Programming Assignment 3: Spam Classifier

Background and Structure

Seemingly, everyone is talking about Machine Learning and Artificial Intelligence these days. Artificial Intelligence (AI), a subfield of Computer Science, is concerned with enabling computers to perform tasks that require rational decision-making. As one of the oldest areas of research in the discipline, AI has played a significant role in driving technological advancements since the 1950s. On the other hand, machine Learning (ML) is a subfield of AI that uses trends from previous examples to make predictions about unseen data using statistical methods. ML algorithms are not magic — they simply guess the most likely outcome based on many, many previous examples. This means that **any ML algorithm's predictions are only as good as the data it was built upon,** which can easily be biased in some way, or just plain wrong.

As computer scientists, it is important to be able to recognize and advocate for appropriate uses of these models, regardless of how miraculous they may seem to the public.

Terminology

There are several machine learning terms used throughout the specification for this assignment that we would like to formally define before you begin. It might even be worth having this slide open in another tab while reading the assignment to make sure you fully understand the terms being given to you.
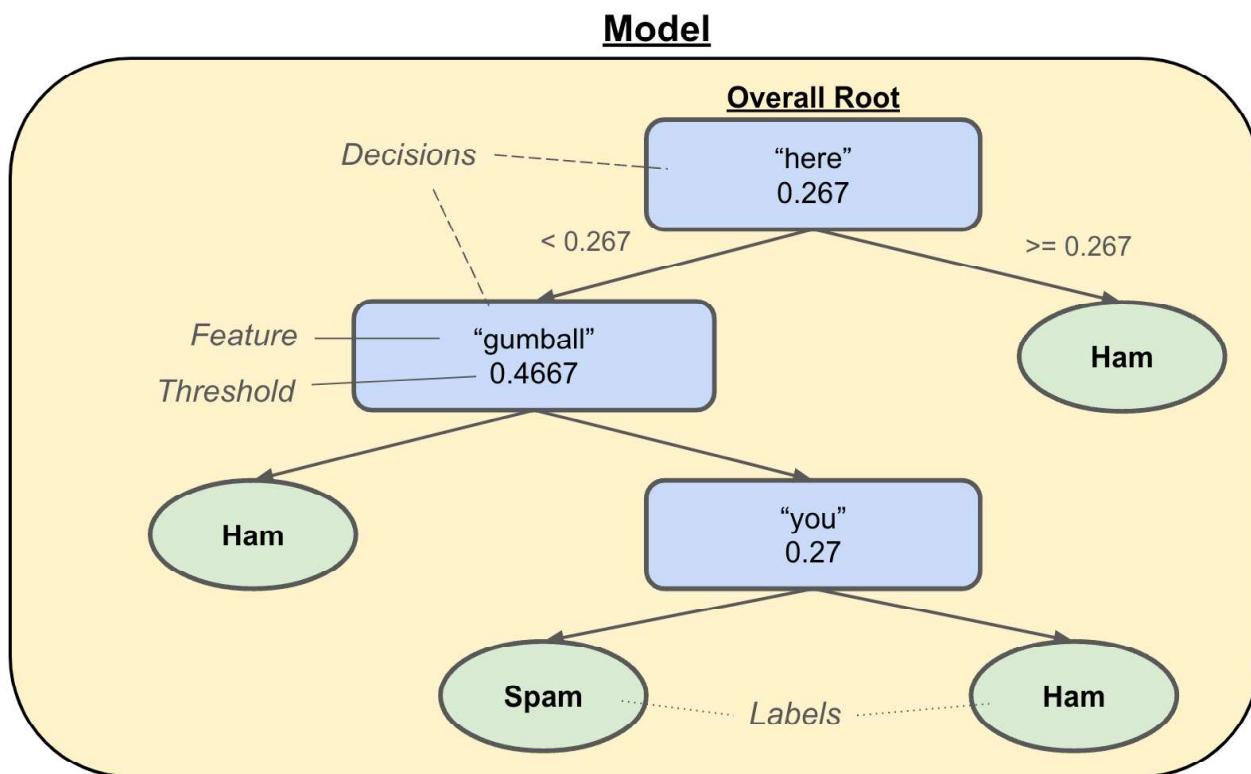
- **Model:** The actual program that makes probabilistic classifications on provided inputs.
- **Training:** Models are "trained" on previously gathered datasets to make future predictions.
- **Label**: How data is classified after being run through the model. In our tree, leaf nodes will house classification labels.
- **Feature**: Important aspects/characteristics of our dataset that we use in classification that correspond to a numeric value. Typically, the hardest part of a machine learning algorithm is determining how to take input data and "featurize" it into something a computer can understand
    - Ex: turning a sentence or image into a series of numbers.
- **Threshold**: The numeric value we compare a feature against at any branch node within our classifier. In our tree, if the current input is less than the threshold we should go left. If it's greater than or equal to, we should go right.

Structure

Your goal for this assignment is to implement a text-based classification tree, a simplistic machine-

learning model that predicts a label when given some text-based data. In this section, we'll familiarize you with the classifier's visual structure. Additionally, this assignment involves a lot of Machine Learning (ML) terminology. For clarity, these terms are underlined within this specification

Below is a visual example of what a classification tree might look like for classifying spam emails:

Model

As seen above, in our classification tree the **leaf nodes represent our predictive labels** ("Spam" or "Ham" – a funny way of writing not spam) while the **branch nodes represent decision nodes** that contains some feature of our data and a threshold to determine what decision to make. For this assignment, the feature will be the word probability of a certain word.

As mentioned earlier, you will be given text-based data to classify. This may include, but is not limited to, emails, academic papers, or even movie reviews! Throughout this assignment, we'll call these **text blocks** and represent them with the `TextBlock` class (more on that in the Specification slide).

To classify a given text block, you start at the root of the tree and determine whether the corresponding feature found in the input falls to the left or right of the current node's threshold (determined by `<` or `>=`). Then, you travel in the corresponding direction. Repeating this process will eventually lead you to a classification for your input.

Below, we'll trace through the classification of a sample input with our example model shown above.

▶ Expand

These classification trees may not always be the same, and may not always operate on identifying `"spam"` or `"ham"`. Below is an alternative example of what a potential classification tree could look

like for Pokemon Types based on text from Pokedex entries.

Legend

Split Label

Overall Root
"hot"
0.5

"swim"
0.89

Fire

"flower"
0.056

Water

Normal Grass

To solidify the different tree behaviors, we'll trace through an input much like the example above.

▶ Expand

This is what you'll be implementing in this assignment! Specifically, you'll be creating a classification tree that's able to predict a label given some text. This could range from predicting "Spam" or "Ham" given the contents of an email (as shown above) to predicting the author of a given Federalist Paper!

Training a Classification Tree

One of our goals is to be able to "train" our model from previously gathered data in order to make future predictions.

In the previous slide, we magically arrived at a constructed classification tree. In this section, we'll explain the algorithm to train a new model. Throughout this section, we'll be using the following file (which can be found in `spec_example.csv`):

```
Category,Message
Ham,here here here four five six seven eight nine ten eleven tweleve thirteen office you
Spam,one two three four five six seven eight nine ten eleven tweleve thirteen office you
Spam,one two three four five six seven eight nine ten eleven tweleve thirteen office you
Spam,here here here office office office office office office office office office office office of
```

Note the structure of this `.csv` file: the second column contains the data, and the first column contains the expected label for that piece of data. For ease of implementation, this file will automatically be parsed for you (using the provided `DataLoader` and `CsvReader` classes) and passed into your constructor as two lists.

```
public Classifier(List<TextBlock> data, List<String> labels)
```

Step 1: Initialize our Model

Since the classification tree is empty at the beginning, we need to add an initial data point so it can start making classifications. This algorithm processes training examples in order, starting from index `0`. So we begin by inserting the first data-label pair (at index 0) into the tree.

Additionally, we want to store the `TextBlock` data associated with the label in our leaves. The reason why will make sense later in our algorithm.

Expand to see the visualization:

▶ Expand

Step 2: Classify Data

Now that we have a classification tree, we can start to classify inputs! Unfortunately, with only one point of data, our model doesn't seem very useful — currently, it classifies every input as `Ham`. What if we try to classify a piece of data that has an expected label of `Spam`?

To handle this, we proceed to the next step of the algorithm: we process the next index. We'll **start**

at the top of the tree and traverse down to find the label our model will predict for
`data.get(index)`. Now, we check whether our model's prediction matches the expected label.

- If the prediction is correct, then our model is accurate up to that point, and we have nothing to do!
- If the prediction **is incorrect**, we need to **update the model** — this is the "learning" part. We modify the tree so that it can correctly classify this new example in the future.

Expand to see visualization:

▶ Expand

Step 2a: Updating our Model

Typically, a large part of the complexity in building a classification tree is determining how to partition the data in case of incorrect predictions. There are many potential ways to accomplish this, but we'll be taking this approach:

- Call the `findBiggestDifference` method on the current `TextBlock` input and the previously stored one (from the misclassified label node).
  - `findBiggestDifference` identifies and returns the feature with the **largest difference in word probabilities** between the two `TextBlock`s.
- We then compute the **midpoint** between the two feature values in the `TextBlock`s and use it as the threshold for a new decision node.
- The decision node should be placed where the original label node currently is.
- After the new decision node is constructed, the original label node and current input should be placed appropriately.

> ℹ️ Notice that we are only ever modifying the leaves of our tree!

Expand to see visualization:

▶ Expand

> ℹ️ **SIDE NOTE:** This algorithm requires you to keep track of both the label and the `TextBlock` datapoint first assigned to this label within every leaf node created in this constructor, as without the previous `TextBlock` datapoint we would be unable to create a new decision node! Ideally we'd like to keep track of all input data that falls under a specific leaf node such that when creating a new decision node, we can make sure it's valid for our entire training dataset. For simplicity, only worry about the first datapoint used to create a label node.

Step 3: Repeat

Repeat step 2 for the rest of the list until we've finished processing the list. At that point, our model is fully trained on our data and is able to predict the right label for the data we inputted.

Expand to see visualization:

▶ Expand

Learning Objectives

By completing this assignment, students will demonstrate their ability to:

- Define data structures to represent compound and complex data
- Write a functionally correct Java class to represent a binary tree.
- Write classes that are readable and maintainable, and that conform to provided guidelines for style, implementation, and performance.
- Produce clear and effective documentation to improve comprehension and maintainability of programs, methods, and classes.

System Structure

Below we describe the **provided** `TextBlock` class that will be used in your implementation of `Classifier.java` . You do not need to (and should not) make changes to this class, but your code will be a client of it. Make sure to understand the purpose of this class and read through the provided documentation.

TextBlock.java

Text data that gets classified via the <u>classifier</u>. It defines four public methods:

```
public double get(String word)
```

- Returns the corresponding word probability for the given word.
    - Although there are classification trees where it would make sense to return something else (imagine a color feature within a real estate dataset), since our implementation is only dealing with thresholds for word probabilities, this must return a double.

```
public Set<String> getFeatures()
```

- Returns a set of all features from this text block.

```
public boolean containsFeature(String word)
```

- Returns true if this dataset contains the given word. False otherwise.

```
public String findBiggestDifference(TextBlock other)
```

- Returns the word with the biggest difference in probability between this `TextBlock` and the

other `TextBlock`.
- Note that there is no difference between calling `a.findBiggestDifference(b)` and `b.findBiggestDifference(a)`. Both will return the same string.

Required Class

> ⚠️ **NOTE:** To earn a grade higher than N on the Behavior and Concepts dimensions of this assignment, **your *core* algorithms for each method must be implemented** *recursively*. **You will want to utilize the** *public-private pair* **technique discussed in class.** You are free to create any helper methods you like, but the core of your implementations must be recursive.

Classifier.java

In this assignment, you implement your classification tree by creating a class called `Classifier.java`. You are provided with a client program that handles user interaction and calls your `Classifier` methods in order to train/load a model and classify data.

```
public Classifier(Scanner input)
```

- Load the classifier from a file connected to the given `Scanner`. The format of the input file should match that of the `save` method (described below).
  - Importantly, in this method, you should only read data from the file using `nextLine` and convert it to the appropriate format using `Double.parseDouble`.
- This method should throw an `IllegalArgumentException` if the provided `input` is null and an `IllegalStateException` if the tree is still empty after processing `input`.

```
public Classifier(List<TextBlock> data, List<String> labels)
```

- Create and train a classifier from the input data and corresponding labels.
- The lists should be processed in parallel in increasing order (i.e. process index 0, then 1, then 2, etc...), where the label corresponding to `data.get(<index>)` can be found at `labels.get(<index>)`. The general construction process should be accomplished via the algorithm described in <u>Training a Classification Tree</u> slide.
- This method should throw an `IllegalArgumentException` if any of the following cases are met:
  - `data` or `labels` is null
  - `data` and `labels` are not the same size
  - `data` or `labels` is empty

> ✅ **HINT:** This algorithm requires you to keep track of the initial `TextBlock` used to create the label node. Without this initial `TextBlock`, we would be unable create a new feature for when our model is inaccurate! Keeping this in mind, what may be one of the fields needed in the `ClassifierNode` class?

```
public String classify(TextBlock input)
```

- Given a piece of data, return the appropriate label that this classifier predicts.
  - This method should model the steps taken in our Email example above: at every feature evaluate our input data and determine if it's less than our threshold. If so, continue left; otherwise, continue right. Repeat this process until a leaf node is reached.
- If the parameter `input` is null, throw an `IllegalArgumentException`.

```
public void save(PrintStream output);
```

- Saves this current classifier to the given `PrintStream`
  - For our classifier tree, **this format should be pre-order**. Every branch node will print two lines of data, one for feature preceded by "Feature: " and one for threshold preceded by "Threshold: ". For leaf nodes, you should only print the label. **Examples of the format can be seen below and through the `trees` directory in the start code.**
- If the parameter `output` is null, throw an `IllegalArgumentException`

Provided Methods

Additionally, we have provided two other methods in `Classifier.java`:

```
public Map<String, Double> calculateAccuracy(List<TextBlock> data, List<String> labels)
```

- Returns the model's accuracy on all labels in a provided testing dataset. This is useful to see how well our model works, and what labels it is struggling with classifying correctly.

```
private static double midpoint(double one, double two)
```

- Helper method to calculate the midpoint between two doubles.
  - **HINT:** This should be used in the `Classifier(List<TextBlock>, List<String>)` constructor to calculate the midpoint!

`ClassifierNode`

As part of writing your `Classifier` class, you should also have a **private static inner class** called `ClassifierNode` to represent the nodes of the tree. The contents of this class are up to you, but must meet the following requirements:

- You must have a single `ClassifierNode` class that can represent both features and labels — you should *not* create separate classes for the different types of nodes.
  - You may find that since we are representing both features and labels in the same node class, some fields may be unused at times. This is completely okay!
- The `ClassifierNode` class must not contain any constructors or methods that are not used by the `Classifier` class.
- The fields of the `ClassifierNode` class must be `public`.

- All data fields should be declared `final` as well. This does not include fields representing the children of a node.

> **i** **NOTE:** You may get a `variable ___ might not have been initialized` error, in which case, you will have to explicitly initialize the values for *all* `final` fields in your node class — even if you do not plan to utilize the value.

File Format

The files that are both created by the `save` method and read by the `Scanner` constructor will follow the same format. These files will contain a pair of lines to represent branch nodes and a single line to represent leaf nodes in the `Classifier`. The first line in each branch node pair will start with "Feature: " followed by the feature and the second line will start with "Threshold: " followed by the threshold. Lines representing the leaf nodes will simply contain the label. The format of the file should be a **pre-order traversal** of the tree.

For example, consider the following sample file (`simple.txt`):

> ▶ Expand

Try out your Classifier!

Once those methods are implemented, you'll have a working classifier! Try it out using `Client.java` and see how well it does (what is its accuracy on our test data). Also, try saving your tree to a file and see what it looks like. Is it creating decisions on features you'd expect? Why or why not? (Note that this is a big area of current CS research called "explainable AI" — how can we interpret the results from these massive probability models that are often difficult for humans to understand).

Client Program & Visualization

We have provided you with a `Client` program to help test your implementation of the methods within `Classifier.java`. The client can create binary trees from the provided `.csv` or `.txt` files and test their accuracy. Note that in order to pass in these files, you should call them by `folderName/fileName`. For example, `trees/simple.txt` .

Click "Expand" below to see sample executions of the client for different situations (user input is **bold and underlined**).

1. This client visualization uses your Scanner constructor, `calculateAccuracy()`, and `classify()`. The constructor loads a pre-trained model from a given text file. The following inputs allow us to test its accuracy using the pre-set `TEST_FILE` (in this example, we initialized it in line 8 of `Client.java` to `"data/emails/test.csv"`) and use the model to predict labels for data points in a given `.csv` file.

> **ⓘ** **NOTE:** When saving the `Scanner` constructor, the contents of the file will be exactly the same as the input `.txt` file used to initially load the pre-trained model.

▶ Expand

2. This client visualization uses the `Classifier` constructor that takes in data and their corresponding labels, `calculateAccuracy()` (implemented for you), and `save()`. You can follow the pattern of inputs below to train the classification model using some `train.csv` file (this calls the constructor `Classifier(List<TextBlock>, List<String>)`, retrieve testing accuracy (similar to above), and save the trained model to a file so that it is in `.txt` format (like the sample input files in the trees/ folder). `TRAIN_FILE` and `TEST_FILE` were set to `"data/federalist_papers/train.csv"` and `"data/federalist_papers/test.csv"` respectively.

▶ Expand

> **ⓘ** **NOTE:** After quitting, the saved file should be available for viewing in the console.

Testing

There are no formal testing requirements for this assignment. However, we'd encourage you to get your hands dirty and see how well your model performs on the provided dataset / investigate the output files to see if you can make sense of what the inner structure is!

Implementation Guidelines

Your program should exactly reproduce the format and general behavior demonstrated in the Ed tests. Our recommended approach is as follows:

1) Design your Node

First, design your node class that represents both the branch and leaf nodes within your classification tree. Think about the information these nodes will be required to store based on the specification. Remember that in our classification tree, branch nodes represent decisions and leaf nodes represent classification labels.

> **i** **NOTE:** You may find that since we are representing both features and labels in the same node class, some fields may be unused at times. This is completely okay!

2) Scanner Constructor

This constructor will be given a `Scanner` that contains data produced by `save()`. In other words, the input for this constructor is the output you produced with `save()`.

Remember that this file is stored in pre-order format, where the feature and threshold for decision nodes are stored on two lines within the file:

```
Feature: here
Threshold: 0.125
```

And labels are present without any additional formatting:

```
ham
```

You may assume that "Feature" and "Threshold" will never be labels within the input file.

Remember that you should only ever call `.nextLine()` on the provided Scanner. You might be tempted to call `nextLine()` to read the feature then `next()` and `nextDouble()` to read the threshold, but remember that mixing token-based reading and line-based reading is not so simple. Assuming you are trying to retrieve the value of the threshold, here is an alternative that uses a method called `parseDouble` in the `Double` class that allows you to use `nextLine()`:

```
double threshold = Double.parseDouble(input.nextLine().substring("Threshold: ".length()));
```

Lastly, you should throw an `IllegalArgumentException` is `input` is null and an `IllegalStateException` if the tree is still empty after processing `input`.

> ✅ **HINT:** It looks like we're processing lines and using that information to *modify* our tree. Keeping in mind our recently learned concept, **what pattern should we employ to help implement this constructor?**

> ⚠️ The tests for your `Scanner` constructor implementation are tied to a working `save` implementation. This means that once you feel comfortable with your solution you should move onto the next part, and test for both implementations at the same time.

Relevant Problem:

- Section 14: Write Tree

3) save()

Once you've implemented the `Scanner` constructor, do the opposite! Namely, given an already constructed classification tree, save it to the provided `PrintStream` via a **pre-order traversal** (described in the Specification slide):

> These files will contain a pair of lines to represent branch nodes and a single line to represent leaf nodes in the `Classifier`. The first line in each branch node pair will start with "Feature: " followed by the feature and the second line will start with "Threshold: " followed by the threshold. Lines representing the leaf nodes will simply contain the label.

You should also throw an `IllegalArgumentException` is `output` is null.

> ⚠️ At this point, test your `Scanner` constructor and `save` implementations. We don't recommend moving forward in this assignment until these two methods are passing the provided tests.

Below, we've provided sample client input and output that should be your expected output at this point (user input is **bold and underlined**):

```
▶ Expand
```

Relevant Problems:

- Lesson 13: `printPreOrder`
- Section 13: Read Tree

4) classify()

Now we can start classifying! This method should traverse through the tree by evaluating decision nodes on the input data to see whether or not the input falls below the current threshold. If so, the traversal should continue into the left subtree, otherwise the right. Once a leaf node is reached the corresponding label should be returned.

For a feature at a given decision node, think about how we could retrieve its word probability from the input data.

Finally, you should throw an `IllegalArgumentException` is `input` is null

> ⚠️ At this point, test your current implementation.We don't recommend moving forward runtil the `classify` method is passing

> ℹ️ **NOTE:** The `classify` tests are another way for us to test that your `Scanner` constructor implementation is correct (since testing using `save` alone doesn't guarantee its correctness). If you find your `classify` tests failing, but believe your `classify` implementation is correct, try taking a look at the logic inside your `Scanner` constructor!

Below, we've provided sample client input and output that should be your expected output at this point (user input is **bold and underlined**):

▶ Expand

5) Two List Constructor

Finally, here is where we actually "train" our model, and will likely be the most difficult part of your implementation. First, you should make sure to throw the proper exceptions:

- `IllegalArgumentException` if any of the following cases are met:
  - `data` or `labels` is null
  - `data` and `labels` are not the same size
  - `data` or `labels` is empty

Next, your implementation should follow the following algorithmic approach (copied from the Training a Classification Tree — diagrams depicted there):

Step 1: Initialize our Model

Since the classification tree is empty at the beginning, we need to add an initial data point so it can start making classifications. This algorithm processes training examples in order, starting from index `0`. So we begin by inserting the first data-label pair (at index 0) into the tree.

Additionally, we want to store the `TextBlock` data associated with the label in our leaves. The reason why will make sense later in our algorithm.

Step 2: Classify Data

Now that we have a classification tree, we can start to classify inputs! Unfortunately, with only one point of data, our model doesn't seem very useful — currently, it classifies every input as `Ham`. What if we try to classify a piece of data that has an expected label of `Spam`?

To handle this, we proceed to the next step of the algorithm: we process the next index. We'll **start at the top of the tree** and traverse down to find the label our model will predict for `data.get(index)`. Now, we check whether our model's prediction matches the expected label.

- If the prediction is correct, then our model is accurate up to that point, and we have nothing to do!
- If the prediction **is incorrect**, we need to **update the model** — this is the "learning" part. We modify the tree so that it can correctly classify this new example in the future.

Step 2a: Updating our Model

Typically, a large part of the complexity in building a classification tree is determining how to partition the data in case of incorrect predictions. There are many potential ways to accomplish this, but we'll be taking this approach:

- Call the `findBiggestDifference` method on the current `TextBlock` input and the previously stored one (from the misclassified label node).
  - `findBiggestDifference` identifies and returns the feature with the **largest difference in word probabilities** between the two `TextBlock`s.
- We then compute the **midpoint** between the two feature values in the `TextBlock`s and use it as the threshold for a new decision node.
- The decision node should be placed where the original label node currently is.
- After the new decision node is constructed, the original label node and current input should be placed appropriately.

> **i** **SIDE NOTE:** This algorithm requires you to keep track of both the label and the `TextBlock` datapoint first assigned to this label within every leaf node created in this constructor, as without the previous `TextBlock` datapoint we would be unable to create a new decision node! Ideally we'd like to keep track of all input data that falls under a specific leaf node such that when creating a new decision node, we can make sure it's valid for our entire training dataset. For simplicity, only worry about the first datapoint used to create a label node.

Step 3: Repeat

Repeat step 2 for the rest of the list until we've finished processing the list. At that point, our model is fully trained on our data and is able to predict the right label for the data we inputted.

> ✅ **HINT:** This algorithm requires you to keep track of the initial `TextBlock` used to create the label node. Without this initial `TextBlock`, we would be unable create a new feature for when our model is inaccurate! Keeping this in mind, what may be one of the fields needed in the `ClassifierNode` class?

> ✅ **HINT:** It looks like we're processing data and using that information to *modify* our tree. Keeping in mind our recently learned concept, **what pattern should we employ to help implement this constructor?**

> ⚠️ At this point, test your current implementation. Once these tests are passing, the assignment should be completed. CONGRATULATIONS!!! Make sure your code adheres to the Code Quality Guide and

Below, we've provided sample client input and output that should be your expected output at this point (user input is **bold and underlined**). Note that `TRAIN_FILE` and `TEST_FILE` were set to `"data/federalist_papers/train.csv"` and `"data/federalist_papers/test.csv"` respectively:

▶ Expand

Relevant Problems:

- Section 14: Remove Leaves in List
- Section 14: Make Full

Code Quality

⚠️ **NOTE:** To earn a grade higher than N on the Behavior and Concepts dimensions of this assignment, **your *core* algorithms for each method in** `Classifier` **must be implemented** *recursively*. **You will want to utilize the *public-private pair* technique discussed in class.** You are free to create any helper methods you like, but the core of your implementations must be recursive.

As always, your code should follow all guidelines in the Code Quality Guide and Commenting Guide. In particular, pay attention to these requirements:

- **Constructors in inner class:**
  - Any constructors created should be used.
  - When applicable, reduce redundancy by using the `this()` keyword to call another constructor in the same class.
  - Clients of the class should never have to manually set fields of an object immediately after construction (when possible) — there should be a constructor included for this situation.
    - For example, if you were the implementor of the `Point` class:

      ```
      Point coord = new Point(); // Poor usage of constructor
      coord.x = 5; // ✗
      coord.y = 7; // ✗

      Point coord = new Point(5, 7); // ▫ Correct usage of constructor
      ```

- **Methods:**
  - All methods present in `Classifier` that are not listed in the specification must be `private`.
  - Make sure that all parameters within a method are used and necessary.
  - Avoid unnecessary returns.
- **`x = change(x)`:**
  - Similar to with linked lists, do not "morph" a node by directly modifying fields (especially

when replacing an branch node with a leaf node or vice versa). Existing nodes can be rearranged in the tree, but adding a new value should always be done by creating and inserting a new node, not by modifying an existing one.

- An important concept introduced in lecture was called `x = change(x)`. This idea is related to the proper design of recursive methods that manipulate the structure of a binary tree. **You should follow this pattern where necessary when modifying your trees**.

- **Avoid redundancy:**
  - If you find that multiple methods in your class do similar things, you should create helper method(s) to capture the common code. As long as all extra methods you create are private (so outside code cannot call them), you can have additional methods in your class beyond those specified here.
  - Look out for including additional base or recursive cases when writing recursive code. While multiple calls may be necessary, you should avoid having more cases than you need. Try to see if there are any redundant checks that can be combined!

- **Data Fields:**
  - Properly encapsulate your objects by making data fields in your `Classifier` class private. (Fields in your `ClassifierNode` class should be public following the pattern from class.)
  - Avoid unnecessary fields; use fields to store important data of your objects but not to store temporary values only used in one place.
  - Fields should always be initialized inside a constructor or method, never at declaration.

- **Commenting**
  - Each method should have a comment including all necessary information as described in the Commenting Guide. Comments should be written in your own words (i.e. not copied and pasted from this spec).
  - Make sure to avoid including *implementation details* in your comments. In particular, for your object class, a *client* should be able to understand how to use your object effectively by only reading your class and method comments, but your comments should maintain *abstraction* by avoiding implementation details.
  - Continuing with the previous point, keep in mind that the client should **not** be aware of what implementation strategy your class/methods utilize.

Spam Classifier

📄 P3_SpamClassifier.zip

Reflection

For this week's reflection, we would like everyone to watch and engage with the following video.

Step 1: Watch How computers learn to recognize objects instantly | Joseph Redmon (7m 37s)

An error occurred.

Try watching this video on www.youtube.com, or enable JavaScript if it is disabled in your browser.

Question 1

Step 2: Write your own comment below, responding to one of the following prompts in its entirety:

Prompt 1: Misclassification in AI systems can have serious consequences. For example, a self-driving car misclassifying a pedestrian as an object could be dangerous. What are some other real-world cases where classification errors could have ethical or safety implications? What sorts of issues could arise in these situations? Please provide at least two cases.

Prompt 2: The classifier in the video can identify objects with high accuracy but can also make incorrect assumptions about an image. How does this compare to errors a human might make in classification? Which types of errors do you think are worse? Why? Please list at least one similarity and one difference and your opinions on each.

Prompt 3: Consider a situation where a human would struggle to classify an image, but a computer would succeed. What about the reverse? Please provide at least one example of each situation. What might cause these differences? What does this tell you about when machine learning algorithms should or shouldn't be used?

Please label which prompt you are answering. For full credit, all questions within a prompt must be clearly answered. In particular, we suggest using the ACE (answer, cite, and explain) format. For a meaningful response, it may be helpful to pose some counterexamples, connect in terms of your own experience, add additional support from course materials, and be as specific as possible in your own reasoning

No response

Question 2

Describe how you went about testing your implementation. What specific situations and/or test cases did you consider? Why were those cases important?

No response

Question 3

What skills did you learn and/or practice with working on this assignment?

No response

Question 4

What did you struggle with most on this assignment?

No response

Question 5

What questions do you still have about the concepts and skills you used in this assignment?

No response

Question 6

About how long (in hours) did you spend on this assignment? (Feel free to estimate, but try to be close.)

No response

Question 7

Was any part of the specification or requirements unclear? If so, which part(s), how was it unclear, and how could it have been made more clear?

No response

Question 8

[OPTIONAL] Do you have any other feedback, questions, or comments about this assignment?

(Note that we may not be able to respond to questions here, so please post on the message board if you would like a response!)

No response

Final Submission

Fill out the box below and click "Submit" in the upper-right corner of the window to submit your work.

Question

I attest that the work I am about to submit is my own and was completed according to the course Academic Honesty and Collaboration policy. If I collaborated with any other students or utilized any outside resources, they are allowed and have been properly cited. If I have any concerns about this policy, I will reach out to the course staff to discuss *before* submitting.

(Type "yes" as your response.)

No response