a MLP?



How many epochs?

How well did we do?

What is a hyperparameter?

- They are not learned directly from the data.
- They control the model architecture or training process.
- **Examples:** number of layers, number of neurons, learning rate, batch size, dropout rate, optimizer, and number of epochs.
- **Hyperparameter optimization** searches for settings that improve validation performance.
- In contrast, **model parameters are learned during training**, such as weights and biases.

Example Hyperparameters

```
# Build a MLP
model = keras.Sequential()
model.add(layers.Flatten(input_shape=(28, 28)))
model.add(layers.Dense(512, activation='relu'))
model.add(layers.Dense(512, activation='relu'))
model.add(layers.Dense(512, activation='relu'))
model.add(layers.Dense(10, activation='softmax'))
model.summary();



# Compile model
model.compile(optimizer='adam',
              loss='sparse_categorical_crossentropy',
              metrics=['accuracy'])

# Train the model
history = model.fit(train_images,
                   train_labels,
                   epochs=150,
                   batch_size=128,
                   validation_split=0.2)
```

Hyperparameter tuning Example

IMAGES AND CONVOLUTIONS


What is an image?



Digital Camera




<https://en.wikipedia.org/wiki/Pointillism>



The Eye

What is a digital image?

- A grid (matrix) of intensity values






=

| | | | | | | | | | | | | | | |
|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|
| 255 | 255 | 255 | 255 | 255 | 255 | 255 | 255 | 255 | 255 | 255 | 255 | 255 | 255 | 255 |
| 255 | 255 | 255 | 255 | 255 | 255 | 255 | 255 | 255 | 255 | 255 | 255 | 255 | 255 | 255 |
| 255 | 255 | 255 | 20 | 0 | 255 | 255 | 255 | 255 | 255 | 255 | 255 | 255 | 255 | 255 |
| 255 | 255 | 255 | 75 | 75 | 75 | 255 | 255 | 255 | 255 | 255 | 255 | 255 | 255 | 255 |
| 255 | 255 | 75 | 95 | 95 | 75 | 255 | 255 | 255 | 255 | 255 | 255 | 255 | 255 | 255 |
| 255 | 255 | 96 | 127 | 145 | 175 | 255 | 255 | 255 | 255 | 255 | 255 | 255 | 255 | 255 |
| 255 | 255 | 127 | 145 | 175 | 175 | 175 | 255 | 255 | 255 | 255 | 255 | 255 | 255 | 255 |
| 255 | 255 | 127 | 145 | 200 | 200 | 175 | 175 | 95 | 255 | 255 | 255 | 255 | 255 | 255 |
| 255 | 255 | 127 | 145 | 200 | 200 | 175 | 175 | 95 | 47 | 255 | 255 | 255 | 255 | 255 |
| 255 | 255 | 127 | 145 | 145 | 175 | 127 | 127 | 95 | 47 | 255 | 255 | 255 | 255 | 255 |
| 255 | 255 | 74 | 127 | 127 | 127 | 95 | 95 | 95 | 47 | 255 | 255 | 255 | 255 | 255 |
| 255 | 255 | 255 | 74 | 74 | 74 | 74 | 74 | 74 | 255 | 255 | 255 | 255 | 255 | 255 |
| 255 | 255 | 255 | 255 | 255 | 255 | 255 | 255 | 255 | 255 | 255 | 255 | 255 | 255 | 255 |
| 255 | 255 | 255 | 255 | 255 | 255 | 255 | 255 | 255 | 255 | 255 | 255 | 255 | 255 | 255 |

(common to use one byte per value: 0 = black, 255 = white)

Color Images



https://en.wikipedia.org/wiki/Bayer_filter




https://en.wikipedia.org/wiki/RGB_color_model#/media/File:Rgb-compose-Alim_Khan.jpg

What is an image?

- We can think of a (grayscale) image as a **function**, f , from \mathbb{R}^2 to \mathbb{R} :
 - $f(x,y)$ gives the **intensity** at position (x,y)



snoop



3D view


- A **digital** image is a discrete (**sampled, quantized**) version of this function

Image transformations

- As with any function, we can apply operators to an image



$$g(x,y) = f(x,y) + 20$$



$$g(x,y) = f(-x,y)$$

- We'll talk about a special kind of operator, *convolution* (linear filtering)

Image filtering

- Modify the pixels in an image based on some function of a local neighborhood of each pixel

| | | |
|----|---|---|
| 10 | 5 | 3 |
| 4 | 5 | 1 |
| 1 | 1 | 7 |

Local image data

Some function




| | | |
|--|---|--|
| | | |
| | 7 | |
| | | |

Modified image data

Linear filtering

- One simple version: linear filtering
(cross-correlation, convolution)
 - Replace each pixel by a linear combination (a weighted sum) of its neighbors
- The prescription for the linear combination is called the “kernel” (or “mask”, “filter”)



Cross-correlation

Let F be the image, H be the kernel (of size $2k+1 \times 2k+1$), and G be the output image

$$G[i, j] = \sum_{u=-k}^k \sum_{v=-k}^k H[u, v] F[i + u, j + v]$$

This is called a **cross-correlation** operation:

$$G = H \otimes F$$

- Can think of as a “dot product” between local neighborhood and kernel for each pixel

Convolution

- Same as cross-correlation, except that the kernel is “flipped” (horizontally and vertically)


$$G[i, j] = \sum_{u=-k}^k \sum_{v=-k}^k H[u, v] F[i - u, j - v]$$

This is called a **convolution** operation:

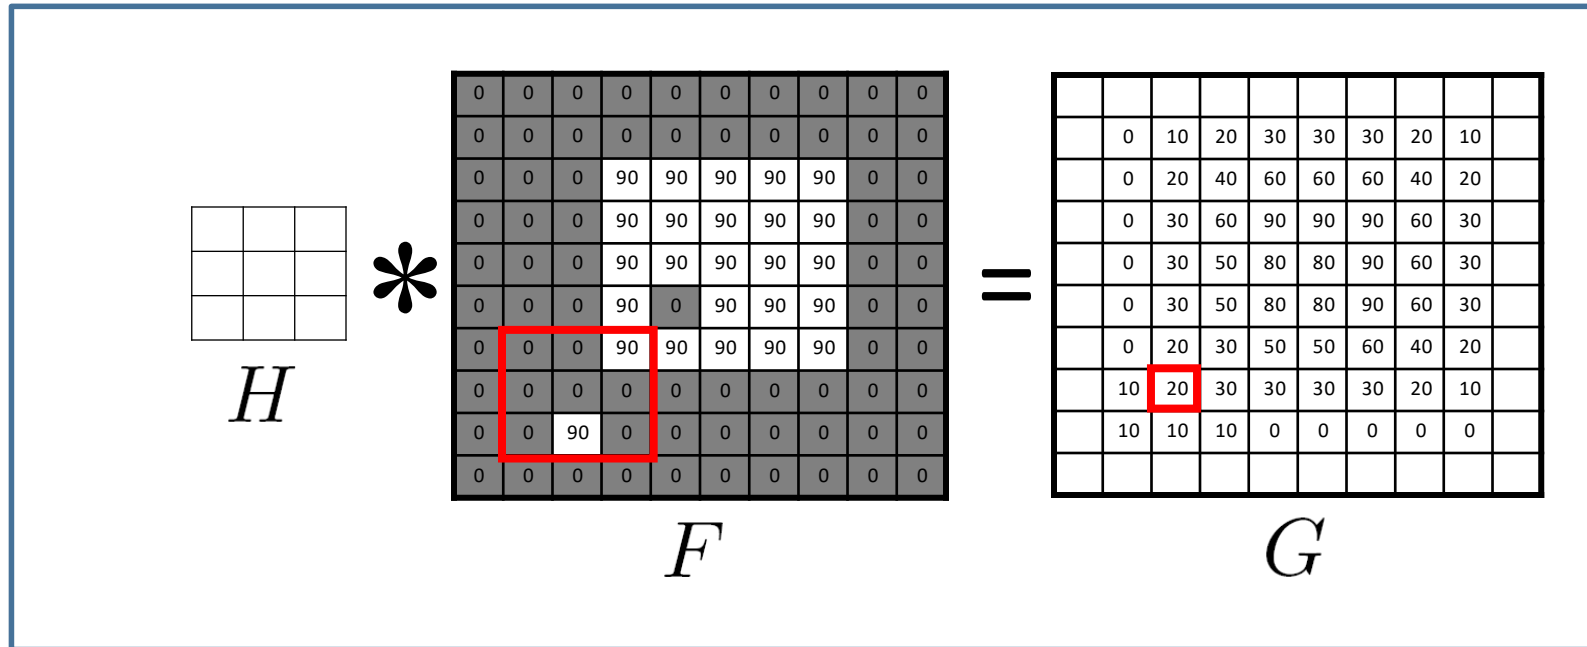
$$G = H * F$$

- Convolution is **commutative** and **associative**
 - Can combine convolution operations


Convolution



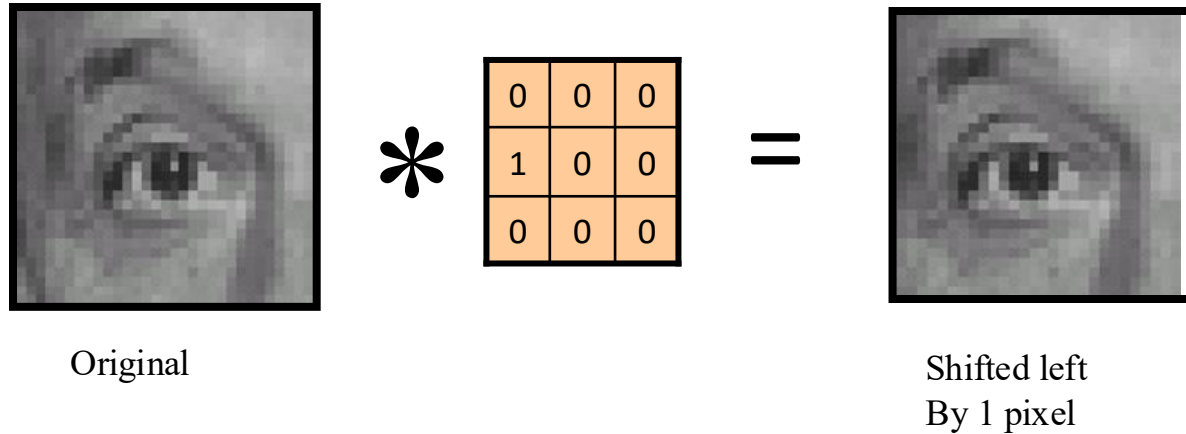
Mean filtering




Linear filters: examples



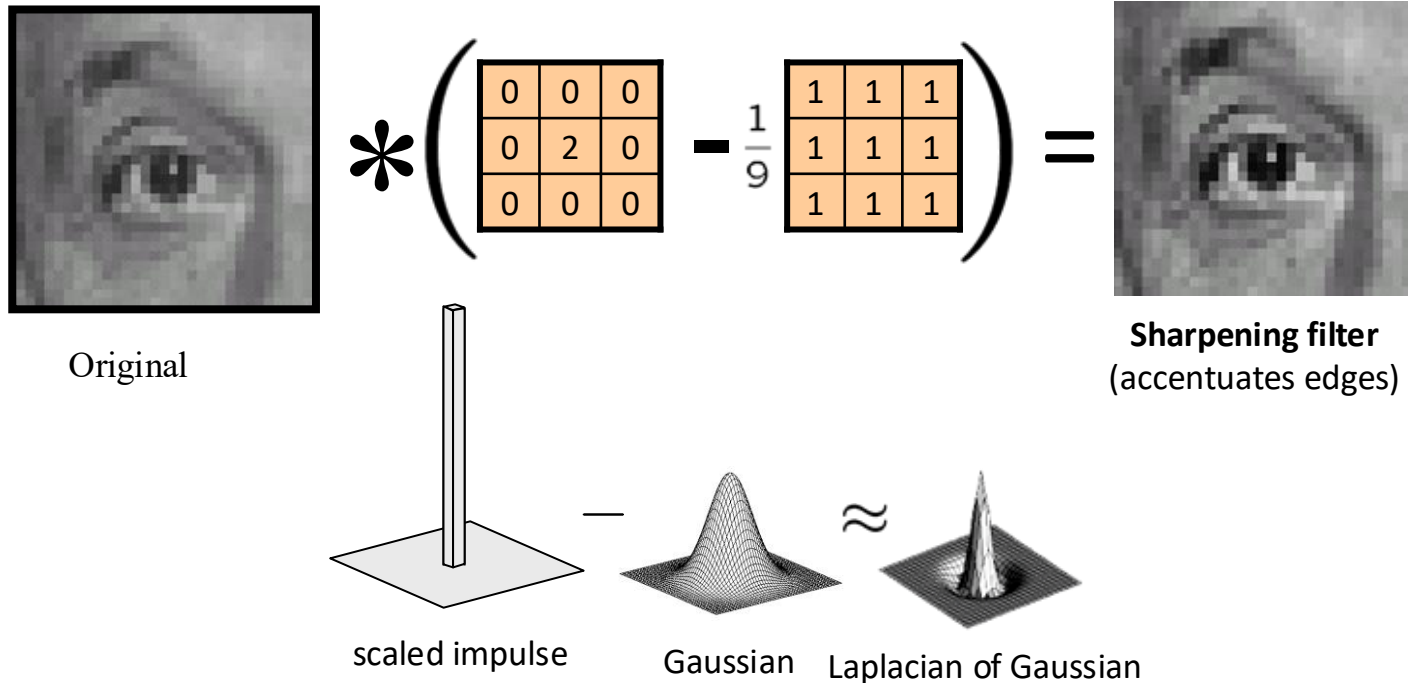
Linear filters: examples




Linear filters: examples




Linear filters: examples



Sharpening





before



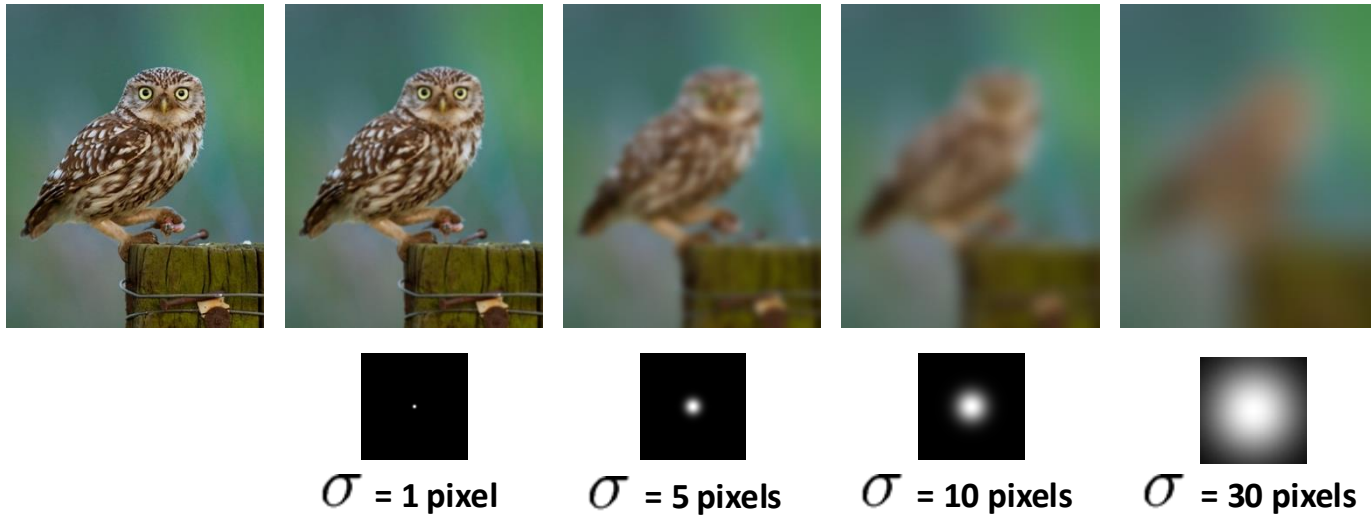
after

Gaussian Kernel

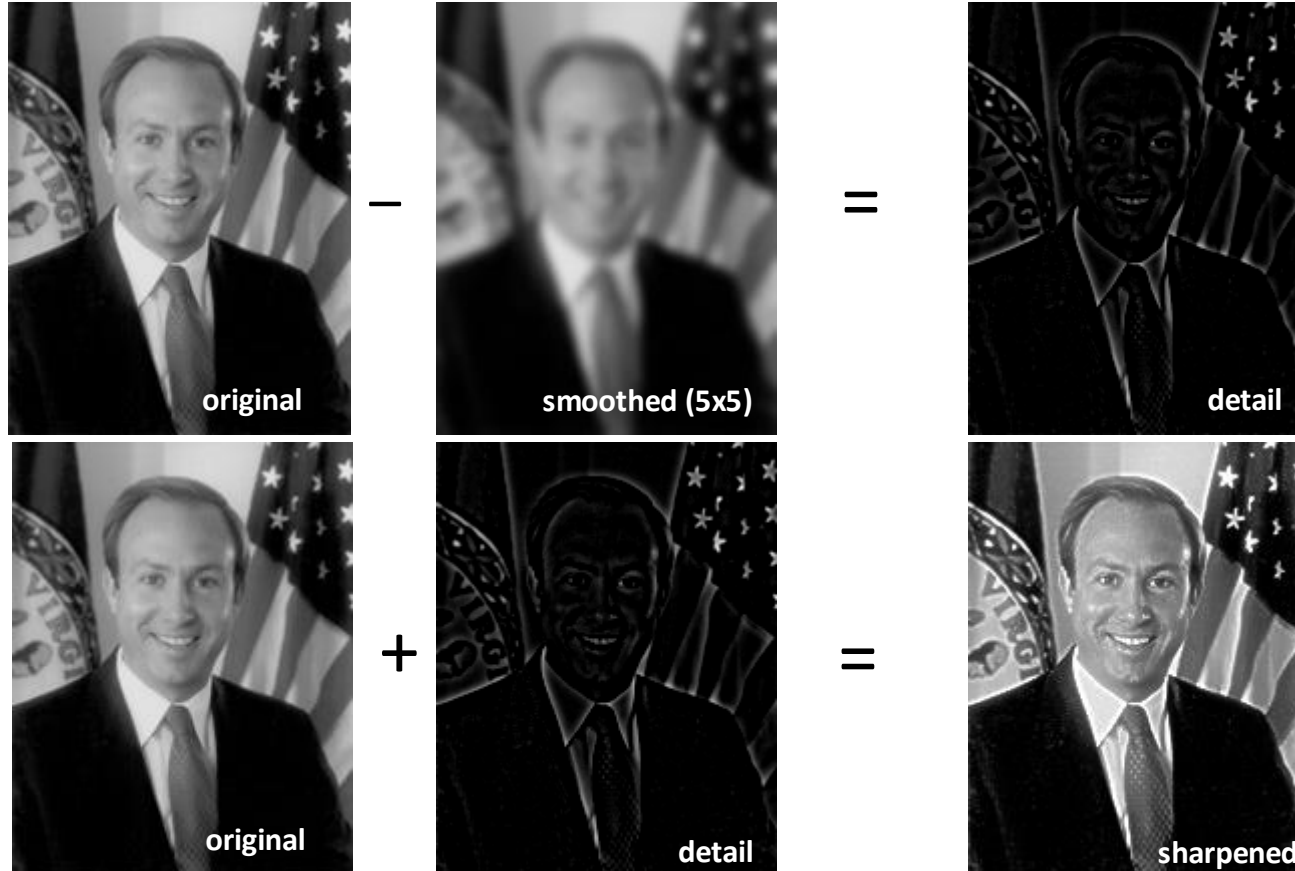


$$G_{\sigma} = \frac{1}{2\pi\sigma^2} e^{-\frac{(x^2+y^2)}{2\sigma^2}}$$

Gaussian filters



Sharpening revisited




What does blurring take away?

CNN Demo with MNIST

Proliferation of Pretrained Models

NEW Create Assistants in HuggingChat



The AI community building the future.

The platform where the machine learning community collaborates on models, datasets, and applications.

Tasks Libraries Datasets Languages Licenses Other

Filter: Tasks by name

Multimodal

- Text-to-Image
- Image-to-Text
- Text-to-Video
- Visual Question Answering
- Document Question Answering
- Graph Machine Learning

Computer Vision

- Depth Estimation
- Image Classification
- Object Detection
- Image Segmentation
- Image-to-Image
- Unconditional Image Generation
- Video Classification
- Zero-Shot Image Classification

Natural Language Processing

- Text Classification
- Token Classification
- Table Question Answering
- Question Answering
- Zero-Shot Classification
- Translation
- Summarization
- Conversational
- Text Generation
- Text2Text Generation
- Sentence Similarity

Audio

- Text-to-Speech
- Automatic Speech Recognition
- Audio to Audio
- Audio Classification
- Voice Activity Detection

Tabular

- Tabular Classification
- Tabular Regression

Reinforcement Learning

- Reinforcement Learning
- Robotics

Models 469,541 Filter by name

- meta-llama/Llama-2-70b
Text Generation • Updated 4 days ago • ± 25.2k • ♥ 64
- stabilityai/stable-diffusion-xl-base-0.9
Updated 6 days ago • ± 2.01k • ♥ 393
- openchat/openchat
Text Generation • Updated 2 days ago • ± 1.3k • ♥ 136
- lillyasviel/ControlNet-v1-1
Updated Apr 26 • ♥ 1.87k
- cerspense/zeroscope_v2_XL
Updated 3 days ago • ± 2.66k • ♥ 334
- meta-llama/Llama-2-13b
Text Generation • Updated 4 days ago • ± 328 • ♥ 64
- tiiuae/falcon-40b-instruct
Text Generation • Updated 27 days ago • ± 288k • ♥ 899
- WizardLM/WizardCoder-15B-V1.0
Text Generation • Updated 3 days ago • ± 12.5k • ♥ 332
- CompVis/stable-diffusion-v1-4
Text-to-Image • Updated about 17 hours ago • ± 448k • ♥ 5.72k
- stabilityai/stable-diffusion-2-1
Text-to-Image • Updated about 17 hours ago • ± 782k • ♥ 2.81k
- Salesforce/xgen-7b-8k-inst
Text Generation • Updated 4 days ago • ± 6.18k • ♥ 57

<https://huggingface.co>

YOLOv8

- YOLOv8 (You Only Look Once, version 8) is the latest iteration in the YOLO series of real-time object detection algorithms. It is designed to achieve high accuracy and speed in detecting multiple objects in images or video streams.

```
from ultralytics import YOLO

# Load a model
model = YOLO("yolov8n.yaml") # build a new model from scratch
model = YOLO("yolov8n.pt") # load a pretrained model (recommended for training)

# Use the model
model.train(data="coco128.yaml", epochs=3) # train the model
metrics = model.val() # evaluate model performance on the validation set
results = model("https://ultralytics.com/images/bus.jpg") # predict on an image
path = model.export(format="onnx") # export the model to ONNX format
```

Classify



Detect



Segment



Track



Pose

