

Tutorial 2

CPSC 340: Machine Learning and Data Mining


Fall 2017

- 1 Decision Tree
 - Decision Stump
 - Decision Tree

- 2 Training, Testing, and Validation Set

Decision Stump

- **Decision stump**: simple decision tree with 1 **splitting rule** based on 1 **feature**.
- Binary example:



- Assigns a label to each leaf based on the **most frequent label**.
- How to find the best splitting rule?
 - score the rules
 - intuitive score: **classification accuracy**.

The Dataset

- X : A matrix
 - each row corresponds to a city
 - first column corresponds to longitude
 - second column corresponds to latitude
- y : class label
 - 1 for blue states, 2 for red states.
- Given new city X_{test}
 - predict label y_{test}

exampleDecisionStump.jl

```
example_decisionStump.jl
1 # Load X and y variable
2 using JLD
3 X = load("citiesSmall.jld","X")
4 y = load("citiesSmall.jld","y")
5
6 # Compute number of objects and number of features
7 (n,d) = size(X)
8
9 #####
10 ### Majority Predictor Model ###
11 #####
12
13 # Fit majority predictor and compute error
14 include("majorityPredictor.jl")
15 model = majorityPredictor(X,y)
16
17 # Evaluate training error
18 yhat = model.predict(X)
19 trainError = sum(yhat .!= y)/n
20 @printf("Error with majority predictor: %.2f\n",trainError);
21
22 #####
23 ### Decision Stump Model #####
24 #####
25
26 # Fit decision stump classifier that uses equalities
27 include("decisionStump.jl")
28 model = decisionStumpEquality(X,y)
29
30 # Evaluate training error
31 yhat = model.predict(X)
32 trainError = sum(yhat .!= y)/n
33 @printf("Error with equality-rule decision stump: %.2f\n",trainError);
34
35 # Plot classifier
36 include("plot2Dclassifier.jl")
37 plot2Dclassifier(X,y,model)
38
39
```

decisionStump.jl

```
decisionStump.jl
1 include("misc.jl") # Includes "mode" function
2
3 # First let's define a "type",
4 # specifying the functions/variables we want the stump to have
5 type StumpModel
6     predict # Function that makes predictions
7     split # Function that splits data
8     baseSplit # Set this to one stump doesn't split
9 end
10
11 function decisionStumpEquality(X,y)
12     # Fits a decision stump based on equality after rounding to nearest integer
13
14     # Get the size of the data matrix
15     (n,d) = size(X)
16
17     # Round all the data to the nearest integer
18     X = round.(X)
19
20     # Initial the best rule with the baseline rule (no split)
21     y_mode = mode(y)
22     minError = sum(y .!= y_mode);
23     splitVariable = [];
24     splitValue = [];
25     splitYes = y_mode;
26     splitNo = [];
27
28     # Search for the best rule
29     # (Uses 0(n^2d) approach to keep code simple)
30     yhat = zeros(n)
31     for j in 1:d
32         # Try unique values of column as split values
33         for val in unique(X[:,j])
34
35             # Test whether each object satisfies equality
36             yes = X[:,j] .== val
37
38             # Find correct label on both sides of split
39             y_yes = mode(y[yes])
40             y_no = mode(y[.!yes])
41
42             # Make predictions
43
44             # Make predictions
45             yhat[yes] = y_yes
46             yhat[.!yes] = y_no
47
48             # Compute error
49             trainError = sum(yhat .!= y)
50
51             # Update best rule
52             if trainError < minError
53                 minError = trainError
54                 splitVariable = j
55                 splitValue = val
56                 splitYes = y_yes
57                 splitNo = y_no
58             end
59         end
60     end
61
62     # Now that we have the best rule,
63     # let's build our splitting function
64     function split(Xhat)
65         (t,d) = size(Xhat)
66         Xhat = round.(Xhat)
67         if isempty(splitVariable)
68             return fill(true,t)
69         else
70             return Xhat[:,splitVariable] .== splitValue
71         end
72     end
73
74     # Now that we have the best rule,
75     # let's build our predict function
76     function predict(Xhat)
77         (t,d) = size(Xhat)
78         yes = split(Xhat)
79         yhat = fill(splitNo,t)
80         yhat[yes] = splitYes
81         return yhat
82     end
83
84     return StumpModel(predict,split,isempty(splitNo))
85 end
```

Decision Tree

- **Decision stumps** have **only 1 rule** based on **only 1 feature**.
 - Very limited class of models: usually **not very accurate** for most tasks.
- **Decision trees** allow **sequences of splits** based on **multiple features**.
 - Very general class of models: can get very **high accuracy**.
 - However, it's computationally **infeasible** to find the **best** decision tree.
- Most common decision tree learning algorithm in practice:
 - **Greedy recursive splitting**.

exampleDecisionTree.jl and decisionTree.jl

example_decisionTree.jl


```
1 # Load X and y variable
2 using JLD
3 X = load("citiesSmall.jld","X")
4 y = load("citiesSmall.jld","y")
5 n = size(X,1)
6
7 # Fit a decision tree and compute error
8 include("decisionTree.jl")
9 depth = 2
10 model = decisionTree(X,y,depth)
11
12 # Evaluate training error
13 yhat = model.predict(X)
14 trainError = sum(yhat .!= y)/n
15 @printf("Error with depth=%d decision tree: %.3f\n",depth,trainError)
16
17 # Plot classifier
18 include("plot2Dclassifier.jl")
19 plot2Dclassifier(X,y,model)
20
```

decisionTree.jl

```
1 include("decisionStump.jl")
2
3 function decisionTree(X,y,depth)
4     # Fits a decision tree using greedy recursive splitting
5     # (recursion to make the code simpler)
6
7     (n,d) = size(X)
8
9     # Learn a decision stump
10    splitModel = decisionStump(X,y)
11
12    if depth <= 1 || splitModel.baseSplit
13        # Base cases where we stop splitting:
14        # - this stump gets us to the max depth
15        # - this stump doesn't split the data
16        return splitModel
17    else
18        # Use the decision stump to split the data
19        yes = splitModel.split(X)
20
21        # Recursively fit a decision tree to each split
22        yesModel = decisionTree(X[yes,:],y[yes],depth-1)
23        noModel = decisionTree(X[.!yes,:],y[.!yes],depth-1)
24
25        # Make a predict function
26        function predict(Xhat)
27            (t,d) = size(Xhat)
28            yhat = zeros(t)
29
30            yes = splitModel.split(Xhat)
31
32            yhat[yes] = yesModel.predict(Xhat[yes,:])
33            yhat[.!yes] = noModel.predict(Xhat[.!yes,:])
34            return yhat
35        end
36
37        return GenericModel(predict)
38    end
39 end
40
41
```



Rewriting a decision tree using if/else statements

- Decision tree:



Rewriting a decision tree using if/else statements

- Decision tree:



- If-else statement:

```
if X(i, 98) ==1
  if X(i, 6) ==1
    return 2
  else
    return 1
  end
else
  if X(i, 89) ==1
    return 2
  else
    return 4
  end
end
```

Training, Testing, and Validation Set

- Given **training data**, we would like to learn a model to **minimize** error on the **testing data**
- How do we decide decision tree depth?
- We care about test error.
- But we can't look at test data.
- So what do we do?????

Training, Testing, and Validation Set

- Given **training data**, we would like to learn a model to **minimize** error on the **testing data**
- How do we decide decision tree depth?
- We care about test error.
- But we can't look at test data.
- So what do we do?????
- One answer: **Use part of your train data to approximate test error.**
- Split training objects into **training set** and **validation set**:
 - **Train model** on the **training data**.
 - **Test model** on the **validation data**.

Cross-Validation

- Isn't it wasteful to only use part of your data?
- **k-fold cross-validation**:
 - Train on $k-1$ folds of the data, validate on the other fold.
 - Repeat this k times with different splits, and average the score.




Figure 1: Adapted from Wikipedia.

- Note: if examples are ordered, split should be random.

Problem: 2-Fold Cross Validation

- Modify the code below to compute the 2-fold cross-validation scores on the training data alone.
- Find the depth that would be chosen by cross-validation.

```
% Load X and y variable
load newsgroups.mat
[N,D] = size(X);
T = length(ytest);
depth = 5;
model = decisionTree(X,y,depth);
yhat = model.predictFunc(model,X);
errorTrain = sum(yhat ~= y)/N;
yhat = model.predictFunc(model,Xtest);
errorTest = sum(yhat ~= ytest)/T;
```

Solution: 2-Fold Cross Validation

```
% Load X and y variable
load newsgroups.mat
[N,D] = size(X);
Xtest = X (floor(N/2) + 1 :N , : );
ytest= y (floor(N/2) +1 : N) ;
X = X ( 1:floor(N/2) , : ) ;
y = y (1: floor(N/2));
mindepth = -1 ; minError = Inf;
for depth =1 :15
    errorTrain = 0; errorTest = 0;
    for i =1:2
        [N,D] = size(X);
        T = length(ytest);
        model = decisionTree(X,y,depth);
        yhat = model.predictFunc(model,X);
        errorTrain = errorTrain +sum(yhat ~= y)/N;
        yhat = model.predictFunc(model,Xtest);
        errorTest = errorTest + sum(yhat ~= ytest)/T;
        [X, Xtest]=mySwap(Xtest, X);
        [y,ytest] = mySwap(ytest,y) ;
    end
    disp(errorTest/2 ) ;
    if errorTest/2 < minError
        minError= errorTest/2;
        mindepth = depth;
    end
end
disp(minError); disp(mindepth);
```