D&D Combat Length Predictions

Erdős Institute
Data Science Boot Camp

Guiding Question

- Given information about the starting state of a combat encounter, how many rounds should the encounter take?

Guiding Question

- Given information about the starting state of a combat encounter, how many rounds should the encounter take?

- Useful for game masters crafting encounters for their players.

Data Preparation

- FIREBALL dataset: ~25,000 files, each depicting a unique combat encounter.
  - Scraped from AVRAE, a Discord bot which helps players run online D&D sessions.

Data Preparation

- FIREBALL dataset: ~25,000 files, each depicting a unique combat encounter.
  - Scraped from AVRAE, a Discord bot which helps players run online D&D sessions.

- Format: Nested dictionaries, containing combat state updates, player state updates, commands issued, etc.

Data Preparation

- FIREBALL dataset: ~25,000 files, each depicting a unique combat encounter.
  - Scraped from AVRAE, a Discord bot which helps players run online D&D sessions.

- Format: Nested dictionaries, containing combat state updates, player state updates, commands issued, etc.

- Features extracted:
  - Party size
  - Player average level
  - Monster party size
  - Monster average level

Data Visualization

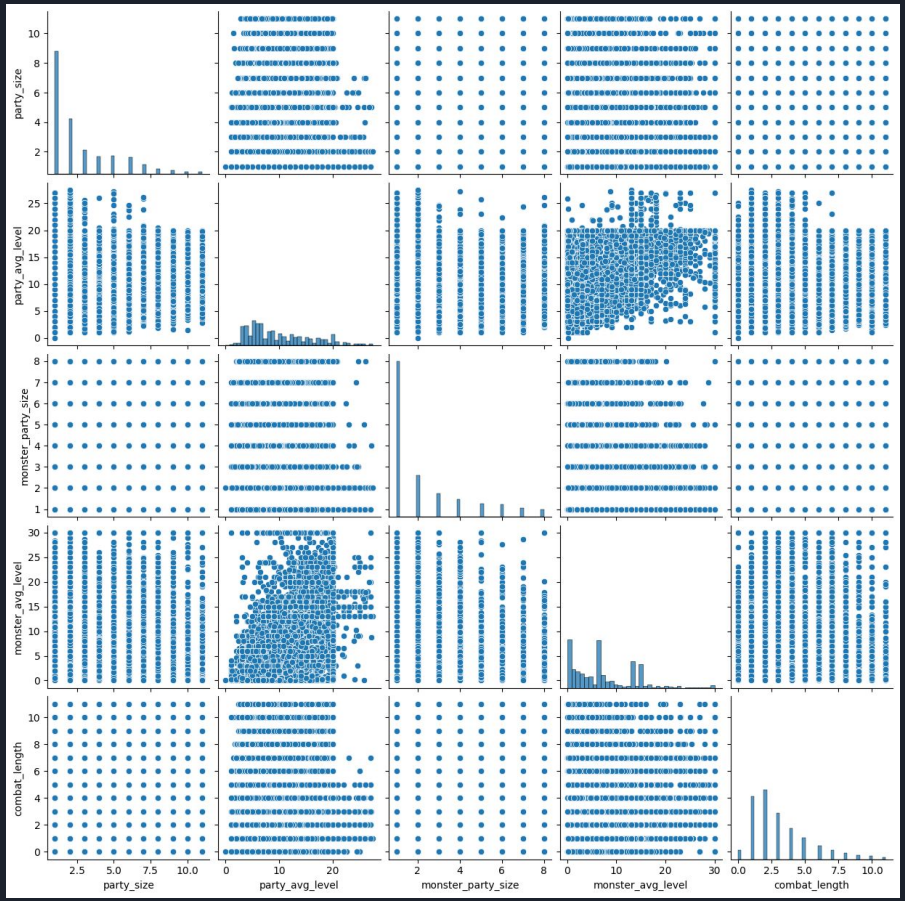

Data Visualization

Purpose

- Visualization of how features we use in the data is distributed

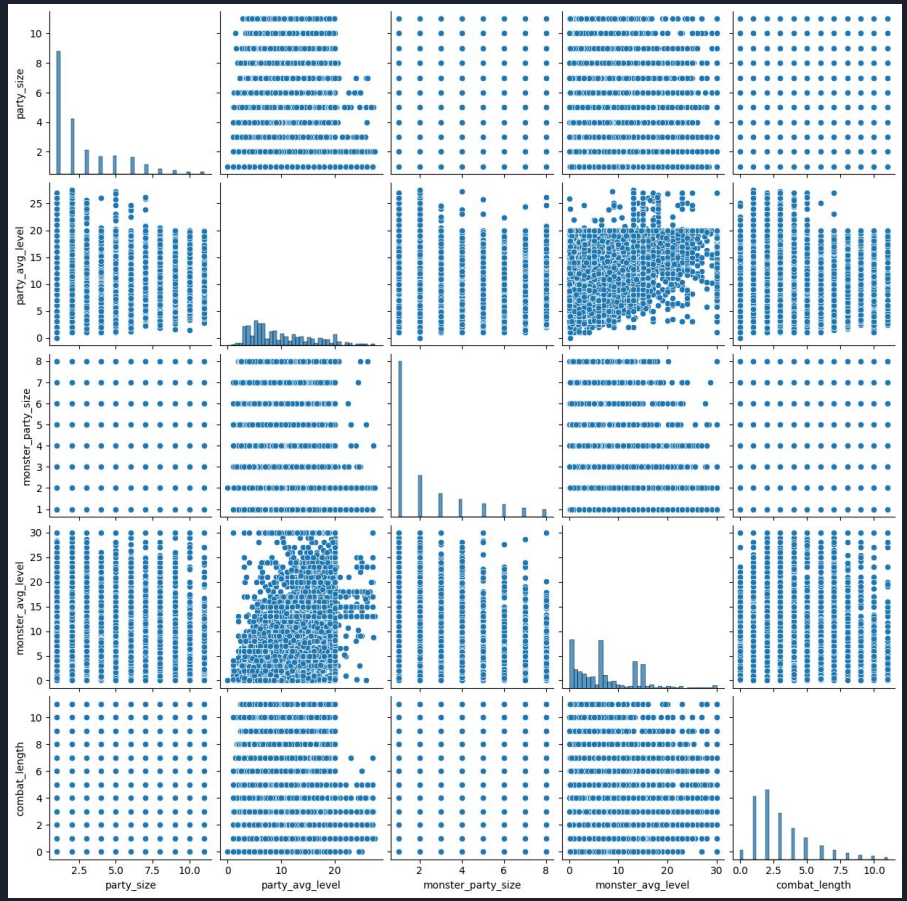

Data Visualization

Purpose

- Visualization of how features we use in the data is distributed
- Visualization of correlations between the features

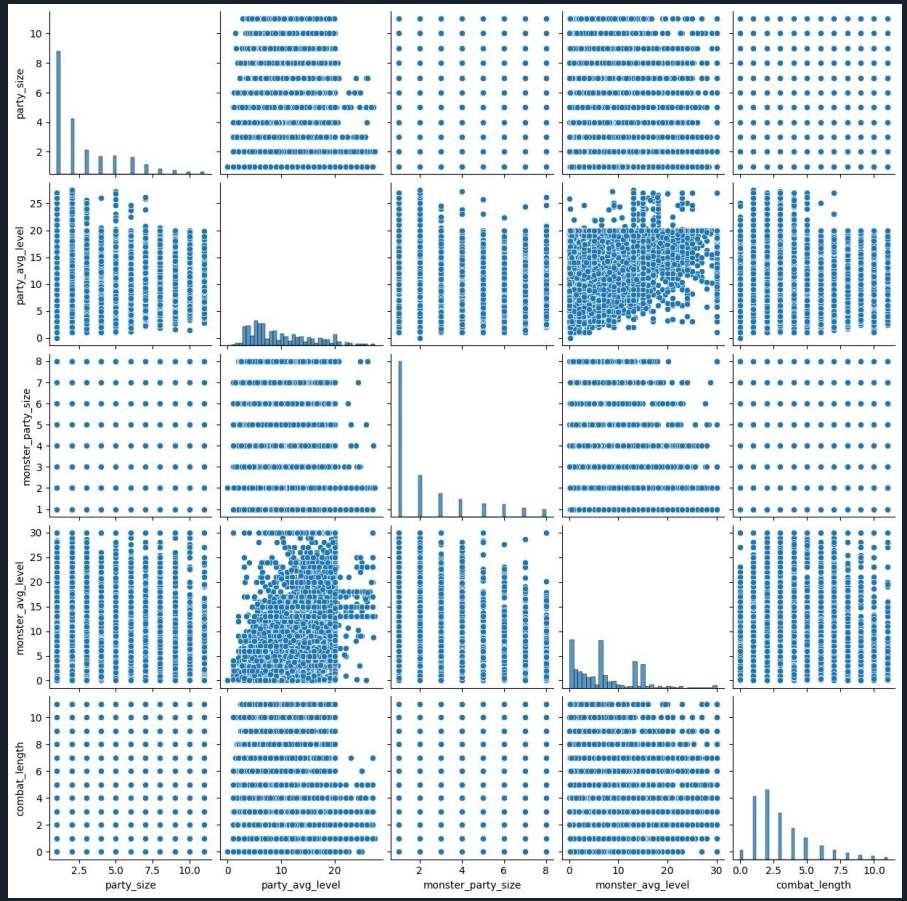

Data Visualization

Purpose

- Visualization of how features we use in the data is distributed
- Visualization of correlations between the features

Results

- Diagonal plots show data distribution

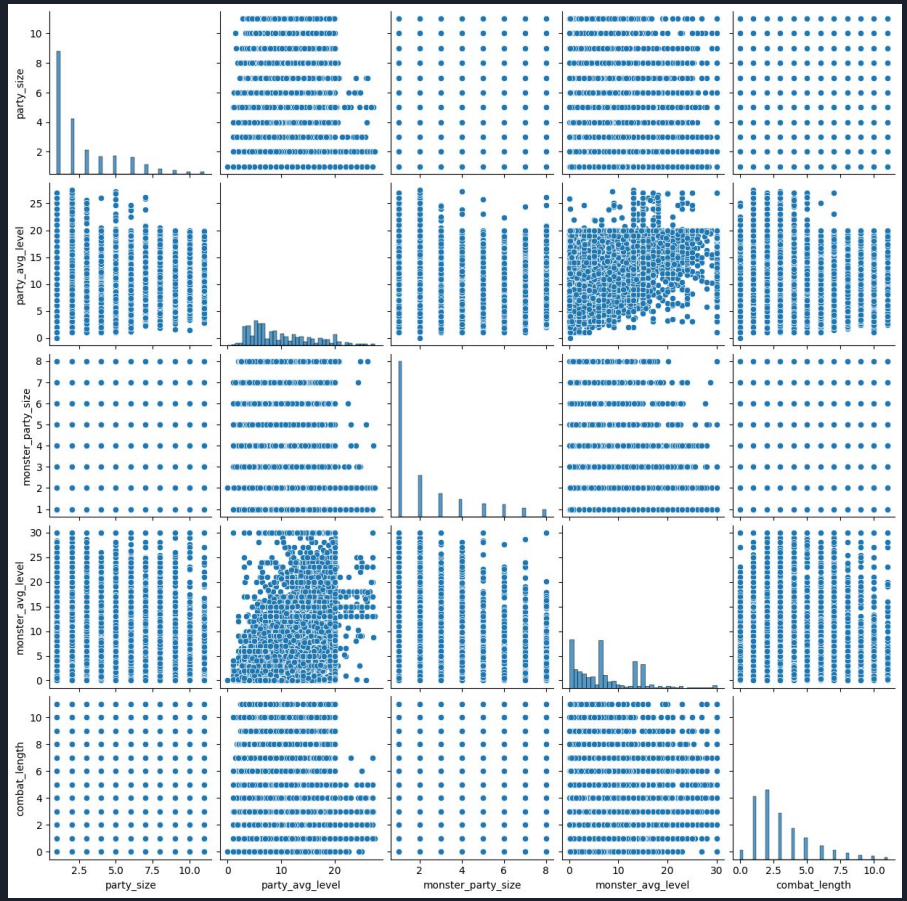

Data Visualization

Purpose

- Visualization of how features we use in the data is distributed
- Visualization of correlations between the features

Results

- Diagonal plots show data distribution
- Off Diagonal plots show that there are not too many correlations

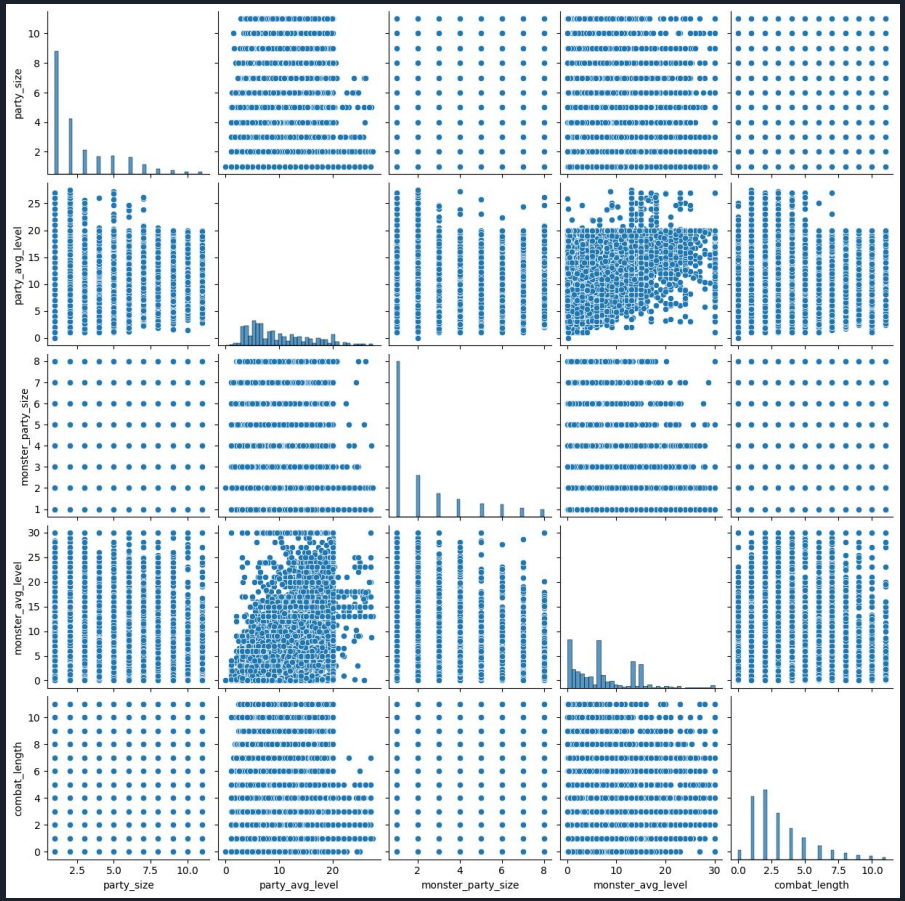

Baseline Models - Mode

Baseline Model	Cross-Validation MSE
• Mode Baseline	5.941

Baseline Models - Median

Baseline Model	Cross-Validation MSE
Mode Baseline	5.941
Median Baseline	5.766

Baseline Models - Mean

Baseline Model	Cross-Validation MSE
Mode Baseline	5.941
Median Baseline	5.766
Mean Baseline	4.822

Initial Advanced Model - Linear Regression

Model	Cross-Validation MSE
Mode Baseline	5.941
Median Baseline	5.766
Mean Baseline	4.822
Linear Regression	4.059

More Advanced Models

- Ensembles of Trees!
  - Random Forest
  - Gradient Boost
  - XGBoost
- Grid Search for Hyperparameter Tuning

Ensembles of Trees

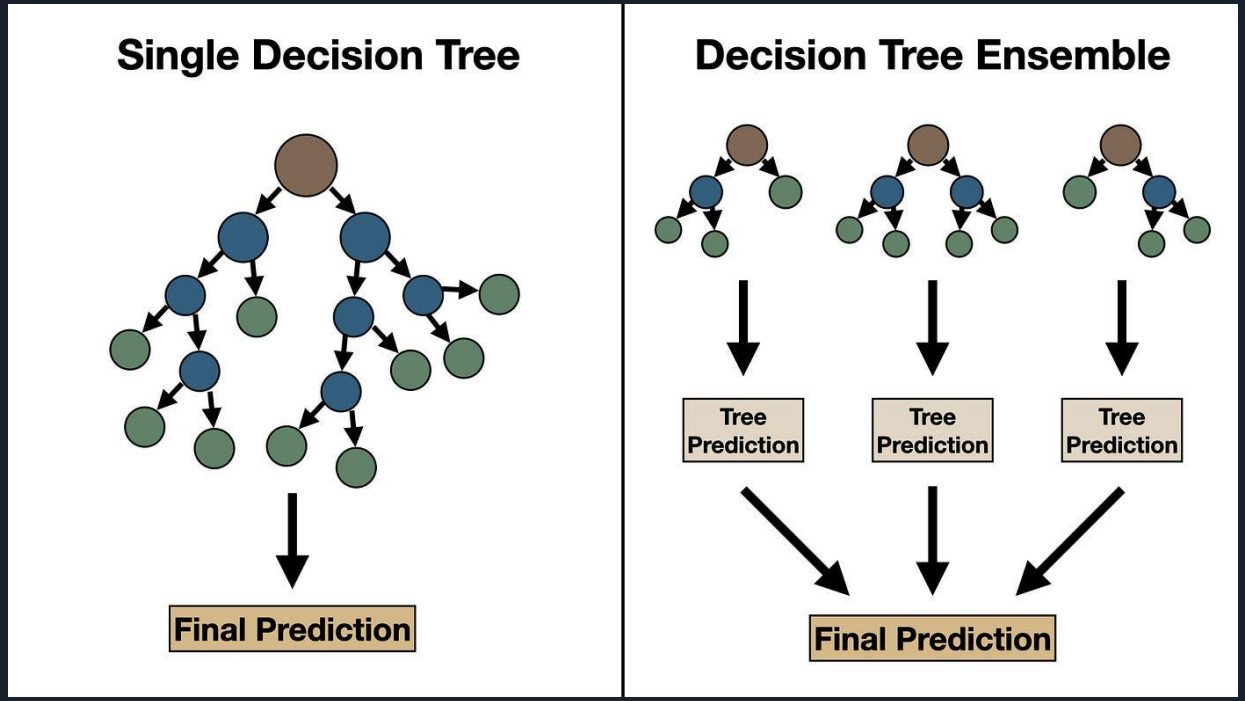

Grid Search for Hyperparameters

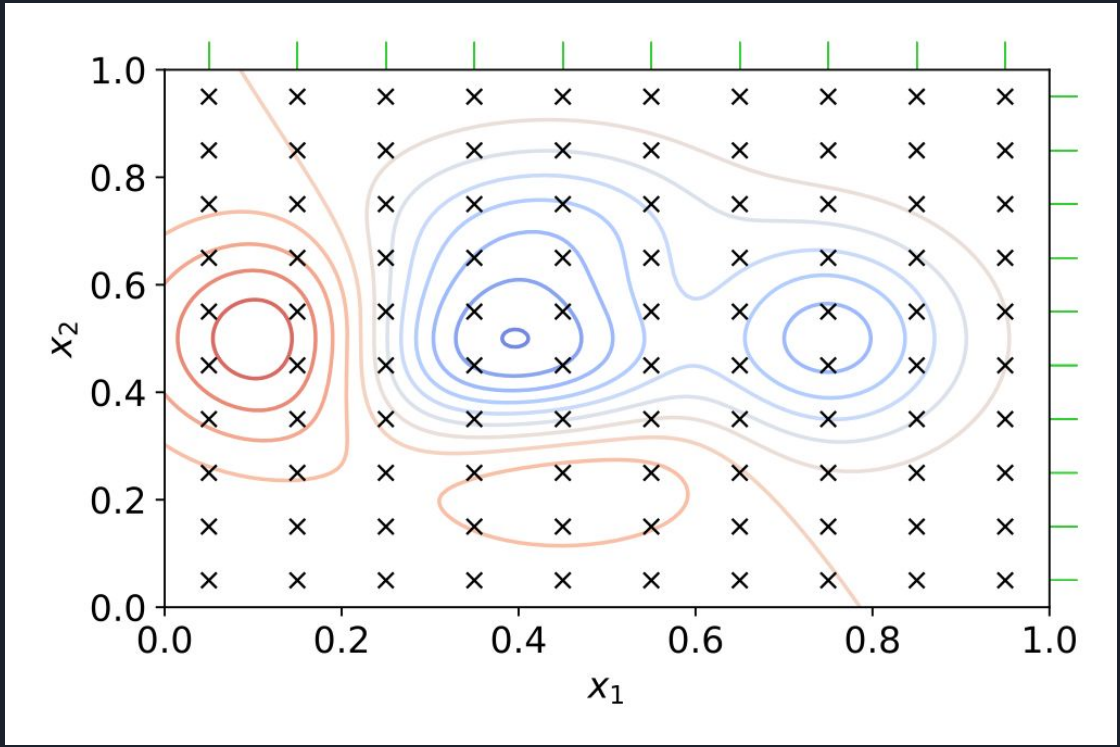

Goal

- Have systematic approach for finding the best hyperparameters for each ensemble learning model

Method

- Use Grid Search for each model

Random Forest Regressor

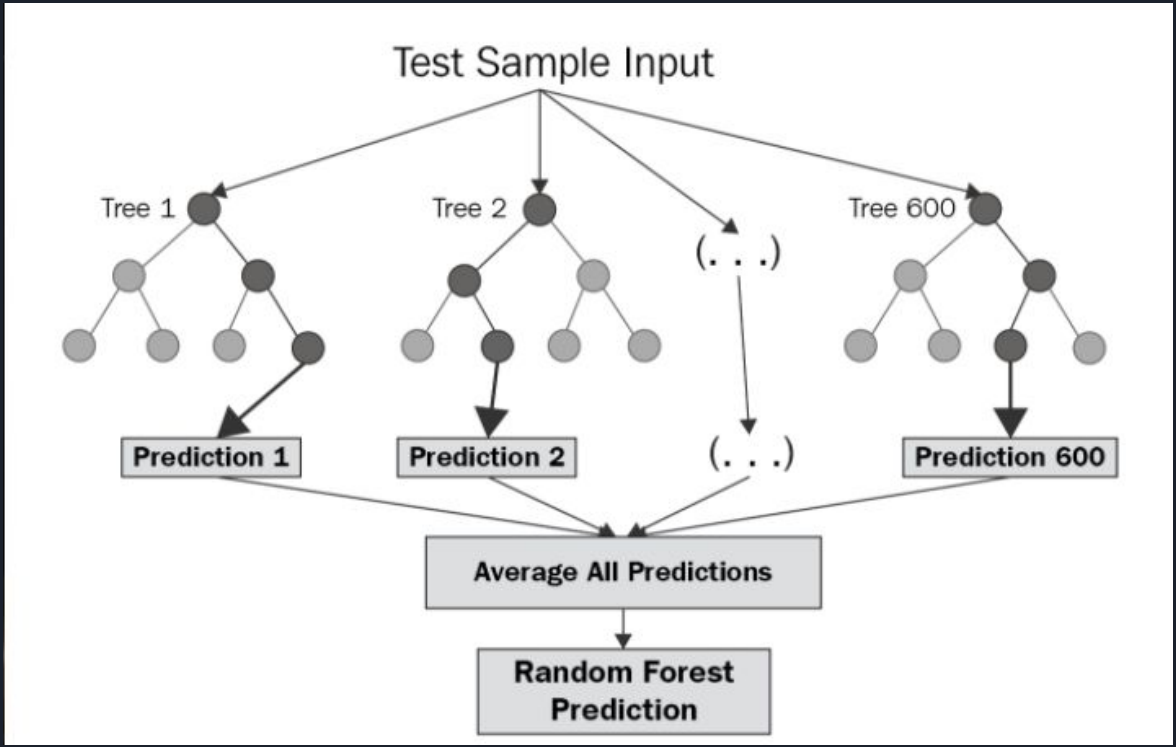

Histogram-based Gradient Boost

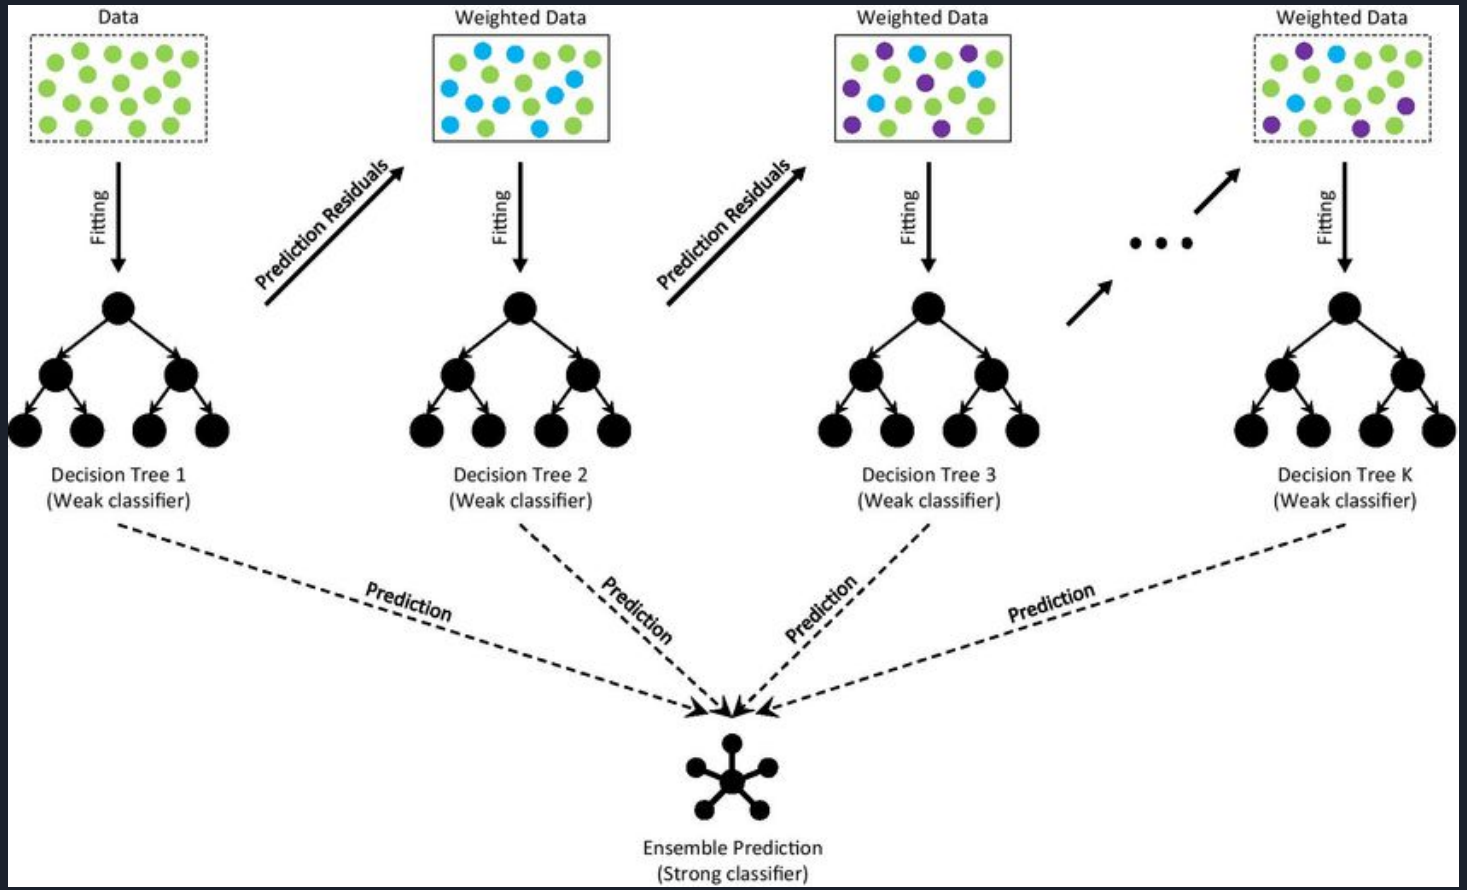

XGBoost

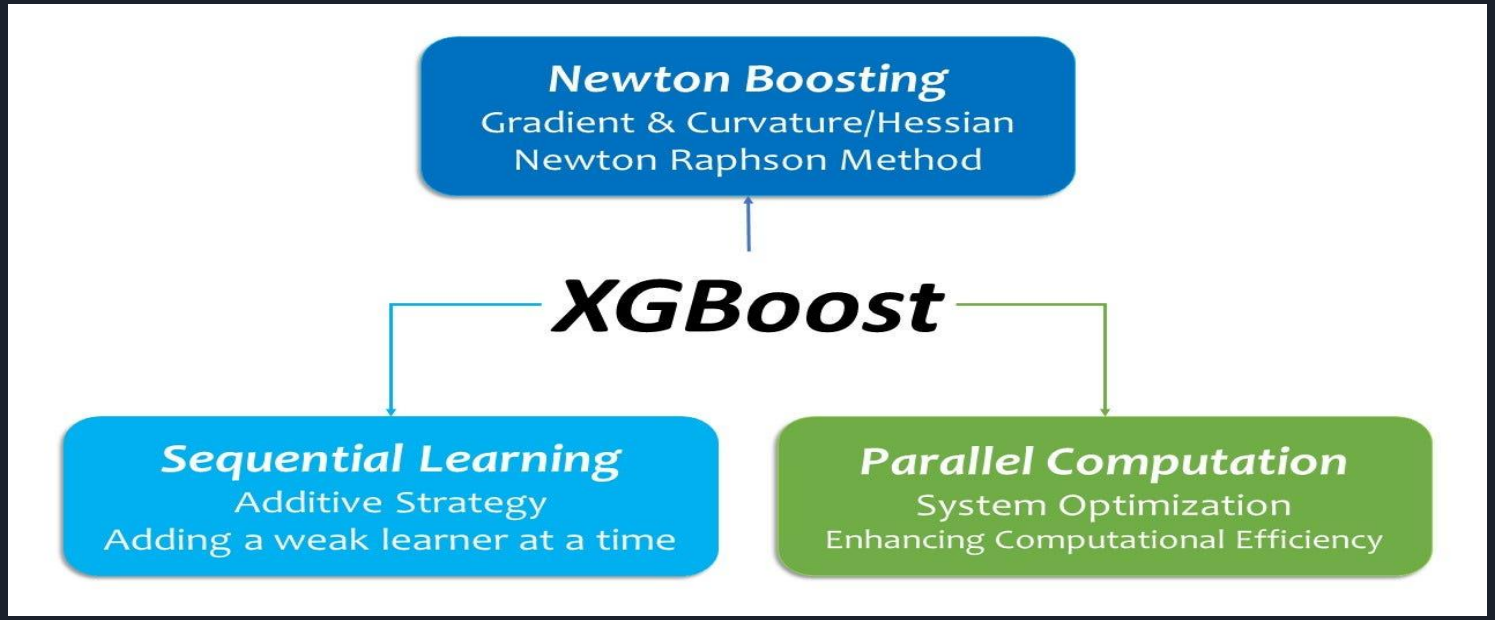
Newton Boosting
Gradient & Curvature/Hessian
Newton Raphson Method
XGBoost
Sequential Learning
Additive Strategy
Adding a weak learner at a time
Parallel Computation
System Optimization
Enhancing Computational Efficiency

Comparing the Models

Model Name	Average Cross-Validation MSE
Mode Baseline Model	5.941
Median Baseline Model	5.766
Mean Baseline Model	4.822

Comparing the Models

Model Name	**Average Cross-Validation MSE**
Mode Baseline Model	5.941
Median Baseline Model	5.766
Mean Baseline Model	4.822
Linear Regression Model	4.059
Tuned Random Forest Model	3.880
Tuned Gradient Boost Model	3.797
Tuned XGBoost Model	***3.788***

Our Final Model - XGBoost + Feature Importance

Fit our tuned XGBoost to whole training set:

Tuned XGBoost	Final MSE = 3.625

Our Final Model - XGBoost + Feature Importance

Fit our tuned XGBoost to whole training set:

Tuned XGBoost	Final MSE = 3.625

Computed Feature Importance on whole set:

Feature	Importance Score
Party Size	0.107
Average Party Level	0.093
Monster Party Size	**0.688**
Average Monster Level	0.112

A Web App

A Web App

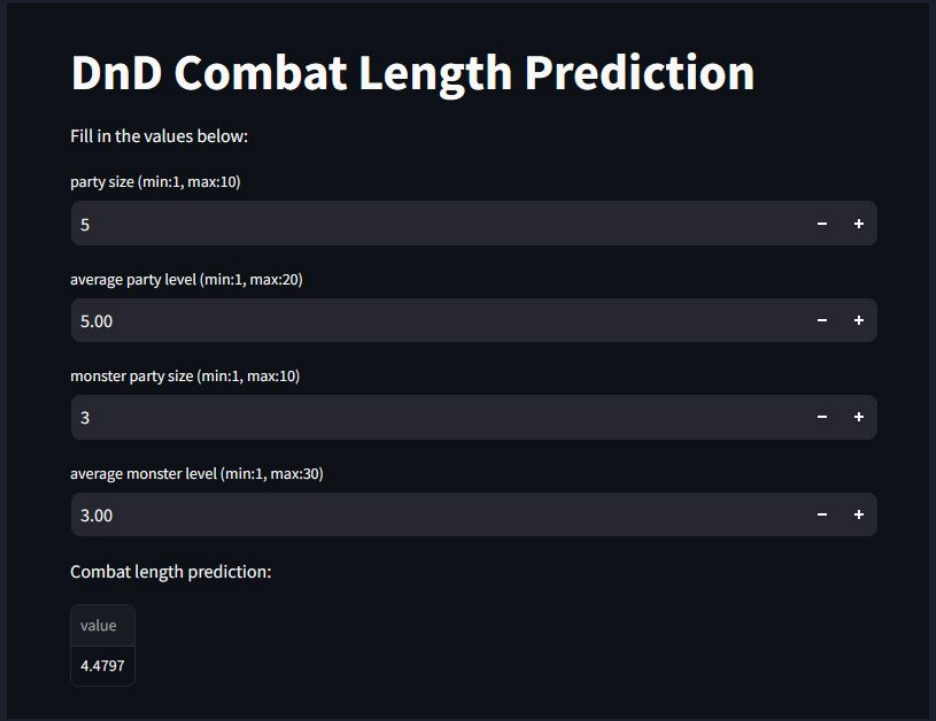
DnD Combat Length Prediction
Fill in the values below:
party size (min:1, max:10)
5
average party level (min:1, max:20)
5.00
monster party size (min:1, max:10)
3
average monster level (min:1, max:30)
3.00
Combat length prediction:
value
4.4797

On the App

Some drawbacks of the app/model:

On the App

Some drawbacks of the app/model:

- The demo shows that the model is suitable for more reasonable data points. That is, a reasonable party size will give a reasonable prediction.

On the App

Some drawbacks of the app/model:

- The demo shows that the model is suitable for more reasonable data points. That is, a reasonable party size will give a reasonable prediction.
- Entering an “unreasonable” party size gives unexpected outputs.

Future Directions

Future Directions

- To further improve the MSE of our final model, the most impactful change we could make is engineering more features.

Future Directions

- To further improve the MSE of our final model, the most impactful change we could make is engineering more features.
- Four features is relatively small, and the model did quite well considering this.

Future Directions

- To further improve the MSE of our final model, the most impactful change we could make is engineering more features.
- Four features is relatively small, and the model did quite well considering this.
- Some possible features to consider which could be extracted from the FIREBALL dataset are various categorical features about the player party, e.g. , capturing the party composition, such as how many spellcasters did the party have.