

E9 205 Machine Learning for Signal Processing

**Introduction to Machine Learning of
Sensory Signals**

14-08-2017

Instructor - Sriram Ganapathy (sriram@ee.iisc.ernet.in)

Teaching Assistant - Aravind Illa (aravindece77@gmail.com).



Overview

- ❖ What are the typical real-world signals
- ❖ What is learning
- ❖ Why should we attempt learning of such signals
- ❖ Roadmap of the course

Real World Signals

- ❖ Signal in general is a function $f : X \rightarrow V$
- ❖ Real World Signals
 - ❖ which we see everyday everywhere
 - ❖ Text, Speech, Image, Videos...
 - ❖ DNA sequence, financial data, weather parameters, neural spike train...
 - ❖ Belonging to / generated by certain category of events.

Real World Signals

- ❖ Types of signals- Continuous and Discrete
- ❖ Observations from real world signals
 - ❖ Information may not be uniform.
 - ❖ Cannot be modeled deterministically.
 - ❖ Affected by noise, sensing equipments.
 - ❖ Missing or hidden variables.

Real World Signals - Examples

- ❖ Text data
- ❖ Discrete sequence of items

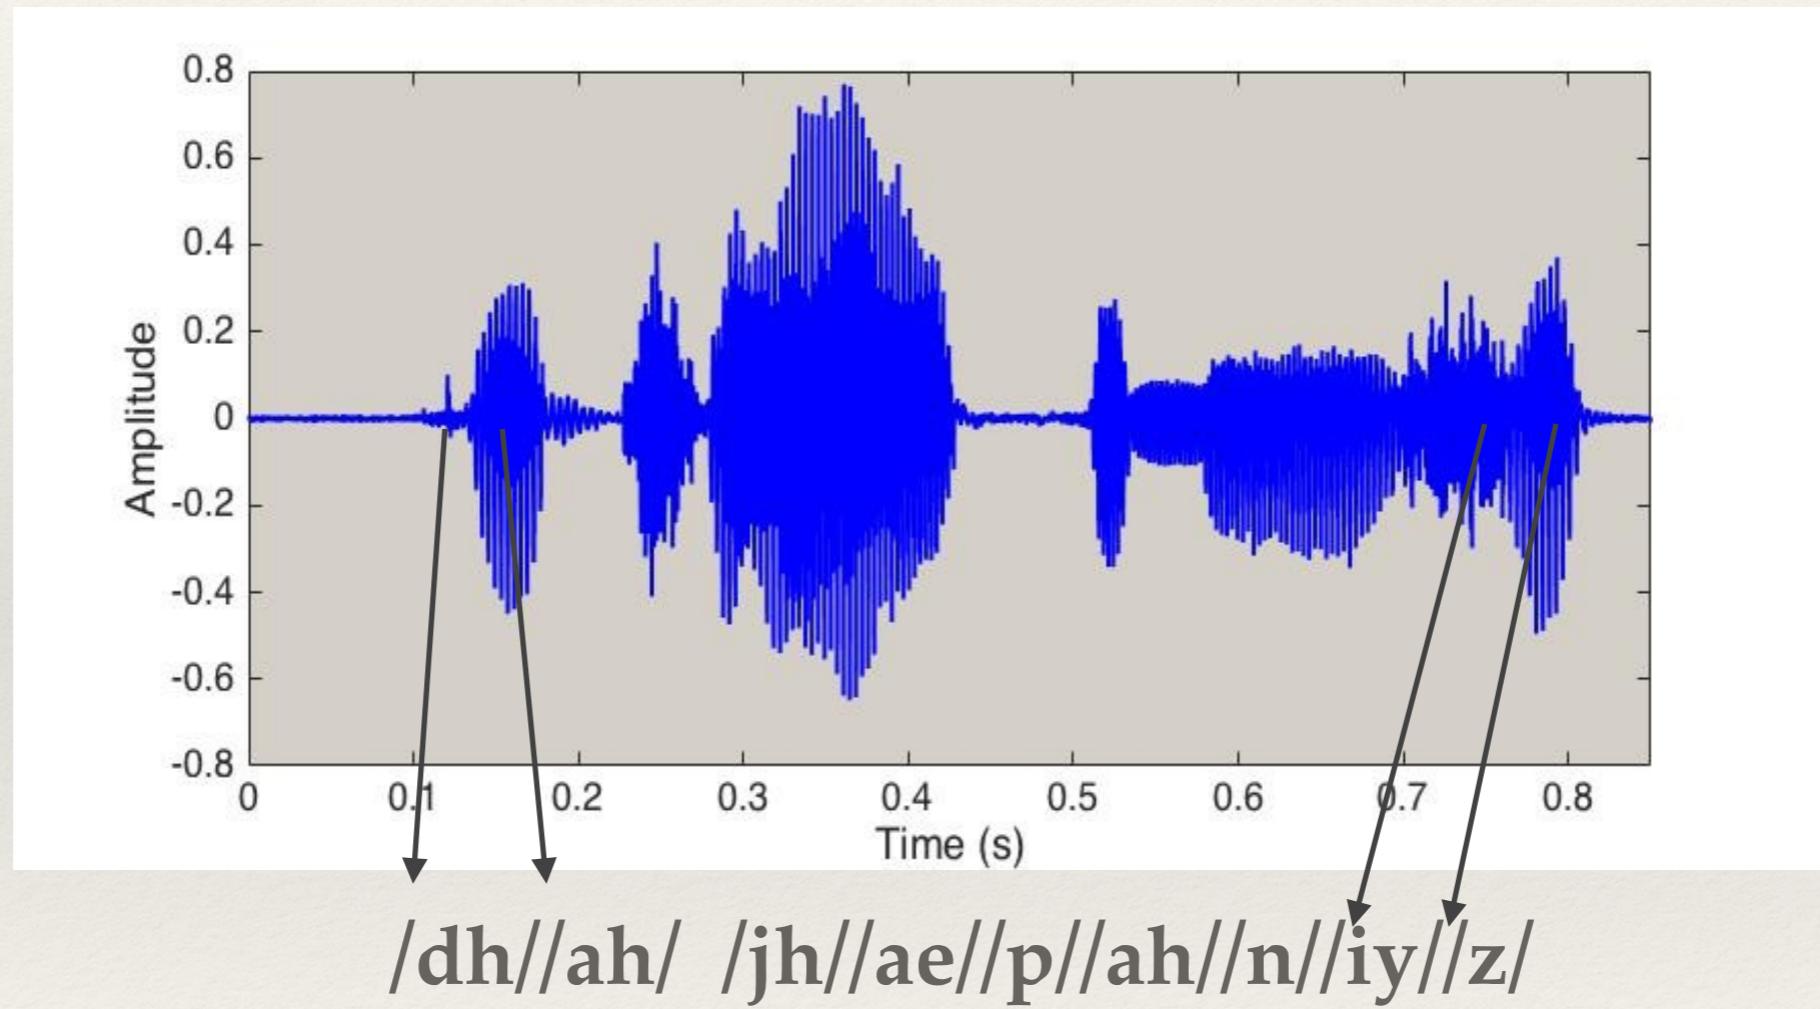
In the last 29 years, sir has never ever said 'well played' to me because he thought I would get complacent and I would stop working hard.

Items - [In] [the] [last] [29] [years]

- ❖ Some items carry more **importance** than others.

Real World Signals - Examples

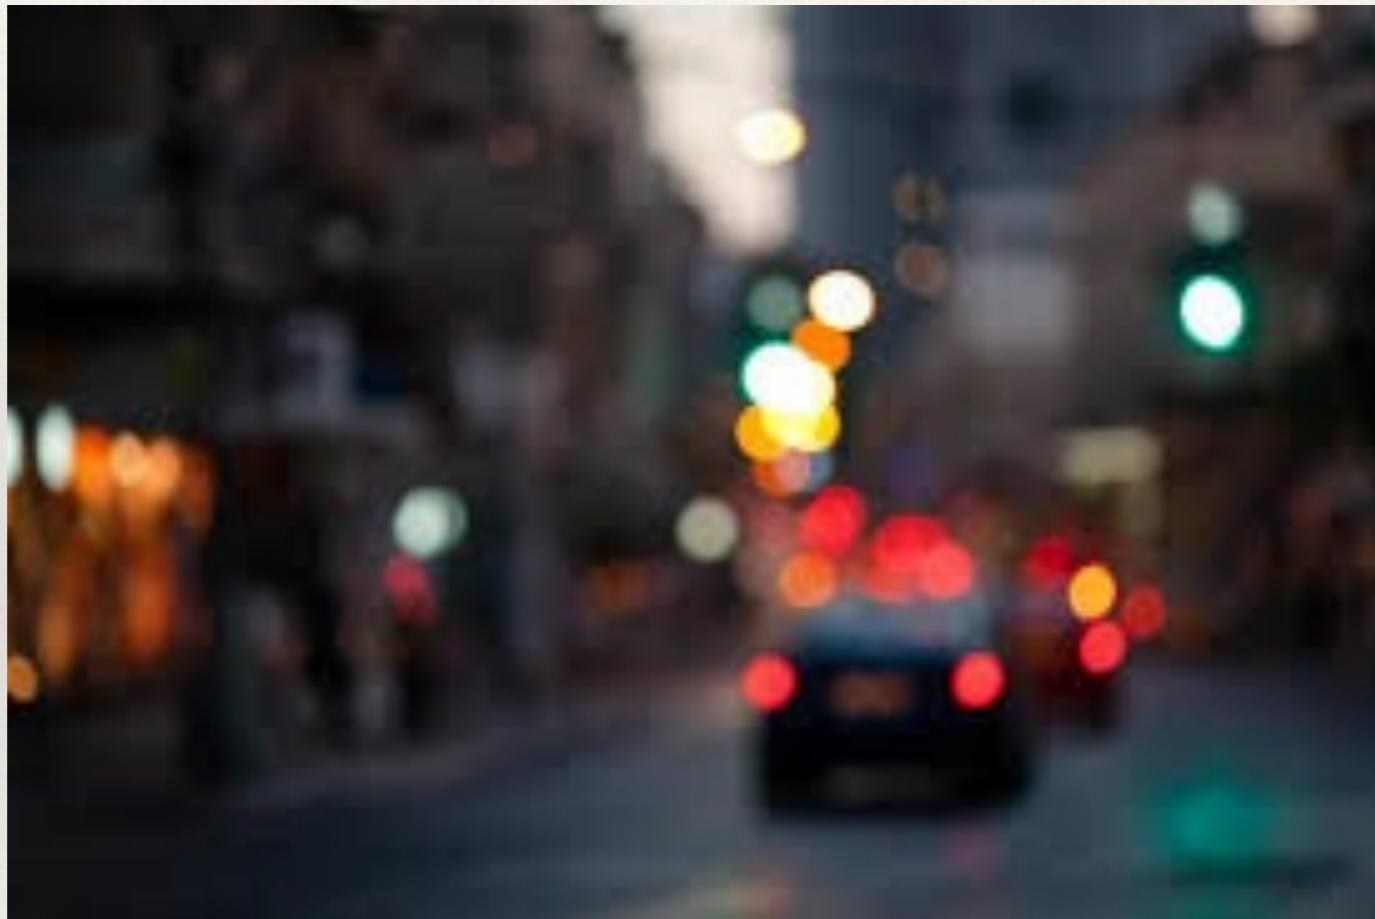
- ❖ Speech data



- ❖ Phonetic units - underlying hidden variables.

Real World Signals - Examples

- ❖ Images



- ❖ Measurement artifacts - noise.

Patterns in Real World Signals

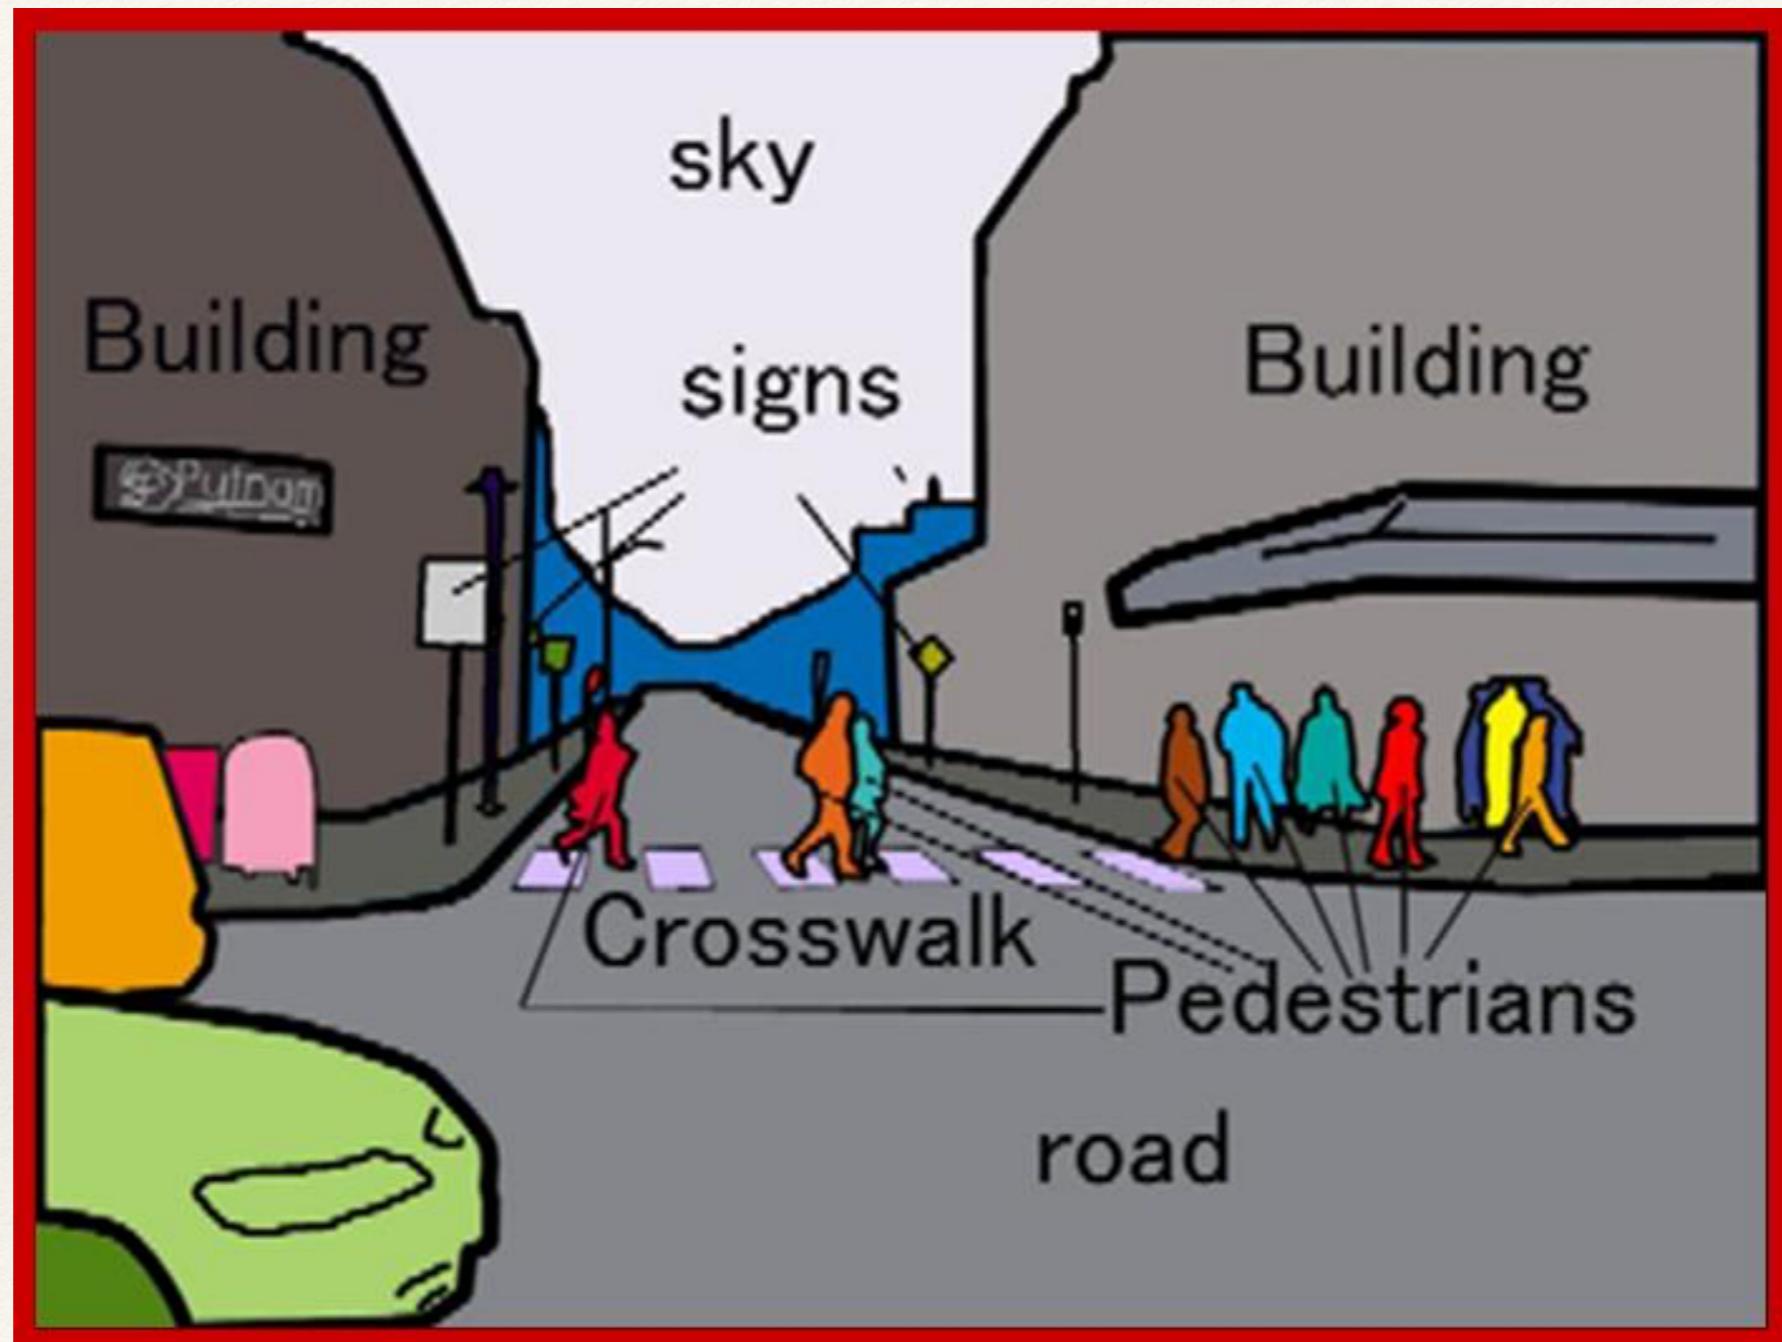
- ❖ Patterns in real world signals
 - ❖ Caused by various generation processes in the real-world signals.
 - ❖ Hidden from the observation.
 - ❖ Value patterns and geometric patterns.
 - ❖ May be hierarchical in nature.
 - ❖ Manifested as pure patterns or transformed / distorted versions.

What is Learning

- ❖ Learning
 - ❖ Process of describing or uncovering the pattern.
 - ❖ Understanding the physical process of generation.
 - ❖ Generalization for prediction, classification, decision making.
 - ❖ Using the data to learn the underlying pattern.
- ❖ Humans are **fundamentally trained** to learn and recognize patterns.

What is Learning

Object
Recognition



What is Learning

Facial Identification



Topic Summarization

The Karnataka government is planning to start an aviation school to help students from lower economic and rural backgrounds become pilots.

Machine Learning

- ❖ Machine Learning
 - ❖ Automatic discovery of patterns.
 - ❖ Motivated by human capabilities to process real world signals.
 - ❖ Mimicking/Extending/Replacing human functions.
 - ❖ Branch of artificial intelligence.
 - ❖ Classification and Regression.

Machine Learning - Examples

Domain Identification - Blog v/s Chat ?

"I tried these Butterscotch Muffins today and they turned out so good. I had half the pack of butterscotch chips that I bought long back so wanted to use it up."

"Hey, it's Geoff from yesterday. How's it going?
Hi there. Don't wanna bother you long, but
you saw this video?"

Machine Learning - Examples

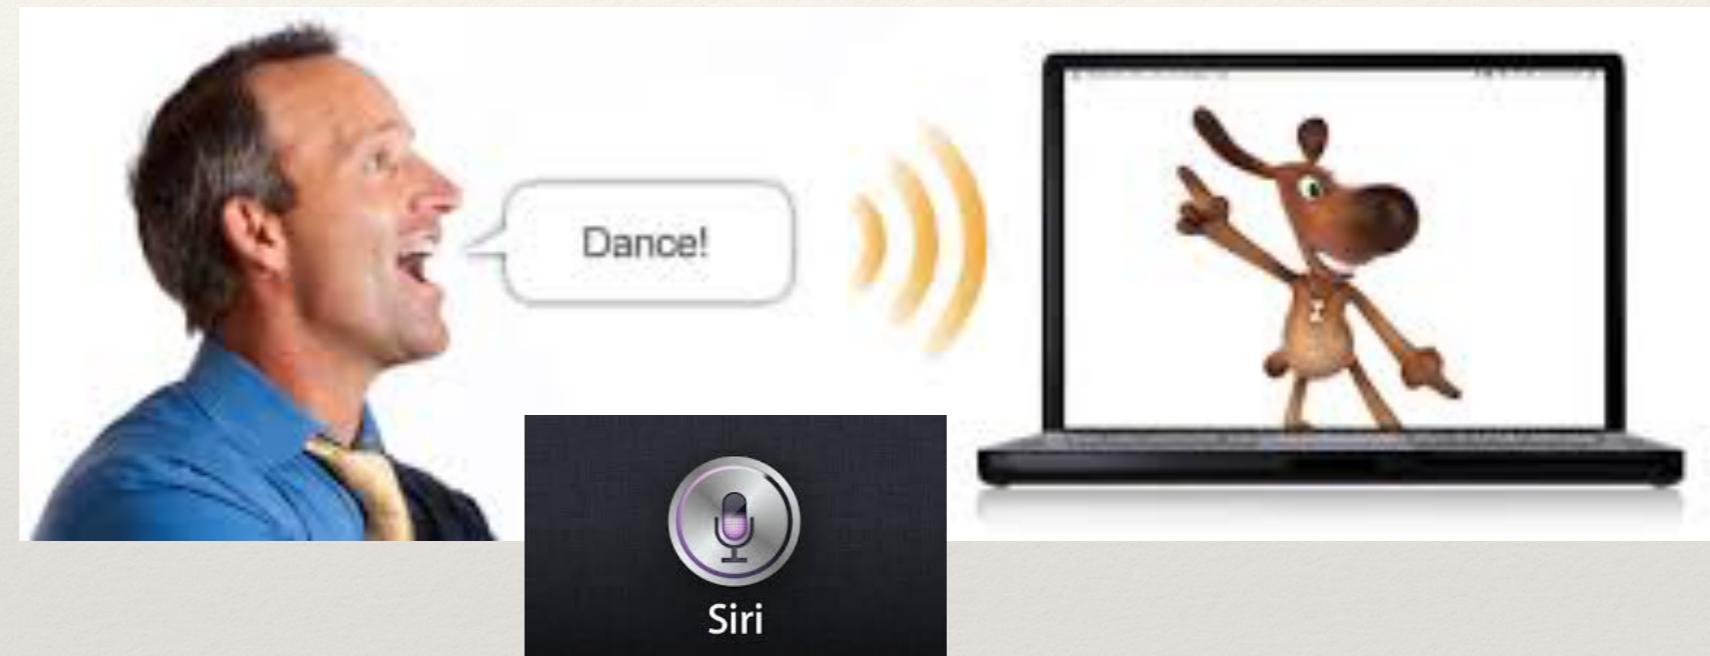
Did a Human or Machine write this ?

“A shallow magnitude 4.7 earthquake was reported Monday morning five miles from Westwood, California, according to the U.S. Geological Survey. The tremor occurred at 6:25 AM, Pacific time at a depth of 5.0 miles.”

“Kitty couldn't fall asleep for a long time. Her nerves were strained as two tight strings, and even a glass of hot wine, that Vronsky made her drink, did not help her. Lying in bed she kept going over and over that monstrous scene at the meadow.”

Machine Learning - Examples

Speech Recognition

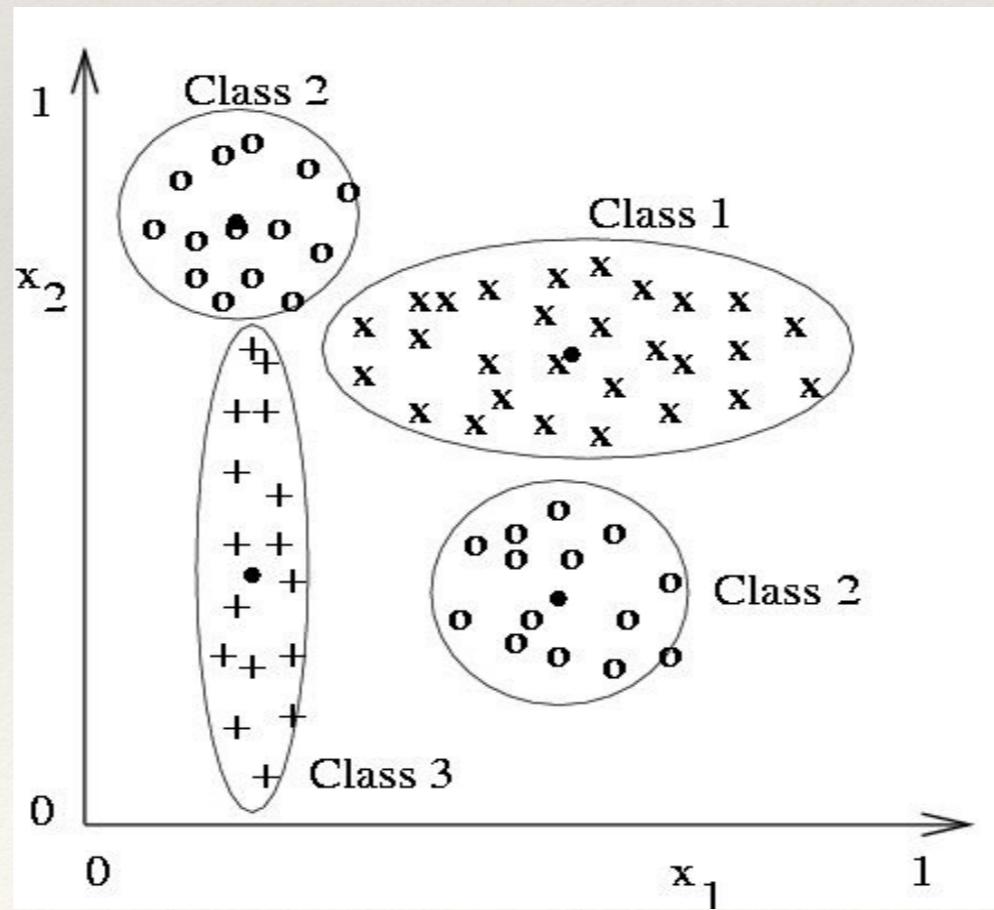


Sound Synthesis

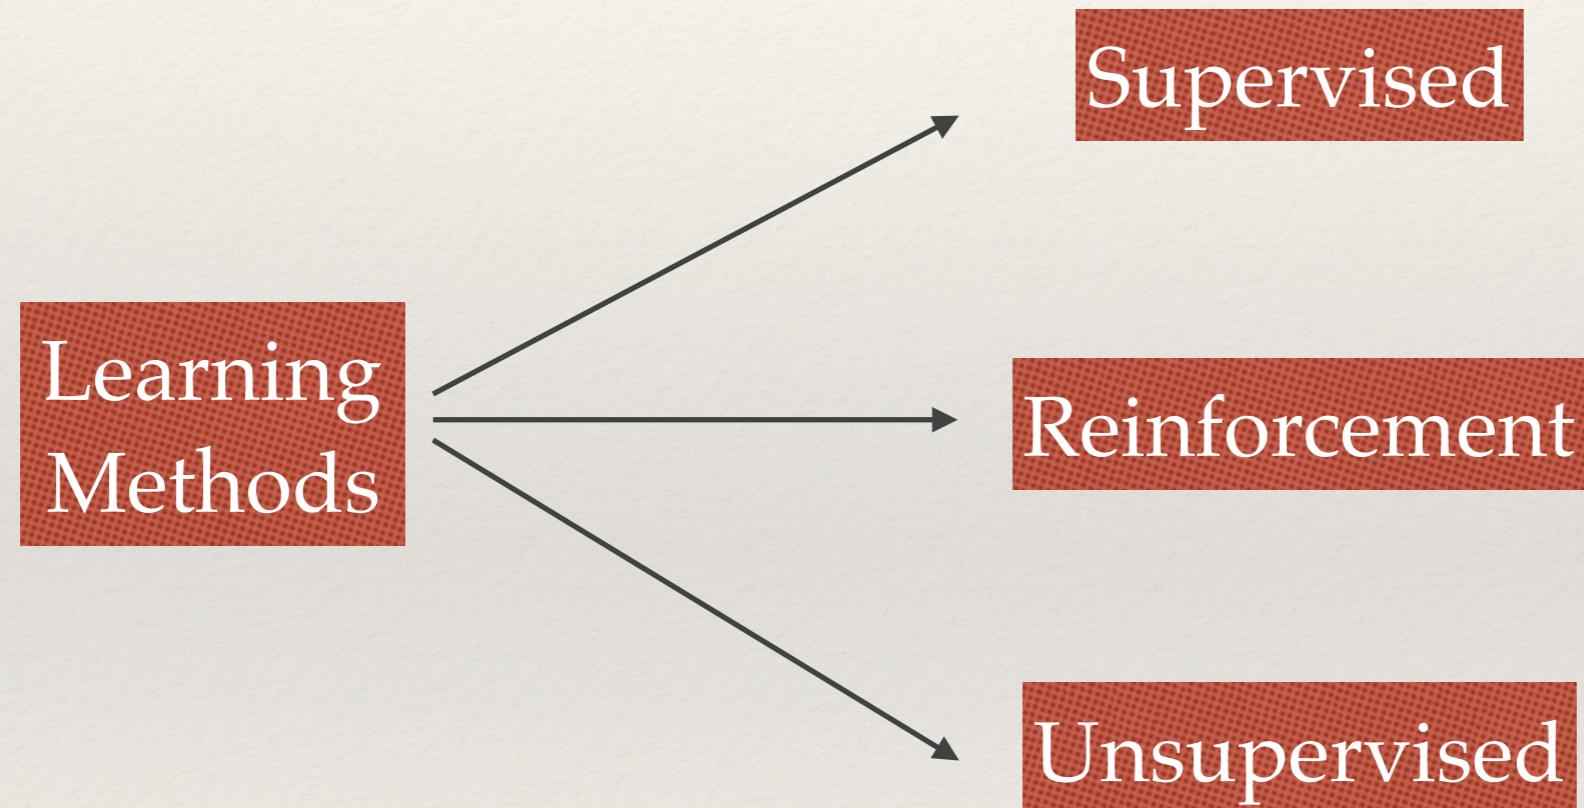
<http://news.mit.edu/2016/artificial-intelligence-produces-realistic-sounds-0613>

Machine Learning

- ❖ Traditional approaches to Machine Learning
 - ❖ Rule and heuristic based methodologies
 - ❖ Using small amounts of data.
- ❖ Recently, most problems are addressed as statistical pattern recognition problem with big data.



Types of Learning



Camstra, Vinciarelli, "Machine Learning for Audio, Image and Video Analysis" 2007.

Unsupervised Learning

- ❖ Data is presented without associated output targets
 - ❖ Extracting structure from the data.
 - ❖ Examples like clustering and segmentation.
 - ❖ Concise description of the data - dimensionality reduction methods.

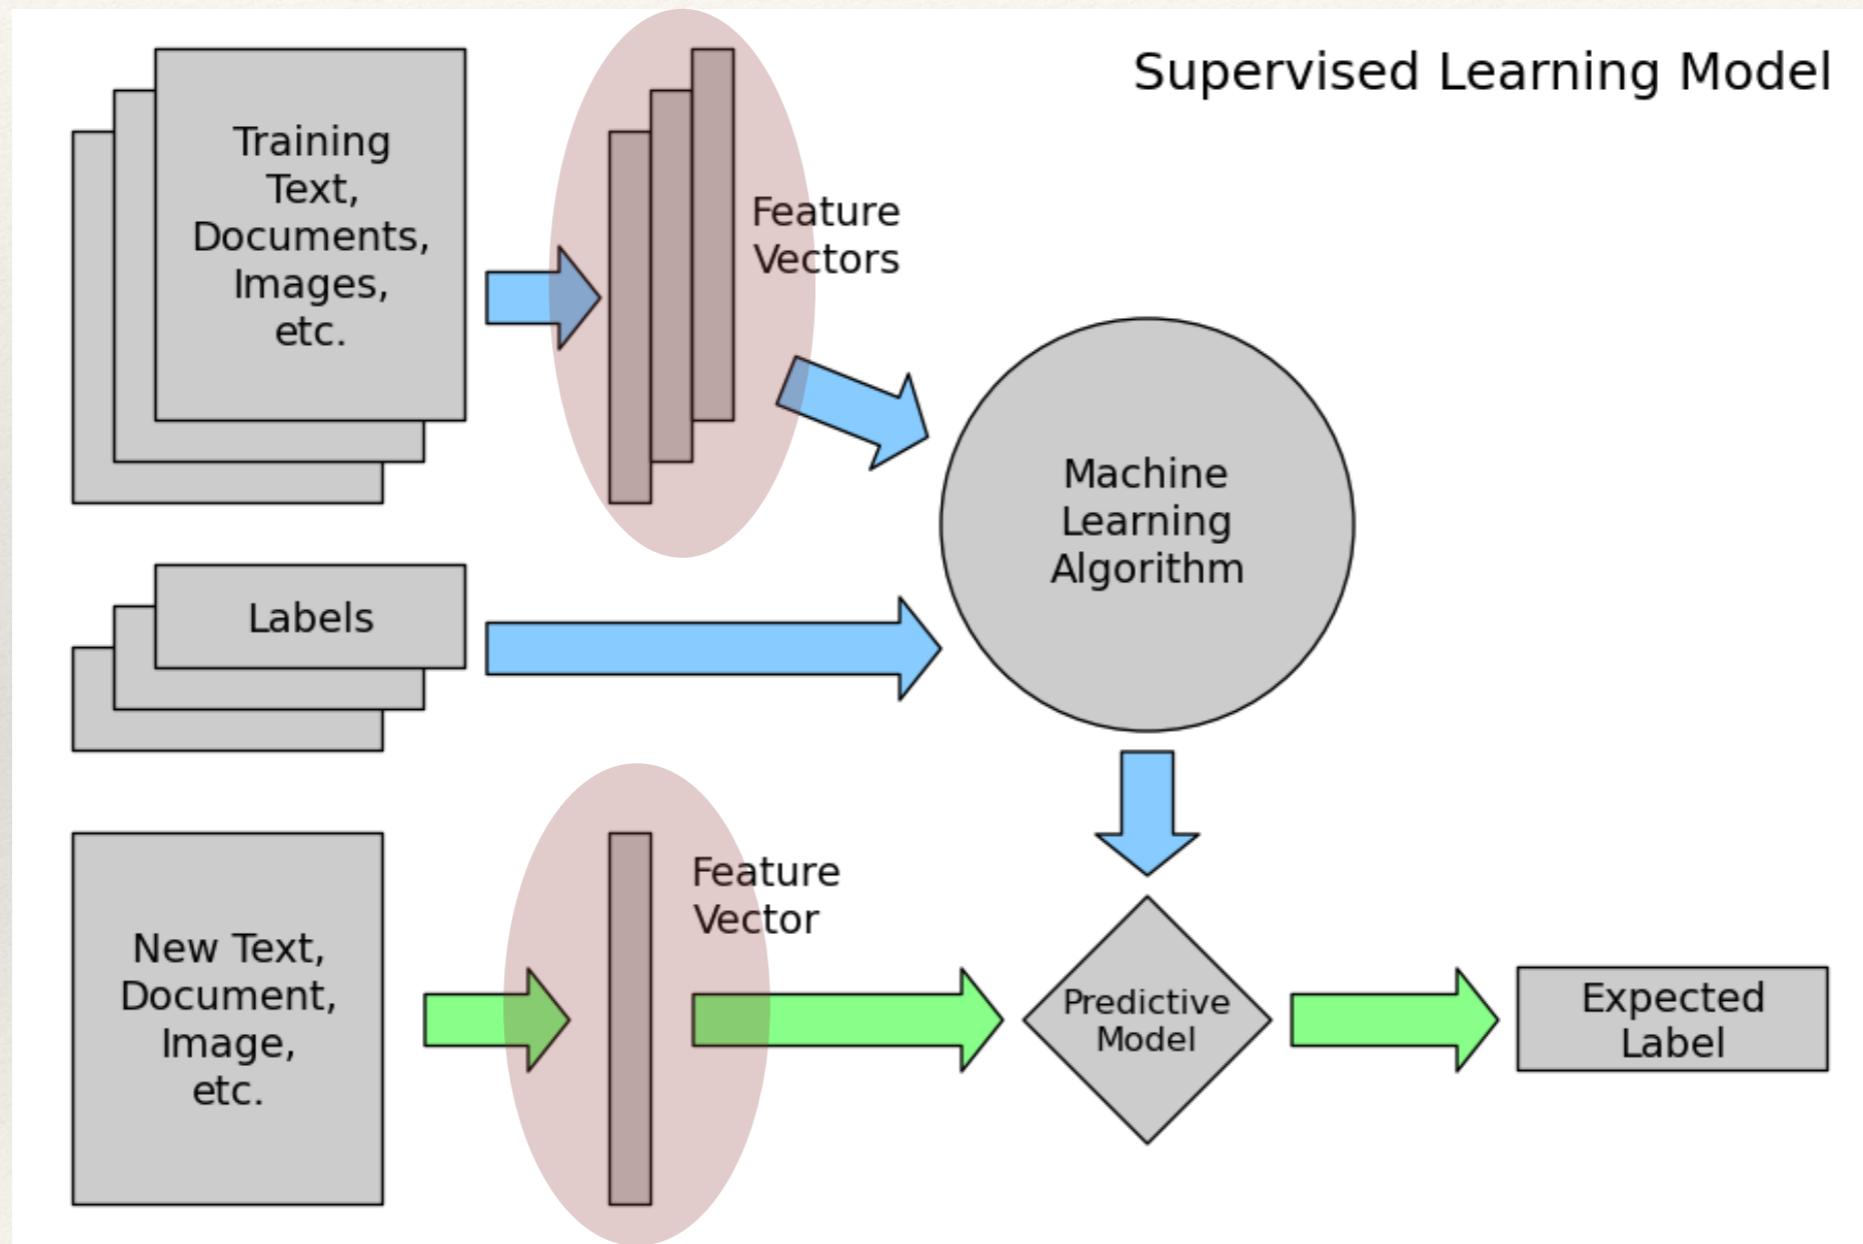
Reinforcement Learning

- ❖ Dynamic environment resulting in triplets - state / action / reward.
 - ❖ No optimal action for a given state
 - ❖ Algorithm has to learn actions in a way such the expected reward is maximized over time.
 - ❖ May also involve minimizing punishment.
 - ❖ Reward / punishment could be delayed - learning based on past actions.

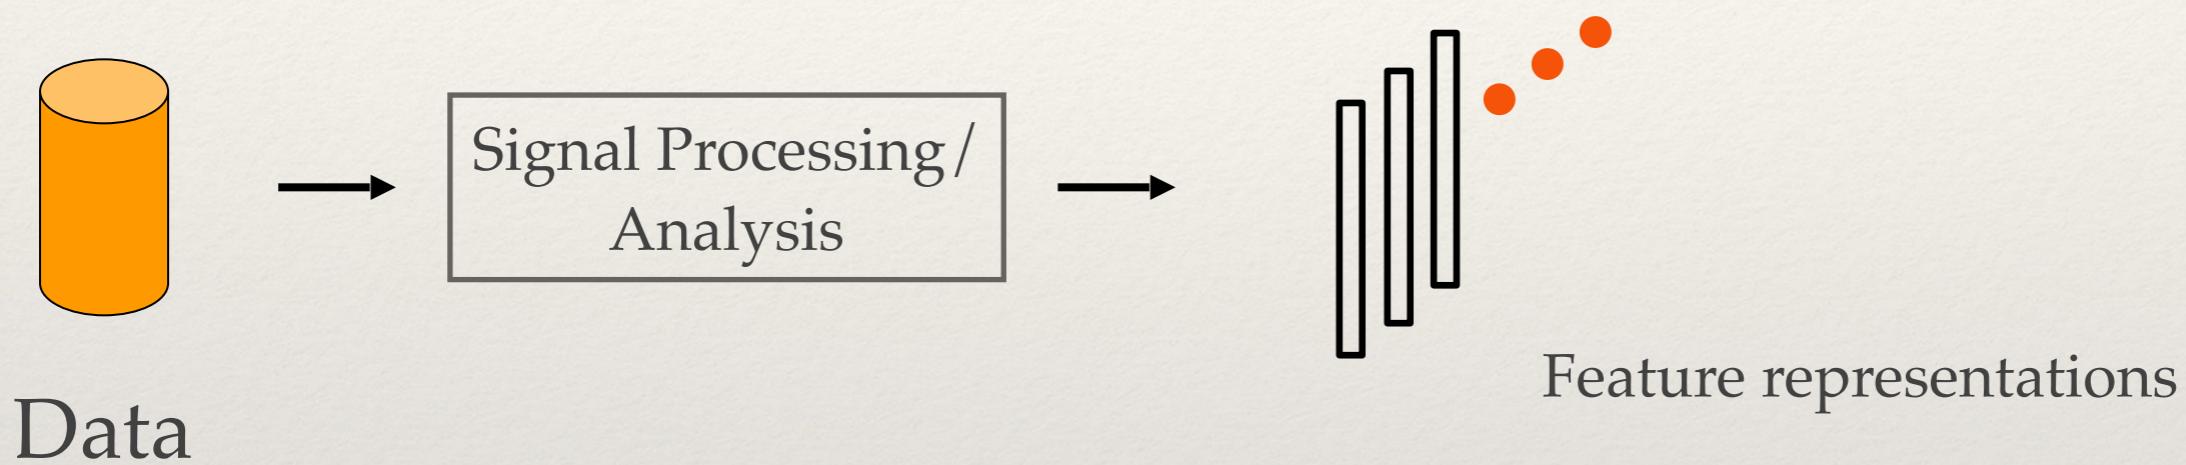
Supervised Learning

- ❖ Training data is provided with along with target values (ground truth).
 - ❖ Goal - to learn the mapping function from data to targets.
 - ❖ Use the mapping function to predict unseen / test data samples.
- ❖ Two types based on the structure of the labels.
 - ❖ Classification - discrete number of classes or categories.
 - ❖ Regression - continuous output variables.

Supervised Learning

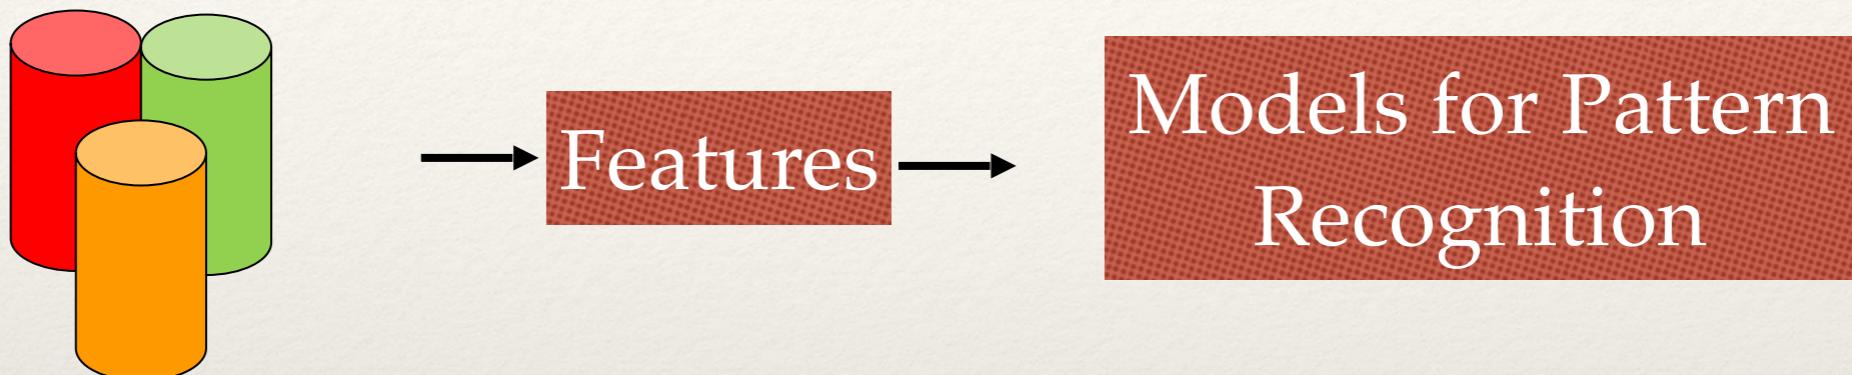


Course Roadmap



- ❖ Feature Extraction from Text, Speech, Image/Video signals.

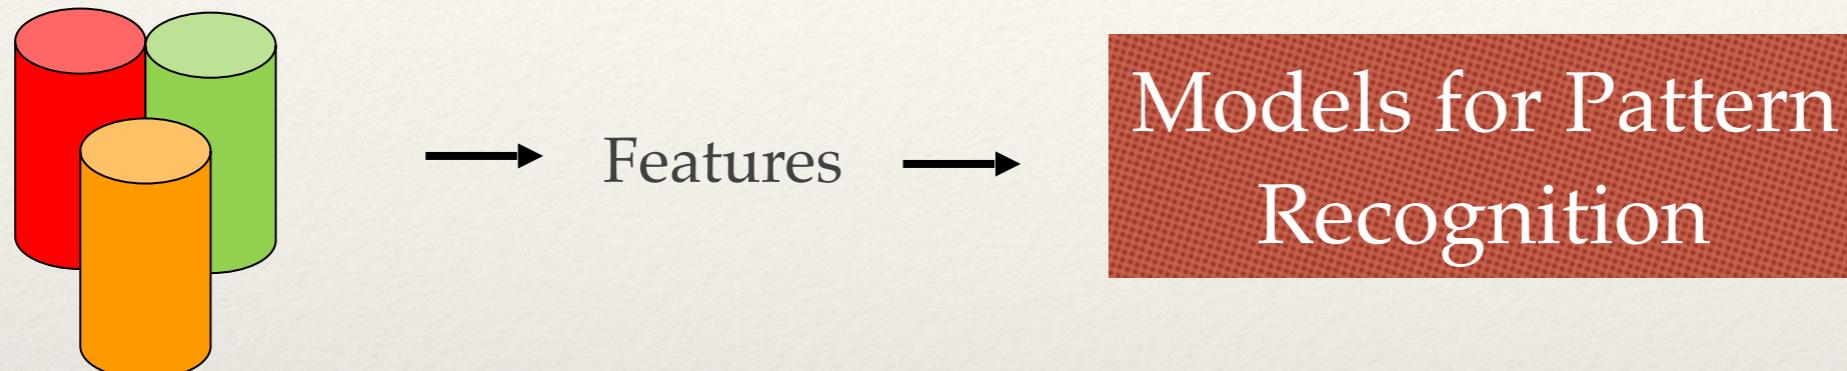
Course Roadmap



Data Set

- ❖ Between features and pattern recognition
 - ❖ Feature selection, dimensionality reduction.
 - ❖ Representation learning.

Course Roadmap



Data Set

- ❖ Modeling the generation of data
 - ❖ Gaussian, Mixture Gaussian, Hidden Markov Models etc.
- ❖ Modeling the separation of data
 - ❖ Support Vector Machines, Deep Neural Networks etc.

Course Structure (Rough Schedule)

- ❖ Signal analysis and processing (2 weeks)
 - ❖ Audio/Speech - spectrograms
 - ❖ Text - TF/IDF, Image feature extraction
- ❖ Basics of Pattern Recognition (1 week).
 - ❖ Dimensionality reduction and feature selection.
- ❖ Generative modeling (3 weeks)
 - ❖ Gaussian and mixture Gaussian modeling, hidden Markov modeling.
- ❖ Discriminative modeling - Support vector machines (2 weeks)
- ❖ Deep Learning (6 weeks)
- ❖ Unsupervised learning (1 week)

Pre-requisites and Course Grading

Requisite	❖ Must <ul style="list-style-type: none">❖ Probability / Random process / Stochastic Models❖ Linear Algebra / Matrix Analysis
Grading	❖ Preferred <ul style="list-style-type: none">❖ Signal Processing❖ Assignments - Theory + Implementation (20%)❖ Mid-terms (20%)
	❖ Project (25%)
	❖ Finals (35%)

Resources and Guidelines

Project and Coding Assignments

- ❖ Coding and submissions in GitHub
 - ❖ Preferred Language - Python.
- ❖ In class demos and example recipes in python.
- ❖ Final Project - GPU platform will be setup

Resources

- ❖ Textbooks -
 - ❖ PRML (Bishop), NN (Bishop).
 - ❖ Deep Learning (Goodfellow)
- ❖ Online resources (papers and other textbooks listed in webpage).

Course Announcements Updates and Dates

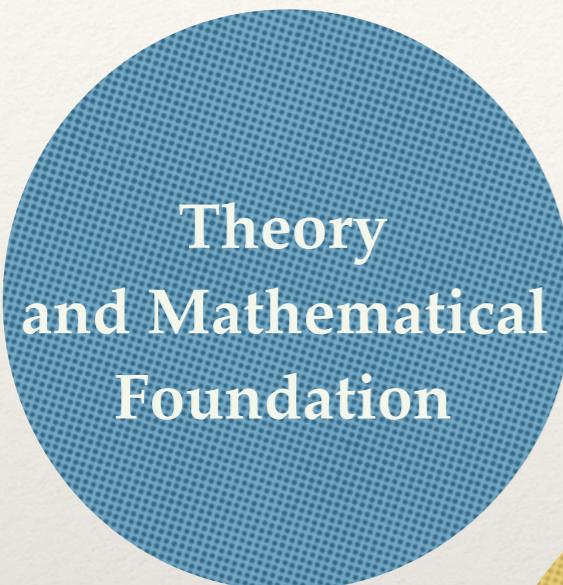
www.leap.ee.iisc.ac.in/sriram/teaching/MLSP_17



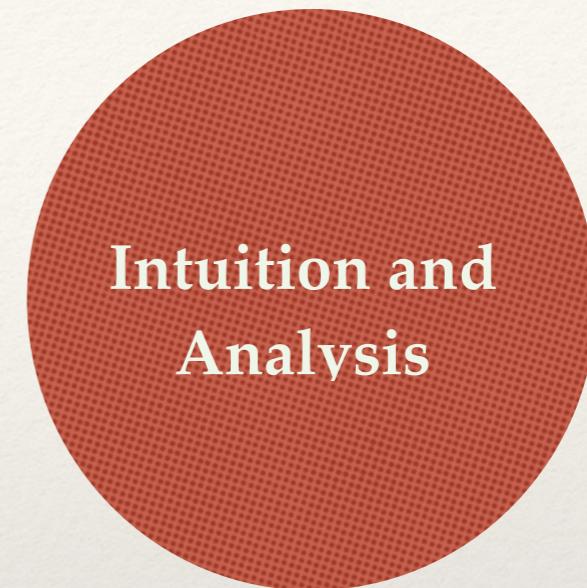
Dates of Rituals

- ❖ 5 assignments spread over 3 months (roughly one assignment every two weeks).
- ❖ September 1st week - project topic announcements
- ❖ September 3rd week - project topic and team finalization. [1 and 2 person teams].
- ❖ September 4th week - 1st Midterm
- ❖ October 1st week - Project Proposal
- ❖ November 1st week - 2nd MidTerm
- ❖ November 3rd week - Project Progress Update
- ❖ December 1st week - Final Exam
- ❖ December 2nd week - Project Final Presentation.

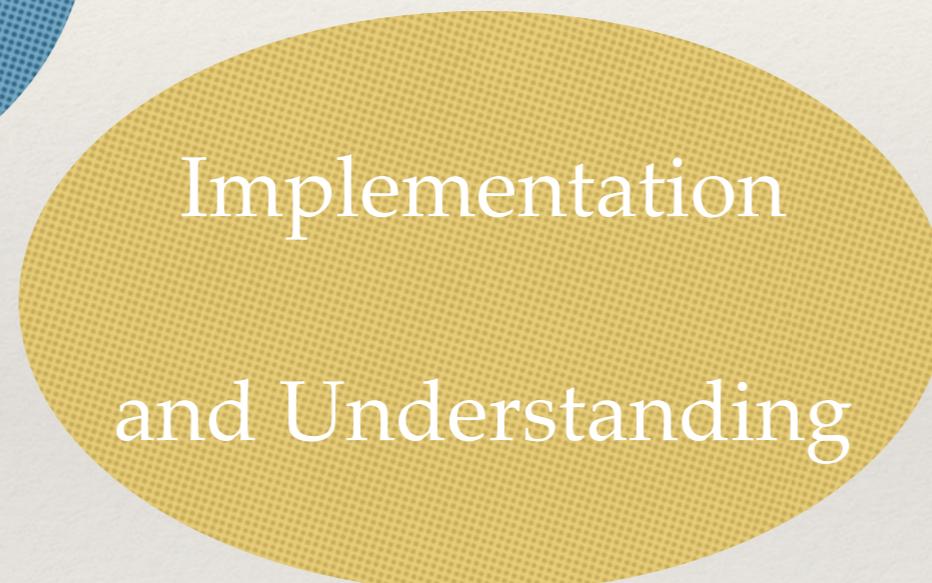
Content Delivery



Theory
and Mathematical
Foundation



Intuition and
Analysis



Implementation
and Understanding

Lecture and
Beyond

- ❖ Teaching Assistant - Aravind Illa (aravindece77@gmail.com).
- ❖ Additional lecture slot on Friday (time ?)
- ❖ Industry research lectures (1-2)

Housekeeping

- ❖ Location shift to C240 (next room).
- ❖ Timing shift to 340-510 pm ?

Questions/Comments ?