

Learning to Evaluate Classification Models: A Step-by-Step Guide to Creating Precision-Recall Curves in Python

Authored by
Mohammed loot

November 2, 2025

RECOMMENDED CITATION

Mohammed loot (2025). *Learning to Evaluate Classification Models: A Step-by-Step Guide to Creating Precision-Recall Curves in Python*. PSYCHOLOGICAL STATISTICS. Retrieved from <https://statistics.arabpsychology.com/?p=8671>

Understanding Classification Model Evaluation

When developing [machine learning](#) models, particularly those focused on binary classification problems, moving beyond simple accuracy is essential for true performance assessment. Two indispensable metrics used to rigorously evaluate the quality and robustness of a classifier are [precision](#) and [recall](#). These statistics offer critical insight into how effectively the model distinguishes between positive and negative classes, which is especially important in high-stakes scenarios characterized by significant data imbalance or asymmetrical costs associated with errors.

In many classification tasks, the primary goal is to correctly identify positive instances while simultaneously minimizing the rate of incorrect predictions. The strategic selection of appropriate evaluation metrics is paramount because a model exhibiting high overall accuracy can still be practically useless if it fails to capture a significant proportion of actual positive cases (resulting in a high volume of false negatives) or if it frequently mislabels negative cases as positive (leading to high false positives). Therefore, relying solely on accuracy often masks critical deficiencies in a model's operational performance.

The decision regarding whether to prioritize [precision](#) or [recall](#) must be directly aligned with the specific business or scientific objective. Consider the field of medical diagnostics: maximizing [recall](#) (ensuring all sick patients are identified, thus minimizing missed diagnoses) is typically the overriding concern, even if it results in slightly lower precision (more false alarms). Conversely, for systems like automated spam filtering, high [precision](#) (only flagging emails that are truly spam) is crucial to prevent the loss of legitimate, important communications, accepting that some spam might slip through.

Defining Core Evaluation Metrics: Precision and Recall

The foundational concepts of precision and recall are mathematically derived from the four possible outcomes cataloged within a [confusion matrix](#): True Positives (TP), False Positives (FP), True Negatives (TN), and False Negatives (FN). A thorough understanding of these components--which represent correct and incorrect predictions relative to the actual class labels--is fundamental for accurately interpreting the subsequent visualization provided by the precision-recall curve.

Precision serves as a measure of the predictive model's exactness. Specifically, it quantifies the proportion of instances identified as positive by the model that were genuinely correct. It answers the fundamental question: "Out of all the data points the model enthusiastically labeled as positive, how many were truly positive cases?" A high precision score is a direct indicator of a low rate of False Positives, meaning the model rarely raises unwarranted alarms.

This critical metric is formally calculated using the following ratio:

Precision = True Positives / (True Positives + False Positives)

Recall, also frequently referred to as Sensitivity or the True Positive Rate, gauges the model's completeness. It measures the proportion of all actual positive cases in the dataset that the model successfully identified. The central question recall addresses is: "Among all the instances that were truly positive, what fraction did the model correctly capture?" Achieving high recall signifies a low rate of False Negatives, demonstrating the model's effectiveness at avoiding missed opportunities or critical omissions.

This metric is calculated using the following formula:


Recall = True Positives / (True Positives + False Negatives)

Visualizing Performance: The Precision-Recall Curve

To fully comprehend the dynamic relationship and the inherent tradeoff between precision and recall across the entire spectrum of classification thresholds, data scientists utilize the powerful graphical tool known as the [precision-recall curve](#). This visualization is particularly invaluable and offers much clearer insight compared to ROC curves when evaluating performance on severely imbalanced datasets, where other aggregate metrics might present a deceptively optimistic view of the model's capabilities.

The precision-recall curve is constructed by plotting the calculated precision values on the vertical (Y) axis against the corresponding recall values on the horizontal (X) axis. This plotting is performed for every possible probability threshold that the model could employ for binary classification. Conceptually, a curve that sweeps closer to the top-right corner of the plot signifies superior model performance, as it indicates that the classifier can maintain high precision even as recall capability significantly increases.

This comprehensive visualization empowers analysts to select the most appropriate operational decision threshold based on specific application requirements. For instance, if the deployment demands higher precision (minimizing false alarms), a threshold corresponding to a point higher up on the Y-axis would be chosen. Conversely, if minimizing missed positive cases is critical, the user would favor a threshold that maximizes recall (a point further right on the X-axis). The following practical example illustrates how to generate this curve using the industry-standard libraries [Scikit-learn](#) and [Matplotlib](#) within the Python environment.



Step 1: Setting up the Environment and Importing Libraries

The first foundational step in any reproducible data science workflow is ensuring that the necessary software libraries are correctly installed and imported into the Python script or notebook environment. For this specific task--generating the precision-recall curve--we rely heavily on the powerful functionalities offered by the Scikit-learn library for data manipulation, model training, and metric calculation, complemented by Matplotlib for the generation of high-quality visualizations.

This initial setup requires the importation of several specialized modules. We need modules for generating synthetic datasets (`datasets`), splitting the data into manageable training and testing subsets (`train_test_split`), the chosen classification algorithm itself (`LogisticRegression`), and, most critically, the specific Scikit-learn function dedicated to calculating the curve points (`precision_recall_curve`). Properly executing these import statements establishes the required programmatic foundation.

The code snippet provided below details the precise import structure necessary to begin the analysis:

```
from sklearn import datasets
from sklearn.model_selection import train_test_split
```

```
from sklearn.linear_model import LogisticRegression
from sklearn.metrics import precision_recall_curve
import matplotlib.pyplot as plt
```

Step 2: Preparing Data and Training the Model

Once the required packages are successfully imported, the next critical phase involves preparing a suitable dataset and training a preliminary predictive model. For the purposes of this demonstration, we leverage Scikit-learn's utility functions, specifically `make_classification`, to efficiently generate 1000 synthetic samples. This dataset includes defined features and controlled complexity, simulating a realistic, albeit simplified, binary classification challenge that is ideal for pedagogical purposes.

We have selected [Logistic Regression](#) as the classification algorithm. This choice is deliberate because Logistic Regression not only performs classification but also outputs calibrated probability scores, which are absolutely necessary for calculating threshold-dependent metrics like the precision-recall curve. A critical methodological step is splitting the generated data into dedicated training and testing sets. This ensures that the model's performance is assessed realistically on data it has never encountered, thereby rigorously preventing the problem of overfitting.

Following the model fitting process on the training data, we employ the powerful `predict_proba` method on the held-out test set. This method returns the probability distribution across all classes. We specifically extract the probability scores corresponding to the positive class (typically indexed as 1). These continuous probability scores form the essential input data upon which the entire precision-recall curve calculation is based, allowing us to systematically test every potential decision threshold.

```
# Create dataset with 4 informative features, simulating a binary classification problem.
```

```
X, y = datasets.make_classification(n_samples=1000,  
n_features=4,  
n_informative=3,  
n_redundant=1,  
random_state=0)
```

```
# Split dataset into 70% training and 30% testing subsets for unbiased evaluation.
```

```
X_train, X_test, y_train, y_test = train_test_split(X, y, test_size=.3, random_state=0)
```

```
# Fit the Logistic Regression model to the training data.
```

```
classifier = LogisticRegression()  
classifier.fit(X_train, y_train)
```

```
# Obtain the probability scores for the positive class (index 1) on the test set.  
y_score = classifier.predict_proba(X_test)
```

Step 3: Generating and Plotting the Curve

The technical core of this visualization process lies in the seamless execution of the Scikit-learn utility function, `precision_recall_curve`. This function is specifically designed to accept two inputs: the ground truth binary labels of the test data (`y_test`) and the continuous probability scores assigned to the positive class (`y_score`). It then efficiently calculates the corresponding precision and recall values for every unique classification threshold implicitly defined by the probability scores.

The `precision_recall_curve` function returns three essential numpy arrays: the precision values, the recall values, and the associated thresholds used to generate those metric pairs. We subsequently feed these calculated precision and recall arrays directly into [Matplotlib](#)'s powerful plotting utilities. The plotting involves first instantiating a figure and axes object, followed by using the plot method to render the curve itself, connecting the calculated (recall, precision) points.

To ensure maximum clarity and interpretability of the final output, we meticulously define the plot's components. This includes setting a descriptive title and ensuring the axis labels are clearly marked: precision is explicitly placed on the Y-axis and recall on the X-axis, adhering to standard conventions for this type of performance metric visualization.

```
# Calculate precision, recall, and thresholds from the test data and probability scores.  
precision, recall, thresholds = precision_recall_curve(y_test, y_score)
```

```
# Create the precision-recall curve plot using Matplotlib.
```


```
fig, ax = plt.subplots()  
ax.plot(recall, precision, color='purple')
```

```
# Add descriptive labels to the plot axes and title.
```

```
ax.set_title('Precision-Recall Curve')  
ax.set_ylabel('Precision')  
ax.set_xlabel('Recall')
```

```
# Display the generated plot visualization.
```

```
plt.show()
```



Interpreting the Precision-Recall Tradeoff

The final generated plot offers a highly condensed and powerful summary of the underlying classification model's behavior across varying sensitivity levels. The horizontal axis maps the achieved recall, while the vertical axis displays the corresponding precision for those points, defined by different decision thresholds. Conceptually, traversing the curve from the extreme left toward the right is equivalent to systematically lowering the probability threshold required for an instance to be classified as positive.

A key observation when analyzing this curve is the inverse relationship: as recall increases, precision almost invariably experiences a corresponding decrease. This dynamic illustrates the fundamental and unavoidable **precision-recall tradeoff** that is inherent in nearly all binary classifiers. To enhance the recall of the model--meaning we become more aggressive in capturing positive cases, thereby mitigating False Negatives--we must necessarily adopt a lower classification threshold. This lower threshold, however, simultaneously increases the volume of False Positives, which in turn drives the overall precision score down.

Grasping the nuances of this tradeoff is absolutely vital for the successful deployment of any predictive system. If the financial or human cost associated with a False Negative (e.g., failing to

detect a critical event) is extremely high, the data scientist would strategically select an operational threshold point on the curve that maximizes recall, even if this choice requires sacrificing some level of precision. Conversely, in situations where False Positives must be rigorously minimized (such as initiating an expensive or disruptive alarm), the user would be compelled to select a point on the curve that maximizes precision, even if it means accepting a slightly lower capacity for recall.

Further Reading and Advanced Resources

For expert practitioners and readers interested in deepening their technical understanding of classification evaluation methodologies, confusion matrices, and advanced derivative metrics such as the F1-score and Area Under the Curve (AUC), the following authoritative resources are highly recommended for detailed study:

Official [Scikit-learn documentation on Classification Metrics](#), which includes detailed explanations of metric calculations, plotting routines, and advanced concepts like Area Under the Curve (AUC) calculations specific to PR curves.

Comprehensive guides detailing the statistical foundations of [Logistic Regression](#) and its crucial application in estimating outcome probabilities for robust binary classification problems.

In-depth academic and technical articles comparing the utility of ROC curves versus Precision-Recall curves, focusing specifically on analytical scenarios involving highly imbalanced data where PR curves offer demonstrably more reliable and less misleading performance insights.