

How to Easily Understand the Difference Between Decision Trees and Random Forests

Authored by
stats writer

December 4, 2025

RECOMMENDED CITATION

stats writer (2025). *How to Easily Understand the Difference Between Decision Trees and Random Forests*. PSYCHOLOGICAL SCALES. Retrieved from <https://scales.arabpsychology.com/?p=104828>

In the realm of machine learning, both the decision tree and the random forest stand out as foundational algorithms for classification and regression tasks. While a decision tree is a straightforward predictive model that maps observations about an item to conclusions about its target value using a tree-like structure, it operates within the constraints of supervised learning and often necessitates careful tuning to determine optimal split points, making it susceptible to high variance.

Conversely, the random forest represents a significant advancement, leveraging the power of ensemble methods. This technique aggregates the results of numerous individual decision trees, utilizing a powerful statistical approach known as bagging (Bootstrap Aggregating). This automated approach inherently minimizes the need for extensive manual intervention, and crucially, it is designed to be significantly more robust and less susceptible to the critical problem of overfitting, thereby yielding superior generalization performance on unseen data.


The Core Mechanics of a Decision Tree

A **decision tree** is a highly intuitive machine learning model, particularly effective when addressing complex, **non-linear relationships** between a set of independent predictor variables and a target response variable. The fundamental objective is to recursively partition the data space into smaller, more manageable subsets, achieving this segmentation through a series of logical, rule-based decisions. Each split is designed to maximize the homogeneity of the resulting subsets concerning the response variable.

The conceptual architecture of the model mirrors a flowchart structure, featuring internal nodes that represent tests on attributes, branches that represent the outcome of the test, and leaf nodes that represent the final classification or predicted value. This process essentially constructs a predictive model, or "tree," where decisions based on predictor variables translate directly into an estimated value or category for the response variable. This inherent clarity makes decision trees highly appealing for preliminary analysis and situations demanding high transparency.

Consider a practical scenario, such as predicting the annual salary of professional baseball players. We might utilize predictor variables like "years played" and "average home runs." The decision tree algorithm systematically evaluates these features to establish optimal thresholds—the split points—that best separate high earners from low earners. The resulting visual model clearly outlines the specific pathway of decisions required to arrive at a prediction for any given player.

Using the structure derived from a training dataset, the resulting decision tree model visualizes the prediction process as follows:



Interpreting the Decision Tree Structure

The visualization provided by the tree diagram offers immediate insight into the underlying decision logic. By following the pathways from the root node down to the leaf nodes, we can determine the predicted salary based on the player's attributes:

Players classified as having less than 4.5 years played fall into the lowest tier, resulting in a predicted salary of **\$225.8k**. This illustrates the importance of experience as a primary determinant. For players exceeding the experience threshold (greater than or equal to 4.5 years played) but demonstrating lower performance (less than 16.5 average home runs), the model predicts a moderate salary of **\$577.6k**.

The highest predicted salary, **\$975.6k**, is reserved for players who meet both the experience requirement (greater than or equal to 4.5 years played) and demonstrate elite performance (greater than or equal to 16.5 average home runs).

A significant benefit of the decision tree is its unmatched **interpretability**. The model can be fitted rapidly to a dataset, and its outcomes are easily visualized and explained using the clear "tree" diagram. This transparency is invaluable in fields where regulatory requirements or ethical considerations mandate understanding how a decision was reached.

However, this inherent simplicity comes with a major vulnerability: the propensity for **high variance**. A single decision tree is highly susceptible to fitting the noise within a training dataset, a phenomenon known as overfitting. Consequently, the model is likely to perform poorly when

applied to previously unseen data, compromising its predictive utility in real-world applications. Furthermore, the splitting structure of the tree can be heavily influenced by anomalies or **outliers** present in the dataset, leading to unstable decision boundaries.

Introducing Random Forests and Ensemble Learning

To mitigate the inherent instability and overfitting issues associated with a single decision tree, researchers developed an advanced methodology: the random forest. This sophisticated technique is a powerful embodiment of ensemble learning, which operates on the principle that combining the predictions of multiple individual models yields a more accurate and stable overall prediction than any single constituent model could achieve alone. A random forest, therefore, is fundamentally a large collection of de-correlated decision trees.

The stability gained by the random forest stems from two key randomization techniques employed during its construction. The first is **Bootstrap Aggregating**, or bagging, which introduces randomness into the data selection. The second randomization step involves feature subset selection, where only a fraction of the total predictor variables is considered at each split point within each tree. These dual processes ensure that the individual trees are diverse and their errors are not systematically correlated, leading to a robust final model when their predictions are aggregated.

The process of constructing a reliable random forest model is systematic and highly automated, following these precise steps:

Data Sampling via Bootstrapping: Multiple, independent bootstrapped samples (sampling with replacement) are generated from the original training dataset. Each sample is used to train one individual tree.

Feature Subsetting for Tree Building: For every bootstrapped sample, a full decision tree is built. Crucially, when determining the best split at any node, only a random subset of the total predictor variables is considered, ensuring the resulting trees are distinct and de-correlated.

Prediction Aggregation: To derive the final predictive model outcome, the predictions from all individual trees are averaged (for regression tasks) or determined by a majority vote (for classification tasks).

Advantages and Disadvantages of Random Forests

The primary benefit of utilizing random forests lies in their superior performance regarding generalization. Because the technique aggregates many diverse trees and utilizes bagging, they are inherently much more robust against the high variance that plagues single decision trees. This robustness translates directly into higher accuracy and reliability when making predictions on unseen datasets. Furthermore, by averaging the results of numerous models, the impact of

isolated data points or **outliers** is significantly diluted, making the random forest model comparatively stable.

However, the complexity that grants random forests their power also introduces notable drawbacks. The most immediate concern is the loss of **interpretability**. Unlike a single decision tree, which can be neatly visualized, a random forest involves potentially hundreds or thousands of individual trees, making it impossible to produce a simple, clear visual representation of the final decision process. The model essentially becomes a "black box," where input goes in and a prediction comes out, without a transparent path to the decision.

Another major constraint revolves around **computational resources**. Training an ensemble model composed of hundreds of trees requires substantial processing power and memory, particularly when dealing with exceptionally large datasets. Consequently, while random forests deliver high accuracy, they demand significantly more time and computational ability to build compared to fitting a single, simple decision tree.

Comparative Analysis: Interpretability vs. Accuracy

The essential trade-off between the two models centers on the fundamental conflict between **model interpretability** and **predictive accuracy**. The following visual summary highlights the key operational differences:

	Decision Tree	Random Forest
Interpretability	Easy to interpret	Hard to interpret
Accuracy	Accuracy can vary	Highly accurate
Overfitting	Likely to overfit data	Unlikely to overfit data
Outliers	Can be highly affected by outliers	Robust against outliers
Computation	Quick to build	Slow to build (computationally intensive)

We can elaborate on the critical distinctions outlined in the comparison table:

Interpretability and Transparency

Decision trees excel in transparency. The logical, conditional path from input features to output prediction is explicitly mapped, allowing practitioners to easily generate a comprehensive tree diagram to visualize and fully comprehend the model's internal decision-making mechanism. This

inherent clarity is often necessary in high-stakes fields like finance or medicine, where explaining the rationale behind a prediction is paramount.

In sharp contrast, the random forest, due to its complex ensemble nature, sacrifices interpretability for power. Since the final prediction is a result of aggregating hundreds of independent trees, visualizing the entire process is infeasible. This often renders the random forest a complex "black box," making it challenging to precisely understand the weight and interaction of specific predictor variables in the final prediction.

Stability: Overfitting and Outlier Resistance

When evaluating model stability, the handling of overfitting becomes the paramount concern. Individual decision tree models inherently seek to achieve perfect classification or prediction on the training data, often by creating overly complex boundaries that capture statistical noise instead of the genuine underlying data patterns. This high specialization leads to poor generalization, meaning decision trees typically yield significantly lower accuracy when deployed on novel, unseen datasets.

The random forest was specifically engineered to overcome this instability. By combining two forms of randomization--data sampling (bagging) and feature selection--the individual decision trees within the forest are intentionally diverse and decorrelated. When the predictions from these weak, yet diverse, learners are averaged, the overall predictive model exhibits drastically reduced variance and is far less prone to overfitting, resulting in substantially higher and more reliable accuracy on testing data.

Furthermore, the resilience of the random forest extends to managing **outliers**. A single decision tree is highly susceptible to the influence of extreme data points, which can drastically alter its split criteria. However, because a random forest model builds many individual decision trees and then takes the average of those trees' predictions, the effect of any single outlier is minimized or cancelled out in the final averaging process. This statistical robustness makes random forests a more durable choice for datasets that may contain anomalies or noise.

Computational Requirements

The computational footprint of the two algorithms presents another key differentiation point. Since a decision tree involves partitioning the dataset sequentially only once, it is inherently fast and efficient to train. The computational expense is relatively low, making it an excellent choice for rapid prototyping, initial data exploration, or environments with severely limited processing power.

Conversely, the random forest demands significantly greater **computational resources**. The process requires generating multiple bootstrapped samples and subsequently training hundreds or

even thousands of individual decision trees in parallel or sequence. This intensive training process translates into a much longer model build time and a greater demand on memory, especially when the input dataset is voluminous. While modern computing infrastructure often handles this efficiently, resource constraints remain a practical limitation compared to the speed of a single tree model.

Practical Application: When to Choose Each Model

The choice between a decision tree and a random forest should be dictated by the specific priorities of the project--namely, whether **interpretability** or **accuracy** is the overriding requirement. As a general rule of thumb, one should opt for a **decision tree** if the primary need is for a fast, simple, non-linear predictive model where the ability to easily explain and visualize the decision rules is non-negotiable. This is often the case in regulated industries or for exploratory data analysis where understanding feature interaction is more valuable than achieving marginal gains in prediction accuracy.

Conversely, the **random forest** is the preferred algorithm when prediction performance is the critical metric. If ample computational resources are available, and the goal is to develop a highly robust model that minimizes variance and resists overfitting, the random forest is the superior choice. This approach sacrifices model transparency for guaranteed high accuracy, making it standard practice in competitive machine learning environments and industrial applications where predictive power trumps explanatory power.

In contemporary practice, machine learning engineers and data scientists predominantly rely on random forests. Advances in modern hardware and distributed computing systems have largely neutralized the historical computational constraints, allowing practitioners to leverage the algorithm's superior accuracy and inherent resistance to instability, even when dealing with massive, complex datasets.

Further Exploration and Resources

For those seeking to delve deeper into the theoretical underpinnings and practical implementation of these algorithms, the following resources provide comprehensive introductions to both decision trees and random forest models:

Introduction to Algorithm Theory and Implementation

In-Depth Comparison of Tree-Based Methods

Additionally, specialized tutorials explaining the mechanics of fitting decision trees and random forests using the R programming language are available for practical application:

Fitting Decision Trees in R

Implementing Random Forests in R

ARABPSYCHOLOGY.COM