

Introduction to Machine Learning

Yixin Lin, Serge Assaad, Rohith Kuditipudi

Duke University

yixin.lin@duke.edu, serge.assaad@duke.edu, rohith.kuditipudi@duke.edu

March 25, 2017


- 1 Overview of basic principles of machine learning
- 2 Introduction to neural networks
- 3 Tutorial on implementing deep learning algorithms

Overview of basic principles of machine learning

- Three components to any ML problem: the **task**, the **performance measure** and the **data**
- Essential definitions
 - **Features**
 - **Model**
 - **Parameters**
 - **Loss**
- Two (broad) kinds of tasks
 - **Supervised learning**: data is labeled/annotated
 - **Unsupervised learning**: data is unlabeled


Example: Linear Regression

- Can we **predict** y given x , using the model $\hat{y} = mx + \mathbf{b}$?




Introduction to Neural Networks

- Neurons are the building blocks of neural networks
- Each neuron is a **function**: $y = f(\mathbf{w}^T \mathbf{x} + \mathbf{b})$




- Note: **neural networks** \neq **neuroscience!!!!**

Neural Networks are Layers of Neurons



What are Neural Networks Good For?

airplane




automobile




bird



cat



deer



dog




frog



horse





ship



truck



Training Neural Networks: Gradient Descent



- Time to **build something!**
- Goals:
 - Give you a basic framework for thinking about what machine learning is and isn't
 - Have a project done or almost done at the end of the workshop
 - Get you set up to continue to tinker and play with machine learning models

Thanks!

Resources and references: yixinlin.net/intro-ml