

Combining Models

Sargur Srihari

srihari@buffalo.edu

Topics in Combining Models

- Overview
- Two methods to combine models
 - Committees and Decision Trees
- Bayesian Model Averaging

Overview of Combining Models

- Many models available in machine learning for classification and regression
- Instead of using one model in isolation improved performance can be obtained by combining different models

Two Methods of Combining


1. Committee: train L different models
 - Make predictions using average of the predictions
 - *Boosting* is a variant of Committee
 - Train multiple models in sequence
 - Error function used to train a model depends on performance of previous models
2. Select one of L models to make the prediction
 - Choice of model is a function of input variables
 - Different models become responsible for different regions of input space
 - *Decision tree* is an example

Decision Tree

- Choice of model is function of input variables
 - Different models become responsible for making predictions in different regions of input space
 - A sequence of binary selections corresponding to traversal of a tree structure
- Decision trees can be applied to both classification and regression problems

Regression Tree with 2 inputs

Two input attributes (x_1, x_2) , a real output



First node asks if $x_1 \leq t_1$

If yes, ask if $x_2 \leq t_2$

If yes, we are at quadrant R_1

Associate a value of y with each region

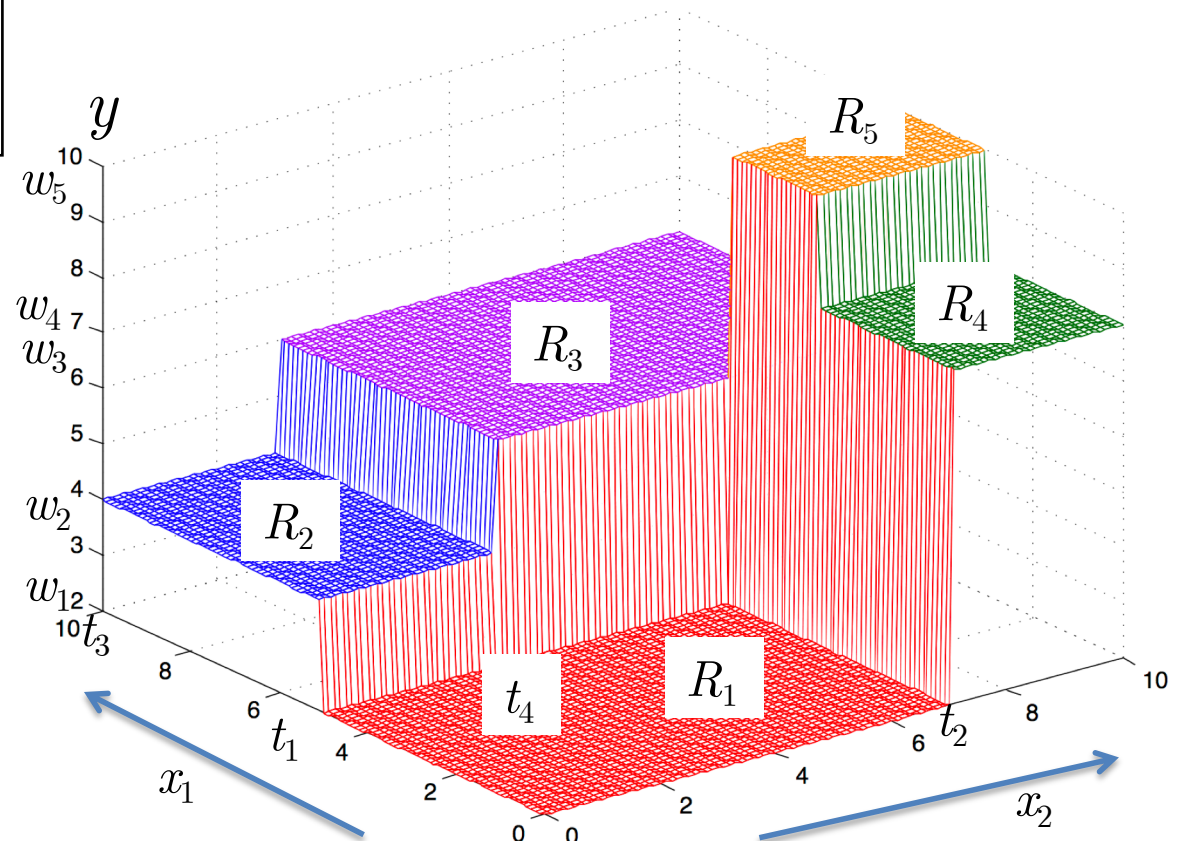
Result of axis-parallel splits:
2-d space partitioned into 5 regions

Each region is associated with a mean response

Result: piecewise constant function

Model can be written as:

$$f(\mathbf{x}) = E[y | \mathbf{x}] = \sum_{m=1}^M w_m I(\mathbf{x} \in R_m) = \sum_{m=1}^M w_m \phi(\mathbf{x}; \mathbf{v}_m)$$



Probabilistic Formulation

- Limitation of Decision Trees: Hard splits of input space
 - Only one model responsible for predictions for any given value of the input variables
- Can be softened by moving to a probabilistic framework for combining models
- If we have K models for the conditional distribution $p(t|\mathbf{x},k)$ where \mathbf{x} is the input and t is the target,

$$p(t|\mathbf{x}) = \prod_{k=1}^K \pi_k(\mathbf{x}) p(t|\mathbf{x},k)$$

where $\pi_k(\mathbf{x}) = p(k|\mathbf{x})$ represents the input-dependent mixing coefficient

- We have a mixture of experts

Model Combination vs Model Averaging

- Data set $X = \{\mathbf{x}_1, \dots, \mathbf{x}_N\}$
- Model Averaging, e.g., Committee
 - Models indexed by $h=1, \dots, L$
 - E.g., one a Gaussian mixture, another a Cauchy mixture
 - Bayesian model averaging

$$p(X) = \sum_{l=1}^L p(X|l)p(l)$$

- Model Combination, e.g., Decision Tree
 - Latent variable z indicates component

$$p(X) = \prod_{n=1}^N \sum_{z_n} p(x_n, z_n)$$