Section B

Thursday, January 23, 2025 6:52 PM

Recitation 2
Decision Trees

10-301/10-601: Introduction to Machine Learning

01/24/2024

1 Programming: Tree Structures and Algorithms

Topics Covered:

- Depth of nodes and trees
- Recursive traversal of trees
  - Depth First Search
    - Pre-order Traversal
    - In-order Traversal
    - Post-order Traversal
  - Breadth First Search (Self Study)
- Debugging in Python

→ node

Questions: Path - The edges traversed btw any 2 nodes

1. Depth of a tree definition

   Longest path length from root to leaf

2. Depth of a node definition

   Length of path btw root and an input node

3. What is the depth of tree A? What is the depth of node X_4 in tree A?

   3 2

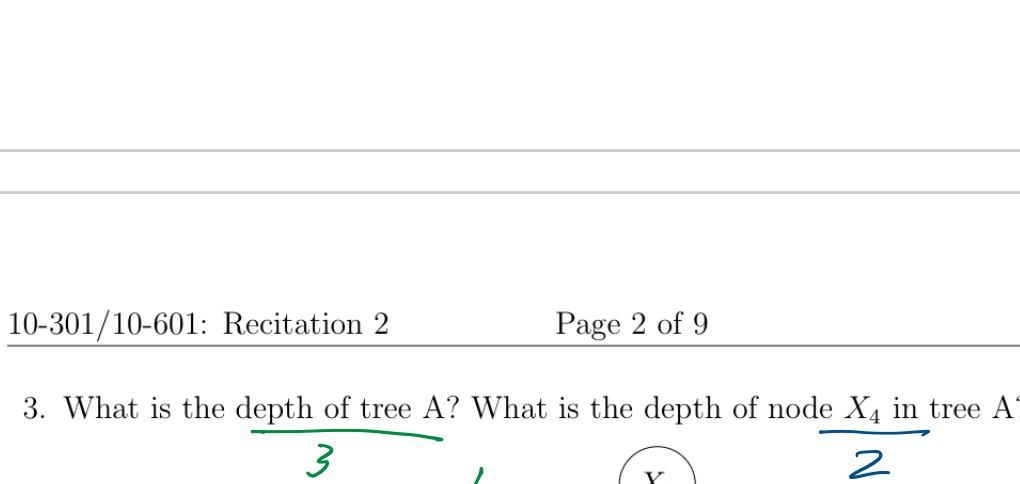

4. What is the depth of tree B?

   1

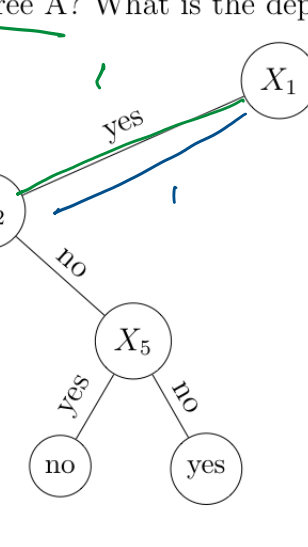

→ Decision Stump

5. What is the depth of tree C? What are the depths of nodes X_1 and X_5 in tree C?

   4 0 2

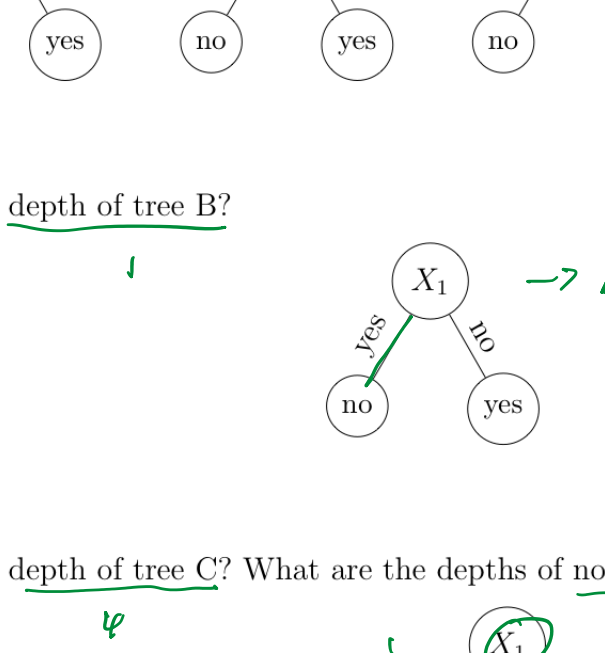

6. In-class coding and explanation of Depth First Traversal in Python.
   Link to the code:
   https://colab.research.google.com/drive/1VvNZUQ4ZikQXcvWL-EY1OPnGdye2P024?usp=sharing

DFS Tree Traversals and Printing

```
# This class represents an individual node
class Node:
    def __init__(self, key):
        self.left = None
        self.right = None
        self.val = key

def traversal1(root):
    if root is not None:
        # First print the data of node
        print(root.val, end='\t')
        # Then recurse on left child
        traversal1(root.left)
        # Finally recurse on right child
        traversal1(root.right)

def traversal2(root, __[a]__):
    if root is not None:
        # First recurse on left child
        traversal2(root.left, __[b]__)
        # Then print the data of node
        print(f'({root.val}, {__[c]__})')
        # Now recurse on right child
        traversal2(root.right, __[d]__)

def build_a_tree():
    root = Node(1)
    root.left = Node(2)
    root.right = Node(3)
    root.left.left = Node(4)
    root.left.right = Node(5)
    return root

if __name__ == '__main__':
    root = build_a_tree()
    print('traversal1 of the binary tree is: ')
    traversal1(root)
    print()
    print('traversal2 of the binary tree is: ')
    traversal2(root, 0)
```

(5) self.val = 5

→ BC → stop at a leaf

[a] depth; [b] depth + 1; [c] depth; [d] depth + 1

D = 0

D = 1

D = 2

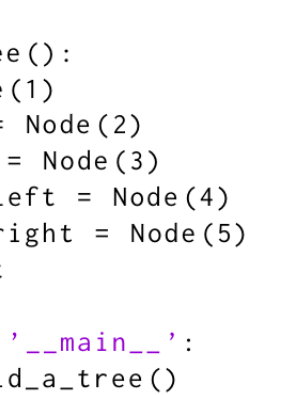

First identify which traversal function is pre-order, in-order, or post-order DFS:

- traversal1() is
- traversal2() is

Fill-in-the-Blanks: Next, fill in the code for traversal2.

Code Output:

pre-order `traversal1 of the binary tree is:` 1 2 4 5 3

in-order `traversal2 of the binary tree is:` (4,2), (2,1), (5,2), (1,0), (3,1)

2 The Need For Speed: Vectorization and Numpy

Performing mathematical operations on vectors and matrices is ubiquitous in most machine learning algorithms. Whether it's a simple similarity measure that works by calculating the dot product between two vectors, or deep neural networks, they all involve repeated matrix operations. This makes it imperative that our underlying code design to perform matrix operations is efficient.

2.1 The Perils of Python

While Python is widely the language of choice for machine learning researchers across the globe (thanks to the speed of development and code readability it offers and the support it enjoys from the open-source community), Python as a high-level language on average is much slower than a lower level language like C++. To combat this, libraries like numpy and scipy implement most of the back-end operations they perform in C/C++, while providing wrappers in Python to be able to call underlying C code seamlessly from a Python script.

2.2 Speed Comparison: Numpy and Python

We highly recommend you to use *numpy* extensively in this course, it will be difficult to pass the programming portion of Homework 4 without writing most of your matrix operations in numpy. In this section, we'll see why.

Consider you have two vectors $\mathbf{a}, \mathbf{b} \in \mathbb{R}^n$. To see how similar they are, as measured by the cosine angle between them, you want to compute their dot product. This translates to the following operation:

$$\mathbf{a} \cdot \mathbf{b} = a_1b_1 + a_2b_2 + \ldots + a_nb_n$$

When translated to code, notice how the dot product in NumPy is a whopping 100x faster than the native Python!

```
from timeit import timeit
import numpy as np
import array

VECTOR_SIZE = int(1e8)

# NumPy arrays
a = np.random.rand(VECTOR_SIZE)
b = np.random.rand(VECTOR_SIZE)

# Python arrays
aArr = array.array('d', a)
bArr = array.array('d', b)
```

np.array([, ----])

numpy

```
def test_np():
    return np.dot(a, b)

# faster than multiprocessing, python lists, or numpy arrays with
   python loops
# faster than using a range and indexing
def test_py_arr():
    return sum(x * y for x, y in zip(aArr, bArr))

def time_dot_product(f):
    return timeit(f, setup=f, number=5) / 5


if __name__ == "__main__":
    print(f"NumPy = {time_dot_product(test_np):.2f}")
    print(f"Python on an array = {time_dot_product(test_py_arr):.2f}")
```

.py

0.05s

5.45s

2.3 Useful Numpy Operations

Some operations in numpy that you will find really useful in your assignments are:

- np.matmul: Matrix multiplication of two matrices
- np.unique: Returns unique elements along an axis. → [1,2,3,4,2,3]
- np.hstack: Stack two arrays horizontally (column-wise) → [[1,2,3,4], [4,2,2,1]]
- np.expand_dims: Convert a row vector of size n into a matrix of size $n * 1$ or $1 * n$
- np.log, np.sum, np.exp, @, .T, and so on.

You can read C vs. Python for more details, and you can also read these two tutorials (beginner, intermediate) from the official numpy website. For instance, understanding broadcasting is recommended. It will help you debug the shape errors you might face in all future homeworks.

vstack → add a row / example

hstack → add a new feature

[1 2 3] 3 × 1 (3, _)

dim=0 ↓ dim=1 →

hstack([1 2 3], [4 5 6]) = [1, 2, 3, 4, 5, 6]

hstack([[1],[2],[3]], [[4],[5],[6]]) = [[1,4],[2,5],[3,6]]
(3,1) (3,1) (3,2)

vstack([1,2,3], [4,5,6]) = [[1,2,3],[4,5,6]]
(3,_) (3,_) 1×3 1×3 2×3

[[1,2,3]]

3 ML Concepts: Construction of Decision Trees

In this section, we will go over how to construct our decision tree learner on a high level. The following questions will help guide the discussion:

1. What exactly are the tasks we are tackling?
2. What are the inputs and outputs at training time? At testing time?
3. At each node of the tree, what do we need to store?
4. What do we need to do at training time?
5. What do we need to do at testing time?
6. What happens if max depth is 0? Majority vote
7. What happens if max depth is greater than the number of attributes? max-depth = 1,000,000, n=5, stop growing tree!

(edge cases: 6, 7)

Learn → Predict → Evaluate
TRAIN | TEST

1. learn DT classifier
   predict labels
   evaluate on a set of metrics

2. Training
   IN: training data
   OUTPUT: learned DT classifier

   Testing
   IN: - learned DT classifier
   - train/test data
   OUTPUT:
   - predicted labels
   - metrics (acc, error)

3. fields
   - attribute chosen
   - children pointers (L/R)

   arguments
   - subsetted data
   - depth
   - set of available features

4. Pseudocode for train
   B.C. "stopping criteria"
   - max depth
   - purity of labels
   - same feature vectors remaining
   Leaf → run majority vote
   Recursive
   Empty node q
   1. Calculate MI of all features; select best feature x_d
      $\forall v \in$ domain(x_d)
      $D_v = \{(x^{(n)}, y^{(n)}) \in D \mid x_d^{(n)} = v\}$ $D_{\{x_d=0\}}$ $D_{\{x_d=1\}}$
      Recursively call construct fn on D_v and store as children of q
      Remove chosen attribute from avail features

   x_1 ... x_d ... x_K y
   | 1 | 1 | 1 | ($D_{\{x_d=1\}}$)
   | 1 | 0 | 1 | 0 | ($D_{\{x_d=0\}}$)
   | 1 | 0 | 1 | 1 |
   | 1 | 1 | 1 | ($D_{\{x_d=1\}}$)

5. Testing
   - predict using trained DT on train/test data
   - write predictions & metrics to output file

4 ML Concepts: Mutual Information

Information Theory Definitions:

{1,1,1,1} ⇒ "pure"

{0,1,0,1} ⇒ "impure"

- entropy: $H(Y) = -\sum_{y \in values(Y)} P(Y = y) \log_2 P(Y = y)$
- cond. entropy: $H(Y \mid X = x) = -\sum_{y \in values(Y)} P(Y = y|X = x) \log_2 P(Y = y|X = x)$
- $H(Y \mid X) = \sum_{x \in values(X)} P(X = x) H(Y \mid X = x)$
- MI: $I(Y; X) = H(Y) - H(Y \mid X) = H(X) - H(X \mid Y)$

Exercises

1. Calculate the entropy of tossing a fair coin.

   $H(Y) = -P(Y=tails)\log_2 P(Y=tails) - P(Y=heads)\log_2 P(Y=heads)$
   $= -\frac{1}{2}\log_2\frac{1}{2} - \frac{1}{2}\log_2\frac{1}{2}$
   $= -\log_2 \frac{1}{2}$
   $= 1$

2. Calculate the entropy of tossing a coin that lands only on tails. *Note:* $0 \cdot \log_2(0) = 0$.

   $H(Y) = -1 \cdot \log_2 1 - 0 \log_2 0$
   $= 0$

3. Calculate the entropy of a fair dice roll.

   $-\sum_y P(Y=y)\log_2 P(Y=y) = -\sum_{i=1}^{6} \frac{1}{6}\log_2\frac{1}{6} = \log_2 6$

   max entropy ⇒ $\log_2 |5|$

4. When is the mutual information $I(Y; X) = 0$?

   $I(Y;X) = 0 \iff X \perp\!\!\!\perp Y$
   ↳ symbol for indep

Used in Decision Trees:

n=8

Outlook (X_1)	Temperature (X_2)	Humidity (X_3)	Play Tennis? (Y)
sunny	hot	high	no
overcast	hot	high	yes
rain	mild	high	yes
rain	cool	normal	yes
sunny	mild	high	no
sunny	mild	normal	yes
rain	mild	normal	yes
overcast	hot	normal	yes

1. Using the dataset above, calculate the mutual information for each feature (X_1, X_2, X_3) to determine the root node for a Decision Tree trained on the above data.

   - What is $I(Y; X_1)$? $H(Y) = -\frac{2}{8}\log_2\frac{2}{8} - \frac{6}{8}\log_2\frac{6}{8} \approx 0.811$; $H(Y|X_1) = \sum_{x_1} P(X_1=x_1) H(Y|X_1=x_1)$
   - What is $I(Y; X_2)$? $I(Y;X_1) \approx 0.811 - 0.344 \approx 0.467$
   - What is $I(Y; X_3)$? $I(Y;X_2) \approx 0.061$ $I(Y;X_3) \approx 0.311$
   - What feature should be split on at the root node?

     max. MI

     X_1 at root

   X_1=sunny: $H(Y|X_1=sunny) = -\frac{2}{3}\log_2\frac{2}{3} - \frac{1}{3}\log_2\frac{1}{3}$

   X_1=rainy: $H(Y|X_1=rainy) = \emptyset$

   X_1=overcast: $H(Y|X_1=overcast) = \emptyset$

   $= P(X_1=sunny) H(Y|X_1=sunny) + P(X_1=rainy) H(Y|X_1=rainy) + P(X_1=overcast) H(Y|X_1=overcast)$

   $= \frac{3}{8}\left[-\frac{2}{3}\log\frac{2}{3} - \frac{1}{3}\log_2\frac{1}{3}\right] + \frac{3}{8}\cdot\emptyset + \frac{2}{8}\cdot\emptyset$

   ≈ 0.344

2. Calculate what the next split should be.

   X_2 or X_3?

   $0.25 \approx I(Y;X_2)$
   $0.918 \approx I(Y;X_3)$ } feature maximizes MI

   X_1: sunny → ?, overcast → ○, rainy → ○

3. Draw the resulting tree.

   X_1: sunny → X_3 (high → no, normal → yes; X_2), overcast → yes, rainy → yes

X_1	X_2	y
0	1	0
0	0	0
1	1	1
1	0	1

$I(Y;X_2) = 0$

X_1	X_2	y
1	1	0
0	0	1
1	1	1
0	0	0

→ majority vote

X_1	X_2	y
0	0	0
0	0	0
0	0	0