

# CS 331

# Computer Vision

10 / 13 / 2026

Instructor: Michael Eckmann

# Today's topics

- Questions/Comments?
- My solutions to those numpy exercises and some more numpy tips.
- Material in Chapter 5 Machine Learning Basics in Goodfellow, et.al.'s book titled *Deep Learning*
- <https://www.deeplearningbook.org/>
- Specifically  
<https://www.deeplearningbook.org/contents/ml.html>
- from page 96 – middle of page 120
- In the bulleted list of tasks on pages 98-101, you can skip all except: Classification and Regression

# First assignment

- Questions about the first assignment?
- Exercise answers from last class
- Some more numpy tips

# Topics

- Performance
- Experience
- Information necessary to understand linear regression
- Linear Regression
- Capacity, Overfitting and Underfitting

# Dataset usage for supervised learning classification

- For supervised learning classification, we'll have many examples of various classes. e.g. images of cats, dogs, frogs, elephants with their correct labels.
- To create a learned model, we'll usually divide our examples into disjoint sets of training (e.g. 80%), validation (10%) and testing (10%). We'll train on our training set, use the validation set to validate the model --- if not good, restart. Then after we think we have a good model, we try it on our test set once at the end.
- We could use cross-validation where we divide our data into some number of folds (e.g. 5) and use various subsets of them for training and testing --- good for small data sets

# Definition of a machine learning algorithm

- A machine learning algorithm is an algorithm that is able to learn from data.
- Mitchell (1997) “A computer program is said to learn from experience  $E$  with respect to some class of tasks  $T$  and performance measure  $P$ , if its performance at tasks in  $T$ , as measured by  $P$ , improves with experience  $E$ .”

# Performance Measure

- Performance = a quantitative measure of performance
- One example: accuracy rate (or conversely error rate)
- Accuracy rate is the number of accurately predicted test set examples / number of examples tested
- Error rate is the number of incorrectly predicted test set examples / number of examples tested
- The performance measure is done on the *test set*
- Top 5 accuracy --- in datasets with large number of classes (e.g. 1000 classes in ImageNet) model outputs what it thinks are the 5 most likely correct classes for a test image and if the correct class is among the 5, it is considered correct

# Experience

- Experience is the dataset that the learning algorithm is provided. For supervised learning this includes the correct label for each.
- Dataset is made up of examples (aka data points)
- The dataset (or some subset of the dataset) is organized into a design matrix.
- Each row in a design matrix contains data for one example
- Each column in a design matrix is one feature of our data
- Vector of labels (known class representing each example)

# Experience

- Each row can be an image with all the individual pixel values in order --- the row would just be the 2d image converted to 1d where each row of the image is put after each other one.
- e.g. an image that is say 200 rows by 50 columns would be stretched into a 1d vector of size  $200 \times 50 = 10000$  pixel values
- if they were 3 channel color images, then 30000 values
- A column would correspond to a channel of a pixel
- This of course assumes each image is of the same size (same resolution)

# Linear Regression

- Let's look at the linear regression algorithm as an example of a learning algorithm that we can work with to highlight many key concepts
- Regression is where we take in a vector of real values and output a scalar real value (that is a value from among a continuous set of reals)
- Classification on the other hand, takes in a vector of real values and outputs one of a discrete set of labels
- Linear Regression learns a linear function from a set of training examples (data points)

# Linear Regression

- Common concepts that we will get a feel for with linear regression that will apply to more complex learning algorithms
- Weights (aka parameters, learned-parameters)
- Minimizing a loss function (aka cost function, error function, objective function) --- how to do this using gradient
- Measuring performance
- Etc.

# Linear Regression

- First, let's recall what a derivative of a function means
- It's a function that gives us the slope at each point on the original function
- Example:
- $f(x) = x * x$
- Derivative of  $f(x)$ ,  $f'(x) = \text{what?}$

# Linear Regression

- First, let's recall what a derivative of a function means
- It's a function that gives us the slope at each point on the original function
- Example:
- $f(x) = x * x$
- Derivative of  $f(x)$ ,  $f'(x) = 2 * x$

# Linear Regression

- First, let's recall what a derivative of a function means
- It's a function that gives us the slope at each point on the original function
- Example:
- $f(x) = x * x$
- $f'(x) = 2 * x$
- When is the original function  $x * x$  minimum?
- Let me draw the graph of  $x * x$

# Linear Regression

- $f(x) = x^2$
- Derivative  $f'(x) = 2x$
- When is the original function  $x^2$  minimum?
- When  $f'(x) = 0$  is when  $x^2$  is minimum
- When a function is convex like  $x^2$  then this always works.