

# How to Perform Linear Discriminant Analysis for Classification

Authored by  
**stats writer**

January 23, 2026

## RECOMMENDED CITATION

stats writer (2026). *How to Perform Linear Discriminant Analysis for Classification*.  
PSYCHOLOGICAL SCALES. Retrieved from <https://scales.arabpsychology.com/?p=127156>

**Linear Discriminant Analysis (LDA)** is a powerful and classical statistical technique utilized primarily for two related purposes: classification and dimensionality reduction. At its core, LDA seeks to determine the optimal linear combination of features that maximizes the separation between distinct classes within a dataset. Unlike methods like Principal Component Analysis (PCA), which focuses on maximizing variance without regard to class separation, LDA operates by explicitly maximizing the ratio of between-class variance to within-class variance. This fundamental objective makes it highly effective for creating robust decision boundaries that cleanly distinguish between groups, whether applied to binary or multiclass problems. Consequently, LDA is a staple in fields like machine learning, pattern recognition, and data analysis, providing an efficient pathway to improving the accuracy and interpretability of predictive models.

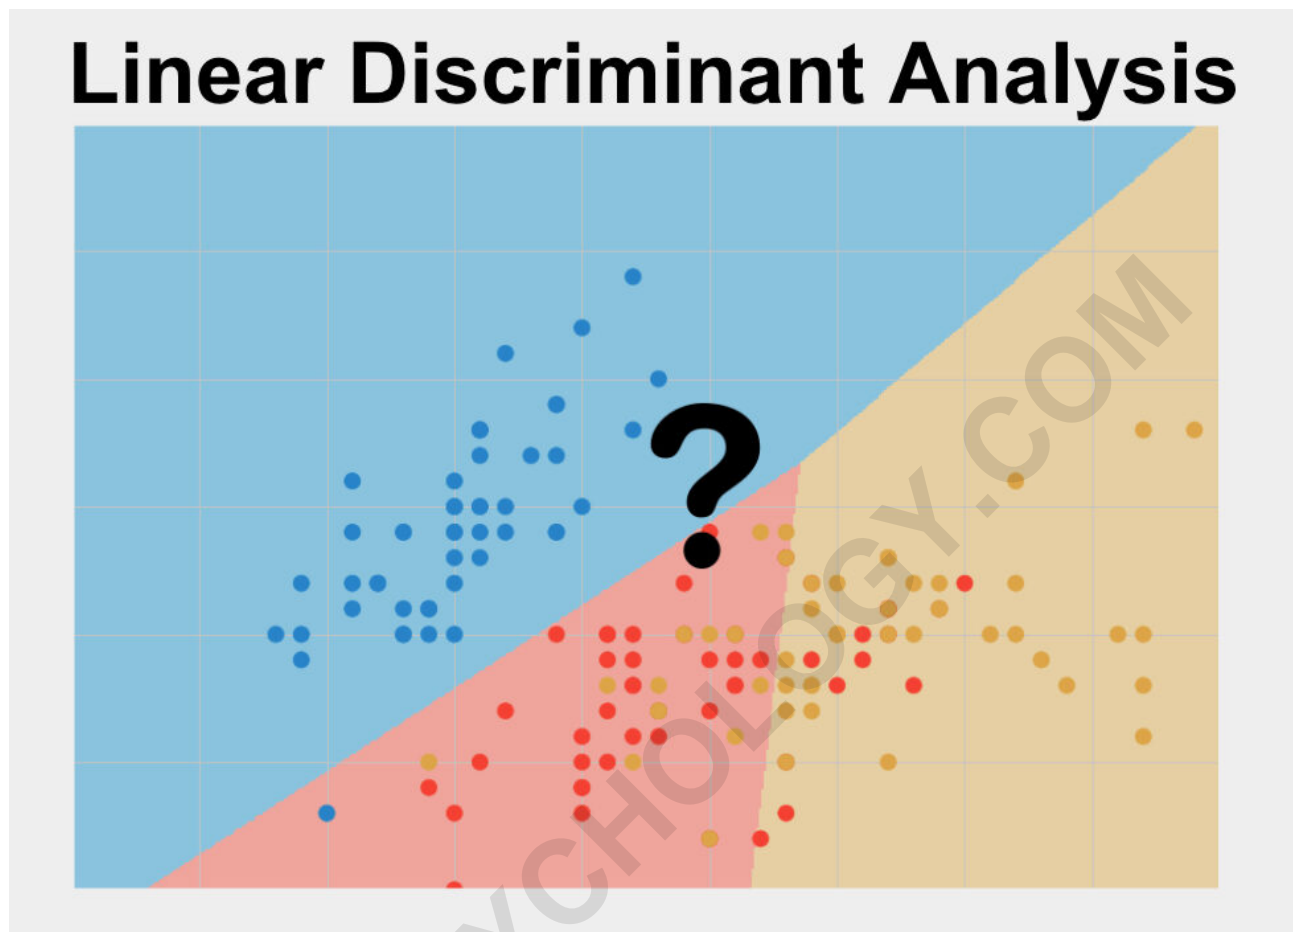
### What is Linear Discriminant Analysis?

Linear Discriminant Analysis is fundamentally a statistical test designed to predict a single categorical variable (the dependent variable) based on the input of one or more continuous variables (the predictor variables). Beyond simple prediction, LDA is also invaluable for determining the underlying numerical relationship and structure that allows for the separation of these variables into discrete groups. The result is a set of discriminant functions, which are linear combinations of the predictors, designed to best differentiate the classes. This makes it an essential tool when the goal is to develop a predictive model where the outcome variable represents distinct, non-ordered categories, such as customer segment, product preference, or species type.

The methodology of LDA hinges on the concept of projecting the high-dimensional feature space onto a lower-dimensional space. The key constraint is that this projection must retain the maximum amount of class separation. To achieve this, LDA calculates the means and variances for each class and then attempts to find an axis upon which the data points are projected such that the distance between the projected means of the classes is large, while the variance (or spread) within each class is minimized. This crucial balance ensures that the resulting model possesses high predictive power and is not overly sensitive to individual data points within the defined categories. This statistical technique is highly utilized because it provides both insight into the variables driving separation and a functional model for future predictions.

While powerful, the application of LDA requires careful adherence to several underlying statistical conditions. For instance, the technique assumes that the data for each class is drawn from a multivariate normal distribution and that the covariance matrices for all classes are equal (homoscedasticity). If the data severely violates these underlying assumptions, the resulting discriminant functions may be biased or inefficient, leading to inaccurate predictions and misleading interpretation of the relationships between the continuous variables and the target categorical variable. Therefore, careful preparation and validation of the dataset against these

preconditions are mandatory steps before applying LDA.



*Linear Discriminant Analysis is sometimes also called normal discriminant analysis (NDA), or discriminant function analysis (DFA).*

## Critical Assumptions for Linear Discriminant Analysis

As with any rigorous statistical method, Linear Discriminant Analysis relies on a set of core assumptions regarding the characteristics of the input data. These assumptions are not merely suggestions; they are properties that your data must satisfy to ensure the statistical validity, accuracy, and interpretability of the LDA results. Failure to meet these conditions can lead to unreliable discriminant functions and ultimately compromise the conclusions drawn from the analysis. Understanding and testing for these assumptions is a crucial precursor to implementing LDA in any research or business setting.

The primary assumptions underlying Linear Discriminant Analysis are related to the distribution and relationship among the continuous variables used as predictors. Violations often necessitate data transformation, the use of robust statistical alternatives, or sometimes, the selection of an

entirely different classification technique, such as Quadratic Discriminant Analysis (QDA) which relaxes the equal covariance assumption, or non-parametric methods. It is standard practice to document the assessment of these assumptions thoroughly when reporting the outcomes of an LDA model.

The fundamental assumptions for successful application of Linear Discriminant Analysis include:

Linearity in the separation of classes.

Absence of influential Outliers.

Independence of observations.

Absence of severe Multicollinearity among predictors.

Equality of Covariance Matrices (Homoscedasticity).

Multivariate Normality of the predictor variables within each class.

Let's delve into the details of each of these requirements and discuss why they are so vital to the integrity of the LDA model.

## Linearity

The core mechanism of Linear Discriminant Analysis dictates that the decision boundaries separating the classes must be linear. This means LDA assumes that a straight line (or a hyperplane in higher dimensions) is sufficient to achieve optimal separation between the categories. For this assumption to hold, the relationship between the predictor variables and the discriminant function--the score used for classification--must be linear. If the true relationship between the predictors and the class membership is complex and inherently non-linear (e.g., curved or parabolic), LDA will fail to capture the optimal separation boundaries, leading to poor model performance. In cases of non-linearity, alternatives like QDA or kernel methods should be considered, as they are capable of modeling more intricate decision surfaces.

To test for linearity, one might examine the standardized residuals or visually inspect scatter plots of the discriminant scores against the predictors. If a non-linear pattern is evident, transformations of the predictor variables might resolve the issue, or, alternatively, the non-linear structure itself indicates that LDA is not the best choice for the specific dataset. Ensuring linearity is paramount because the entire mathematical foundation of LDA is based on constructing the most effective linear boundary.

## No Outliers

The presence of influential outliers in the predictor variables can severely compromise the stability and accuracy of Linear Discriminant Analysis. Outliers are data points that exhibit unusually large or small values far removed from the bulk of the data distribution. Since LDA relies on calculating

means and covariance matrices--both of which are highly sensitive to extreme values--a single outlier can disproportionately skew the class mean vectors and inflate the within-class variance. This distortion leads to inaccurate estimation of the optimal discriminant function and thus poorer separation between classes.

Identifying outliers is typically done through visual inspection using box plots or scatter plots, or statistically using metrics like Mahalanobis distance. If outliers are detected, careful investigation is necessary to determine if they are errors (which should be corrected or removed) or genuine but extreme observations. If they are genuine, robust versions of LDA or data capping/winsorizing techniques may be employed to mitigate their influence before running the primary analysis. Maintaining data cleanliness by minimizing the impact of outliers ensures that the calculated discriminant functions accurately represent the typical structure of the classes.

## Independence of Observations

The assumption of independence mandates that each observation (data point) in the dataset must be statistically independent of all other observations. In simple terms, the value of one case should not be influenced by, or related to, the value of another case. Violation of this assumption is common in data that involves repeated measurements, nested structures, or temporal dependencies. For instance, if data points are collected sequentially over time from the same subject, those observations are likely correlated, thus violating the independence requirement.

When observations are dependent, the standard errors calculated by LDA are often underestimated, leading to inflated Type I error rates (i.e., falsely declaring a significant result). To address potential dependence issues, researchers must carefully review the data collection protocol. If dependence is unavoidable (e.g., longitudinal data), specialized statistical techniques designed for correlated data, such as mixed-effects models or time-series analysis, should be used instead of standard LDA.

## No Multicollinearity

Multicollinearity describes a scenario where two or more of the continuous variables (predictors) are highly correlated with each other. While the presence of multicollinearity does not necessarily affect the predictive accuracy of the overall model (i.e., how well it fits the data), it critically undermines the reliability and interpretability of the individual regression coefficients within the discriminant functions. When multicollinearity is present, the model struggles to isolate the unique contribution of each correlated predictor, making the coefficients unstable, prone to large standard errors, and statistically untrustworthy.


To diagnose multicollinearity, researchers often examine the correlation matrix among the predictors or calculate Variance Inflation Factors (VIFs). High VIF values suggest a severe

violation. Resolution strategies include removing one of the highly correlated predictors, combining the correlated variables into a composite score, or utilizing methods like Principal Component Regression (PCR) which are designed to handle correlated features. Maintaining low multicollinearity ensures that the interpretation of which variables drive class separation remains clear and accurate.


### Equality of Covariance Matrices (Homoscedasticity)

In the context of LDA, the assumption of equal spread across the range is known as the equality of covariance matrices, or homoscedasticity. This crucial assumption requires that the variance and covariance structure of the predictor variables must be roughly equivalent across all the different categorical classes. In other words, the scatter or shape of the data cloud for the predictors should look similar regardless of which group (class) is being observed.

If the covariance matrices are significantly unequal (a state known as heteroscedasticity), the standard LDA model--which averages the variances across groups--will produce suboptimal classification boundaries, potentially favoring the class with the smaller variance. The Box's M test is a common statistical procedure used to assess the equality of covariance matrices. If this assumption is violated, the appropriate alternative is to use Quadratic Discriminant Analysis (QDA), which does not assume equal covariance and instead calculates a separate covariance matrix for each class, resulting in quadratic (curved) decision boundaries rather than linear ones.



These data have a similar spread across their range.



These data have greater spread at higher values.

### Multivariate Normality

The most demanding assumption for Linear Discriminant Analysis is that the predictor variables, when considered together, follow a multivariate normal distribution within each class. Univariate normality (where each predictor variable individually follows a bell curve shape) is a necessary but insufficient condition for multivariate normality. While LDA is considered relatively robust to minor departures from normality, severe non-normality--especially skewness or heavy tails--can negatively impact the stability and predictive performance of the discriminant functions.

Testing for multivariate normality is complex, often relying on graphical checks like Q-Q plots of Mahalanobis distances or specific statistical tests like the Mardia's test. If non-normality is detected, data transformations (such as logarithmic or square root transformations) can sometimes restore approximate normality. If transformations fail, researchers may need to consider non-parametric alternatives or highly flexible models that do not rely on distributional assumptions.

### When to Strategically Use Linear Discriminant Analysis

Selecting the appropriate statistical method depends entirely on the nature of the research question and the type of data involved. Linear Discriminant Analysis is specifically tailored for

scenarios involving group differentiation and classification, making it ideal when the goal is to develop a predictive model that assigns new observations to predefined categories. It differs significantly from correlation (which measures association strength) and standard regression (which often predicts continuous outcomes).

You should elect to use [Linear Discriminant Analysis](#) when your research objectives and data structure align perfectly with these three requirements:

The primary goal is either **prediction** or quantifying the numerical distinction between groups.

The variable you aim to predict (the [dependent variable](#)) is strictly **categorical**.

All variables used for prediction (the independent variables) are **continuous** measurements.

A deeper examination of these criteria will clarify the specific niche that LDA fills within the predictive modeling landscape.

## Prediction and Discrimination

The fundamental strength of LDA lies in its ability to classify. When the objective is to build a model that can accurately assign a new, unlabeled observation into one of several known groups, LDA is a prime candidate. This is a prediction question, differentiating it from purely explanatory analyses aimed solely at understanding mechanisms. The technique constructs a series of discriminant functions that serve as efficient rules for allocating future cases. These functions are derived precisely to maximize the separation between group centers, thereby minimizing the misclassification rate. LDA provides a holistic view of group differentiation, moving beyond simply examining differences between means to constructing optimal classification boundaries.

Furthermore, LDA also provides insight into which predictor variables contribute most significantly to the separation of classes. By analyzing the standardized coefficients of the discriminant functions, researchers can understand the relative importance and directionality of the relationship between predictors and group membership. This dual capability--both effective prediction and insightful interpretation--makes it a preferred choice in fields requiring strong discriminatory power, such as risk assessment or diagnostic modeling.

## Categorical Dependent Variable Requirement

A non-negotiable requirement for using [Linear Discriminant Analysis](#) is that the outcome variable, or [dependent variable](#), must be [categorical](#). A categorical variable classifies observations into discrete groups or categories that typically do not possess intrinsic numerical order. Classic examples include nominal data such as eye color (blue, brown, green), political affiliation, or city of birth. LDA is optimally designed for these nominal or multiclass outcomes.

It is critical to distinguish categorical data from other data types that are not suitable for LDA's dependent variable role. Data that is NOT categorical includes: ordered or ordinal data (like finishing place in a race, or satisfaction rankings), binary data (true/false, purchase/no purchase), or continuous data (height, temperature, income). If the dependent variable were continuous, the appropriate method would likely be Simple Linear Regression. If the dependent variable is binary (dichotomous), Simple Logistic Regression is the standard choice, as LDA performs best when dealing with three or more distinct classes, although it can be applied to binary classification.

*If your dependent variable is continuous, you should use Simple Linear Regression, and if your dependent variable is binary, then you should use Simple Logistic Regression.*

## Continuous Predictor Variables

Linear Discriminant Analysis mandates that the variables used to predict the categorical outcome--the independent variables--must all be continuous. Continuous variables are numerical measurements that can take on any value within a given range, offering a high level of detail and granularity (e.g., age, financial income, time spent on a website). LDA leverages the variance and covariance structure of these continuous measures to establish the discriminant functions.

If the independent variables include mixed data types (e.g., both continuous and categorical predictors), the use of LDA becomes problematic without specific adaptations. While specialized techniques exist to incorporate categorical predictors (usually through dummy coding), the standard and most robust application of LDA requires that all predictors provide metric-level data. If you have a mixture of variable types, or if you have many categorical predictors, alternative models such as Multinomial Logistic Regression might offer a more statistically straightforward approach, as they naturally handle categorical inputs.

*If your independent variables are a mix of continuous and categorical types, or if the dependent variable is multiclass, then you can use Multinomial Logistic Regression.*

## Detailed Linear Discriminant Analysis Example

Consider a marketing research scenario where a company wants to understand what characteristics differentiate consumers who prefer three distinct website formats (A, B, or C). The company suspects that demographic factors play a significant role in format preference. This scenario perfectly aligns with LDA requirements:

**Dependent Variable:** Website format preference (e.g. format A, B, C) -- A categorical variable.

**Independent Variable 1:** Consumer age (measured in years) -- A continuous variable.

**Independent Variable 2:** Consumer income (measured annually in dollars) -- A continuous variable.

The statistical starting point is always the null hypothesis ( $H_0$ ), which proposes that the treatment or variables have no effect. In this case,  $H_0$  states that there is no meaningful relationship between consumer age and income and their preference for website format. The LDA test is designed to assess the probability of this null hypothesis being true given the observed data. If the results show strong statistical significance, we reject  $H_0$ , concluding that age and income, in combination, are effective discriminators of website preference.

During data preparation, the researchers gather their sample and observe that consumer income exhibits a right-skewed distribution, meaning it does not meet the normality assumption. To correct this violation, they apply a log transformation to the income variable, successfully normalizing its distribution. After confirming all assumptions (including homoscedasticity), they proceed with running the LDA. The analysis yields one or more statistically significant discriminant functions that optimally separate the three groups (A, B, and C) based on the input variables.

One essential output of Linear Discriminant Analysis is a set of classification rules or formulas that mathematically describe the decision boundaries between the website format preferences as a function of consumer age and income. These functions can be visualized in a lower-dimensional space, clearly showing the separation. Most importantly, the results of this analysis can be used prospectively to predict the website preference for any new consumer based solely on their age and income data, providing the marketing team with a powerful tool for customer segmentation and targeted design optimization.