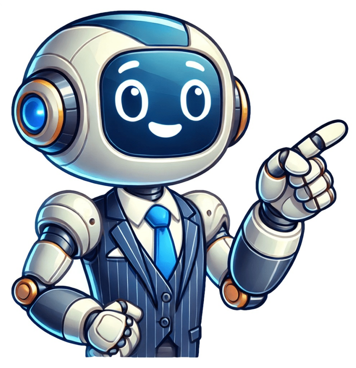

One of the most common assumptions for statistical tests is that the data used are normally distributed. For example, if you want to run a t-test or an ANOVA, you must first test whether the data or variables are normally distributed. The assumption of normal distribution is also important for linear regression analysis, but in this case it is important that the error made by the model is normally distributed, not the data itself. If the data are not normally distributed, the above procedures cannot be used and non-parametric tests must be used. Non-parametric tests do not assume that the data are normally distributed. How is the normal distribution tested? Normal distribution can be tested either analytically (statistical tests) or graphically. The most common analytical tests to check data for normal distribution are the: Kolmogorov-Smirnov Test Shapiro-Wilk Test Anderson-Darling Test For graphical verification, either a histogram or, better, the Q-Q plot is used. Q-Q stands for quantile-quantile plot, where the actually observed distribution is compared with the theoretically expected distribution. Statistical tests for normal distribution To test your data analytically for normal distribution, there are several test procedures, the best known being the Kolmogorov-Smirnov test, the Shapiro-Wilk test, and the Anderson Darling test. In all of these tests, you are testing the null hypothesis that your data are normally distributed. The null hypothesis is that the frequency distribution of your data is normally distributed. To reject or not reject the null hypothesis, all these tests give you a p-value. What matters is whether this p-value is less than or greater than 0.05. If the p-value is less than 0.05, this is interpreted as a significant deviation from the normal distribution and it can be assumed that the data are not normally distributed. If the p-value is greater than 0.05 and you want to be statistically clean, you cannot necessarily say that the frequency distribution is normal, you just cannot reject the null hypothesis. In practice, a normal distribution is assumed for values greater than 0.05, although this is not entirely correct. Nevertheless, the graphical solution should always be considered. Note: The Kolmogorov-Smirnov test and the Anderson-Darling test can also be used to test distributions other than the normal distribution. Disadvantage of the analytical tests for normal distribution Unfortunately, the analytical method has a major drawback, which is why more and more attention is being paid to graphical methods. The problem is that the calculated p-value is affected by the size of the sample. Therefore, if you have a very small sample, your p-value may be much larger than 0.05, but if you have a very very large sample from the same population, your p-value may be smaller than 0.05. If we assume that the distribution in the population deviates only slightly from the normal distribution, we will get a very large p-value with a very small sample and therefore assume that the data are normally distributed. However, if you take a larger sample, the p-value gets smaller and smaller, even though the samples are from the same population with the same distribution. With a very large sample, you can even get a p-value of less than 0.05, rejecting the null hypothesis of normal distribution. To avoid this problem, graphical methods are increasingly being used. Graphical test for normal distribution If the normal distribution is tested graphically, one looks either at the histogram or even better the QQ plot. If you want to check the normal distribution using a histogram, plot the normal distribution on the histogram of your data and check that the distribution curve of the data approximately matches the normal distribution curve. A better way to do this is to use a quantile-quantile plot, or Q-Q plot for short. This compares the theoretical quantiles that the data should have if they were perfectly normal with the quantiles of the measured values. If the data were perfectly normally distributed, all points would lie on the line. The further the data deviates from the line, the less normally distributed the data is. In addition, DATAtab plots the 95% confidence interval. If all or almost all of the data fall within this interval, this is a very strong indication that the data are normally distributed. They are not normally distributed if, for example, they form an arc and are far from the line in some areas. Test Normal distribution in DATAtab When you test your data for normal distribution with DATAtab, you get the following evaluation, first the analytical test procedures clearly arranged in a table, then the graphical test procedures. If you want to test your data for normal distribution, simply copy your data into the table on DATAtab, click on descriptive statistics and then select the variable you want to test for normal distribution. Then, just click on Test Normal Distribution and you will get the results. Furthermore, if you are calculating a hypothesis test with DATAtab, you can test the assumptions for each hypothesis test, if one of the assumptions is the normal distribution, then you will get the test for normal distribution in the same way. many illustrative examples ideal for exams and theses statistics made easy on 454 pages 6th revised edition (March 2025) Only 8.99 € Free sample "It could not be simpler" "So many helpful examples" Scientists, engineers, and business people need to understand what kind of patterns their data follows. The normal distribution is one of the most common. But before using techniques that assume normal patterns, it's good to check if the data matches the normal distribution shape. The Anderson-Darling normality test is a useful tool for checking how well data fits with the normal distribution. The Anderson-Darling test gives a number showing how much a data set deviates from the normal distribution shape. This lets researchers decide if normal theory should be used in their data analysis or not. Anderson-Darling normality test examines datasets checking if they follow special patterns. Most often it's checking for normal distribution matching. This normality test helps decide if data significantly diverges from normalcy's form. It provides a detailed measure showing how close data aligns with predicted distribution models, usually normal distribution. Wide use exists among quality controllers, production supervisors, and distribution supervisors trying patterns correctly. Overall it prevents faulty assumptions misdirecting efforts. Anderson-Darling normality test checks whether datasets follow certain patterns, especially the normal distribution shape. It was created by Theodore W. Anderson and Donald A. Darling in 1952 to assess how much a data set matches the normal distribution shape. The test has become very popular for checking normality properly, especially toward the edges of distributions. It calculates a statistic judging divergence between the observed info and assumed distribution. Usually, it's sized up against the normal distribution. The Anderson-Darling statistic builds upon the Kolmogorov-Smirnov test, emphasizing deviations toward distributions' ends more sensitively. This makes it adept at detecting tail differences, an edge propelling its prevalence across sampling sizes and applications from manufacturing to healthcare. Normality testing is a crucial step in many statistical analyses and applications. The assumption of normality is a prerequisite for numerous parametric tests, such as t-tests, ANOVA, and regression analysis. Violating the normality assumption can lead to invalid conclusions and unreliable results. Moreover, many natural phenomena and processes are assumed to follow a normal distribution, making normality testing essential in various fields, including manufacturing, finance, biology, and physics. By verifying the normality assumption, researchers and practitioners can make informed decisions and draw accurate inferences from their data. The Anderson-Darling test offers several advantages over other normality tests, such as the Kolmogorov-Smirnov test and the Shapiro-Wilk test: Increased sensitivity: The Anderson-Darling test is more sensitive to deviations from normality, especially in the tails of the distribution, making it better suited for detecting non-normal distributions. Versatility: The test can be applied to complete samples or truncated data, allowing for greater flexibility in data analysis. Efficiency: The Anderson-Darling test is efficient for detecting various types of non-normal distributions, including skewed, heavy-tailed, and bimodal distributions. Ease of interpretation: The test provides a clear decision criterion based on the calculated test statistic and the corresponding critical values or p-value, facilitating a straightforward interpretation of the results. By utilizing the Anderson-Darling normality test, researchers and practitioners can make informed decisions about the appropriateness of parametric tests and ensure the validity of their statistical analyses. The Anderson-Darling normality test is a statistical procedure used to determine whether a given dataset follows a specified probability distribution, such as the normal distribution. It falls under the category of normality tests, which are essential tools in many areas of statistics and data analysis. In the context of normality testing, the hypothesis being tested is whether the sample data comes from a normally distributed population. The null hypothesis (H0) states that the data follows a normal distribution, while the alternative hypothesis (H1) suggests that the data deviates from normality. Hypothesis testing for normality involves calculating a test statistic and comparing it to a critical value or computing a p-value. If the test statistic exceeds the critical value or if the p-value is below the chosen significance level (typically 0.05), the null hypothesis of normality is rejected, indicating that the data is unlikely to be normally distributed. The Anderson-Darling normality test statistic is a measure of the distance between the empirical distribution function (EDF) of the sample data and the cumulative distribution function (CDF) of the hypothesized distribution, in this case, the normal distribution. The Anderson-Darling test statistic is calculated using the following formula: A^2 = -n – (1/n) * Σ[(2i-1) * (ln(F(y_i)) + ln(1-F(y_(n+1-i))))] Where: n is the sample size y_i are the ordered data points F(y_i) is the cumulative distribution function of the hypothesized distribution evaluated at y_i The Anderson-Darling test statistic gives more weight to the tails of the distribution compared to other normality tests, making it more sensitive to deviations from normality in the tails. To determine whether the null hypothesis of normality should be rejected or not, the calculated Anderson-Darling normality test statistic is compared to critical values or a p-value is computed. Critical values for the Anderson-Darling test are tabulated and depend on the sample size and the chosen significance level (e.g., 0.05 or 0.01). If the test statistic exceeds the critical value, the null hypothesis of normality is rejected. Alternatively, a p-value can be calculated, which represents the probability of observing a test statistic as extreme or more extreme than the calculated value, assuming the null hypothesis is true. If the p-value is less than the chosen significance level, the null hypothesis of normality is rejected. The p-value for the Anderson-Darling test can be computed using various methods, such as interpolation from tabulated values, approximation formulas, or Monte Carlo simulations. By comparing the test statistic or p-value to the chosen significance level, researchers can make a decision about whether to reject or fail to reject the null hypothesis of normality for the given data set. The Anderson-Darling normality test, like other statistical tests, has certain assumptions and requirements that must be met for the results to be valid. Understanding these assumptions is crucial to ensure proper application and interpretation of the test. The sample size plays a significant role in the performance of the Anderson-Darling normality test. While there is no strict lower limit on the sample size, the test is generally recommended for moderate to large sample sizes (e.g., greater than 20 observations). With very small sample sizes, the test may have reduced power to detect deviations from normality. However, it is important to note that the Anderson-Darling test is more powerful than some other normality tests, such as the Kolmogorov-Smirnov test, for small to moderate sample sizes. The Anderson-Darling normality test requires a single continuous variable or a set of independent and identically distributed (i.i.d.) observations. The data should be numerical and measured on at least an interval scale. Categorical or ordinal data cannot be directly tested for normality using the Anderson-Darling normality test. The Anderson-Darling normality test is specifically designed to assess whether a given set of data follows a normal (or Gaussian) distribution. It tests the null hypothesis that the data comes from a normal population against the alternative hypothesis that the data does not follow a normal distribution. If the normality assumption is violated, the test results may be invalid, and alternative methods for analyzing non-normal data may be required. It is important to note that the Anderson-Darling test assumes that the parameters of the normal distribution (mean and standard deviation) are unknown and must be estimated from the sample data. This assumption distinguishes it from other normality tests, such as the Kolmogorov-Smirnov test, which assumes that the parameters are known. The Anderson-Darling normality test for normality can be performed manually by following these steps: Arrange the data in ascending order: x(1), x(2), ..., x(n). Calculate the sample mean and standard deviation. Standardize the data using the formula: (x(i) – sample mean) / sample standard deviation. Calculate the normal scores for each standardized data point using the inverse cumulative distribution function of the standard normal distribution. Calculate the Anderson-Darling normality test statistic using the formula:A^2 = -n – (1/n) * Sum[(2i – 1) * (ln(Φ(x(i))) + ln(1 – Φ(x(n+1-i)))]where n is the sample size, Φ is the cumulative distribution function of the standard normal distribution, and x(i) are the ordered data points. Compare the calculated A^2 value with the critical value from the Anderson-Darling table for the chosen significance level (e.g., 0.05). If A^2 exceeds the critical value, reject the null hypothesis that the data follows a normal distribution. While manual calculation is possible, it can be tedious and error-prone, especially for larger sample sizes. Most statistical software packages provide built-in functions or tools to perform the Anderson-Darling normality test with just a few clicks or lines of code. This eliminates the need for manual calculations and reduces the risk of errors. Excel: The NORM.DIST function can be used to calculate the normal scores, and then the Anderson-Darling statistic can be computed using the formula. However, Excel does not provide a built-in function for the critical values, so you may need to refer to published tables or use an add-in. Minitab: Minitab offers a dedicated menu option for the Anderson-Darling normality test under the "Stat > Basic Statistics > Normality Test" menu. It provides the test statistic, p-value, and a graphical display of the results. R: The "nortest" package in R includes the "ad.test" function, which performs the Anderson-Darling normality test and returns the test statistic, p-value, and other relevant information. Python: The "scipy.stats" module in Python provides the "Anderson" function, which calculates the Anderson-Darling normality test statistic and critical values for a given significance level. Using software simplifies the process and ensures accurate calculations, especially for larger datasets or when performing multiple tests. The interpretation of the Anderson-Darling test results depends on the test statistic (A^2) and the chosen significance level (α, typically 0.05). Calculate the p-value associated with the test statistic. The p-value represents the probability of observing a test statistic as extreme as or more extreme than the calculated value, assuming the null hypothesis (data follows a normal distribution) is true. Compare the p-value with the chosen significance level (α): If the p-value is less than α, reject the null hypothesis. This suggests that the data does not follow a normal distribution at the given significance level. If the p-value is greater than or equal to α, fail to reject the null hypothesis. This indicates that there is insufficient evidence to conclude that the data deviates from a normal distribution. Alternatively, compare the test statistic (A^2) with the critical value from the Anderson-Darling table for the chosen significance level: If A^2 exceeds the critical value, reject the null hypothesis, suggesting that the data is not normally distributed. If A^2 is less than or equal to the critical value, fail to reject the null hypothesis, indicating that the data is consistent with a normal distribution. It's important to note that the Anderson-Darling normality test has higher power than some other normality tests, particularly for detecting deviations from normality in the tails of the distribution. However, like any statistical test, the results should be interpreted in the context of the specific application and data characteristics. The Anderson-Darling normality test is one of several statistical tests used to assess whether a sample of data follows a normal distribution. It's important to understand how it compares to other popular normality tests in terms of strengths, limitations, and appropriate use cases. The Shapiro-Wilk test is another widely used normality test, especially for smaller sample sizes (n2000), the K-S may be preferred due to the difficulties in calculating critical values for the Anderson-Darling distribution. But for most practical applications with smaller to moderate samples, Anderson-Darling is favored. The Anderson-Darling test has several key strengths: It is a powerful test for detecting many types of non-normal distributions It gives more weight to the tails than the K-S test, making it sensitive to outliers Critical values are widely available for a range of sample sizes It makes use of the specific distribution being tested (normal, lognormal, etc.) However, some limitations include: For very small samples (n2000), calculating critical values becomes difficult Like all tests, it has reduced power for certain alternative distributions The assumptions of randomness and continuity must be met Overall, the Anderson-Darling is an excellent general-purpose option for normality testing across many fields when the sample is not extremely large or small. Understanding how it compares to other tests is important for proper application and interpretation. The Anderson-Darling normality test finds extensive applications in quality control and process monitoring. In manufacturing and production environments, it is crucial to ensure that the processes are operating within specified limits and producing outputs that conform to quality standards. The Anderson-Darling test can be used to verify the normality assumption of the process data, which is a common requirement for many statistical process control (SPC) techniques, such as control charts. By checking the normality of process data, manufacturers can detect deviations from the expected distribution, which may indicate the presence of assignable causes or special causes of variation. This information can then be used to take corrective actions and bring the process back into a state of statistical control, ultimately improving product quality and reducing waste. Another important application of the Anderson-Darling test is in distribution fitting. Many statistical models and analyses assume that the underlying data follows a specific probability distribution, such as the normal, exponential, or Weibull distribution. The Anderson-Darling test can be used to assess the goodness-of-fit of the data to these theoretical distributions. By determining the most appropriate distribution for the data, analysts can make more accurate inferences, predictions, and decisions. For example, in reliability engineering, the Anderson-Darling test can be used to fit a Weibull distribution to failure time data, which can then be used to estimate the reliability of a product or system. Closely related to distribution fitting is the use of the Anderson-Darling test for goodness-of-fit testing. In this context, the test is used to evaluate whether a sample of data follows a hypothesized distribution, such as the normal or exponential distribution. This is particularly important in fields like finance, where the assumption of normality is often made for asset returns or risk models. By performing goodness-of-fit tests, researchers and analysts can validate the assumptions underlying their statistical models and ensure that the chosen distribution is appropriate for the data at hand. If the data does not follow the assumed distribution, alternative models or transformations may need to be considered. To illustrate the practical applications of the Anderson-Darling normality test, let's consider a few real-world examples: Quality control in semiconductor manufacturing: In the semiconductor industry, the thickness of thin-film layers is a critical quality characteristic.The Anderson-Darling test can be used to verify the normality assumption of the thickness measurements, which is necessary for implementing effective statistical process control techniques. Reliability analysis of wind turbines: In the renewable energy sector, the Anderson-Darling test can be used to assess the goodness-of-fit of wind speed data to a Weibull distribution.This information is crucial for estimating the power output and reliability of wind turbines. Financial risk modeling: In finance, many risk models, such as Value-at-Risk (VaR) calculations, assume that asset returns follow a normal distribution.The Anderson-Darling test can be used to validate this assumption and ensure the accuracy of risk estimates. Environmental data analysis: In environmental studies, the Anderson-Darling test can be applied to various data sets, such as pollutant concentrations or meteorological measurements, to determine if they conform to specific theoretical distributions.This information can aid in modeling and decision-making processes related to environmental management and policy. These examples illustrate the versatility and practical relevance of the Anderson-Darling normality test across various industries and domains, highlighting its importance as a fundamental tool in statistical analysis and decision-making. Many statistical software packages include built-in functions or modules for performing the Anderson-Darling normality test. These packages offer a convenient way to conduct the test and often provide additional features and visualizations. In Excel, the Anderson-Darling test can be performed using add-ins or user-defined functions. One popular add-in is the SPC XL Add-in, which includes an Anderson-Darling normality test function. Minitab, a widely used statistical software, has a dedicated function for the Anderson-Darling test. The "ANDDARLING" function in Minitab calculates the test statistic and p-value, making it easy to perform the test and interpret the results. For users of the R programming language, the "nortest" package provides a comprehensive set of tools for normality testing, including the Anderson-Darling test. The "ad.test" function in this package allows you to perform the test and obtain the test statistic and p-value. Similarly, in Python, the "scipy.stats" module offers the "anderson" function, which calculates the Anderson-Darling test statistic and critical values for a given data set. These software packages often provide additional features, such as graphical representations, hypothesis testing, and integration with other statistical analyses, making them valuable tools for conducting the Anderson-Darling normality test and interpreting the results. What exactly is an Anderson-Darling test? Testing for normality is often the first step in analyzing your data. Many statistical tools you might use have normality as an underlying assumption. If you fail that assumption, you may need to use a different statistical tool or approach. This article will explore what the normality of the data means and how the AD test can be used to confirm whether your data will satisfy the assumption of normality. We will also explain the benefits of the AD test and offer a few best practices for understanding when and how to use the AD test. Overview: What Is the Anderson-Darling Normality Test (AD test)? The Anderson-Darling test is used to test if a sample of data comes from a population with a specific distribution. Its most common use is for testing whether your data comes from a normal distribution. But, what does that mean? Normality refers to a specific statistical distribution called a normal distribution, or sometimes the Gaussian distribution, or a bell-shaped curve. The normal distribution is a symmetrical continuous distribution defined by the mean and standard deviation of the data. The normal distribution is theoretical. What you are testing with the AD test is not whether your data is exactly consistent with a normal distribution, but whether your data is close enough to normal that you can use your statistical tool without concern. In some cases, a statistical tool may be robust to the normality assumption, which means the statistical tool is not overly sensitive to some level of violation of the normality assumption. The normal distribution is popular because it describes many real-life situations, such as the distribution of people's heights, weights, and income. The AD test is a hypothesis test. The null hypothesis (Ho) is that your data is not different from normal. Your alternate or alternative hypothesis (Ha) is that your data is different from normal. You will make your decision about whether to reject or not reject the null based on your p-value. The test statistic for the AD test is: Yes, I know it looks really scary, but don't worry. All the computations can be done by statistical software on your computer. The output you get will include a p-value. Assuming you selected your alpha risk to be 0.05, you will reject the null if the p-value is less than 0.05. That allows you to claim that your data is statistically different from a normal distribution. On the other hand, if your p-value is higher than 0.05, you can state that your data is not statistically different from a normal distribution. Here is an example of a probability plot that provides the results for the AD test. Note that the value of the AD statistic is 0.2307 and the p-value is 0.805. The 0.2307 was calculated from the AD formula above. With a p-value of 0.805, you would fail to reject the null and conclude that your data is not different than normal. This would satisfy any assumption of normality you might need for a statistical test. Let's look at another example. This time, notice that the p-value is 0.0047 based on the AD statistic of 1.1697. In this case, you would reject the null hypothesis and say that your data is different than normal. Benefits of the Anderson-Darling Normality Test (AD test) Knowing the underlying distribution of your data is important so you can apply the most appropriate statistical tools for your analysis. The AD test will help you determine if your data is not normal rather than tell you whether it is normal. Since the normal distribution is a hypothetical distribution, you can't prove that the data is normal. The AD test will tell you if it is not normal or if it is not different from normal, but it cannot tell you if the data is normal. The p-value, which is based on the value of the AD statistic, will provide you guidance on whether to reject or not reject your null hypothesis. In many cases, the computer software you use will provide you with a graphical representation of the data along with the AD value and p-value. This will give you some visual and logical confirmation about your data. Different statistical tools for analysis have different assumptions regarding the underlying distribution of the data that you are analyzing. For example, the t-test has an assumption that the data is normally distributed. Linear regression assumes that the underlying distribution of the residuals is normal. Binary logistic regression has an assumption of the binomial distribution. Others might have an assumption of the F or Chi-Square distributions. You need to understand what these assumptions are regarding your data. Since the AD test is a form of hypothesis testing, you want to correctly state your null and alternative or alternate hypotheses. In the case of the AD test, the null is that your data is not different from a normal distribution. This is what you would want since it is the underlying distribution for your desired statistical tool. The alternate is that it is different from the normal distribution. As the sample size of your data increases, your chances of discovering non-normality increase. Small sample sizes may give you a false reading of normality. If you are using a probability plot, don't be deceived by the impact of the sample size. Let your decision be guided by the p-value. The p-value of your AD test will indicate, with your desired level of risk, whether you can reject your null hypothesis. You need to know what that means so the next action you take is appropriate. A manufacturing manager wanted to confirm whether the recent overhaul of his printing press increased the production rate as promised by the vendor. He had daily run speed data for 15 runs before the overhaul and 17 runs after the overhaul. Further, he wanted to compare the average run speed pre- and post-overhaul. He decided to consult with his Lean Six Sigma Black Belt on how to analyze the data. The LSSBB advised that, since the manager was interested in comparing two sets of continuous data, the appropriate test was the 2-sample t-test. An underlying assumption was that the sample data be normally distributed. The LSSBB was concerned that the sample size of 15 and 17 was small, so the normality assumption couldn't just be ignored. Upon checking the normality of the data with the Anderson-Darling test, the LSSBB found the data not to be normally distributed. Therefore, he was not comfortable just doing the 2-sample t-test. He then also ran a 2-sample Mood's Median test, which tests for the difference between two medians and has no assumption of distribution. Both the t-test and the Mood's Median test resulted in p-values greater than 0.05, which indicated that the overhaul did have an impact — and run speed had increased. It's unlikely you'll do the hand calculations for the AD test. The important issue will be how you collect the data and interpret the results of your AD test. Here are a few thoughts to keep in mind. Your data may not be normal, so have a plan B, or alternate analytical tool, that will still answer your statistical question but doesn't have the same underlying assumptions of the data distribution. In the event of you failing the assumptions for the t-test, you might consider using a median test instead. A random sampling of a statistically valid size will help you get a truer picture of what your data distribution is. This will give you more confidence in the results of your AD test. Often, a simple plot of the data on either a histogram or probability plot will provide you with enough insight into how your data looks. This will keep you from having to do more complicated analysis. We've talked quite a bit about the AD test, but what about the other tools of the trade? You'll likely want a refresher on normality, at least within the context of statistical analysis. Understanding how this standard distribution applies to data as a whole is a crucial concept. Further, other hypothesis tests go more in-depth. The aforementioned Two-Sample T-Test is one of the most common statistical tools used in the Six Sigma methodology. As such, understanding how it works will help you get to the bottom of any perceived differences in your data. Many statistical tools have an assumption that your data is approximately normally distributed. If it's not, you must use a different tool to answer your statistical question. The AD test starts with a null statement that your data is not statistically different than normal. The alternate statement is that it is different from normal. The results you will get will suggest you can either reject the null or fail to reject the null. From there, you can decide how to proceed. The Anderson-Darling statistic measures how well the data follow a particular distribution. For a specified data set and distribution, the better the distribution fits the data, the smaller this statistic will be. For example, you can use the Anderson-Darling statistic to determine whether data meets the assumption of normality for a t-test. The hypotheses for the Anderson-Darling test are: H0: The data follow a specified distribution H1: The data do not follow a specified distribution Use the corresponding p-value (when available) to test if the data come from the chosen distribution. If the p-value is less than a chosen alpha (usually 0.05 or 0.10), then reject the null hypothesis that the data come from that distribution. Minitab does not always display a p-value for the Anderson-Darling test because it does not mathematically exist for certain cases. You can also use the Anderson-Darling statistic to compare the fit of several distributions to determine which one is the best. However, in order to conclude one distribution is the best, its Anderson-Darling statistic must be substantially lower than the others. When the statistics are close together you should use additional criteria, such as probability plots, to choose between them. Distribution Anderson-Darling P-value Exponential 9.599 p < 0.003 Normal 0.641 p < 0.089 3-parameter Weibull 0.376 p < 0.432 These probability plots are for the same data. Both the normal distribution and the 3-parameter Weibull distribution provide a good fit for your data. Minitab calculates the Anderson-Darling statistic using the weighted squared distance between the fitted line of the probability plot (based on the chosen distribution and using either maximum likelihood estimation method or least squares estimates) and the nonparametric step function. The calculation is weighted more heavily in the tails of the distribution. The Anderson-Darling test is a statistical test used to determine if a sample comes from a population with a specific distribution. It is particularly useful for assessing whether a dataset follows a normal distribution. The Anderson-Darling test calculates a statistic that measures how well the data fits the specified distribution. The smaller the statistic, the better the data fits the distribution. The test provides a p-value that can be used to determine if the null hypothesis (that the data follows the specified distribution) should be rejected. If the p-value is less than a chosen significance level (usually 0.05 or 0.10), the null hypothesis is rejected. It is more sensitive to deviations from normality, especially in the tails of the distribution, compared to other normality tests like the Kolmogorov-Smirnov test. It can be applied to complete samples or truncated data, allowing for greater flexibility in data analysis. It is efficient for detecting various types of non-normal distributions, including skewed, heavy-tailed, and bimodal distributions. The test provides a clear decision criterion based on the calculated test statistic and the corresponding critical values or p-value, facilitating a straightforward interpretation of the results. You can perform the Anderson-Darling test in Python using the anderson() function from the scipy.stats module. Here's an example: import numpy as np from scipy.stats import anderson # Generate sample data data = np.random.normal(loc=0, scale=1, size=100) # Perform the Anderson-Darling test result = anderson(data) # Print the test statistic and critical values print(f'Statistic: {result.statistic:.3f}') print('Critical values:') for i, cv in enumerate(result.critical_values): print(f'Significance level: {result.significance_levels[i]}%, Critical value: {cv:.3f}') The output will display the test statistic and the critical values for various significance levels. If the test statistic is greater than or equal to a critical value, the null hypothesis (that the data follows the specified distribution) can be rejected at that significance level. If the p-value is less than the chosen significance level (e.g., 0.05), the null hypothesis is rejected, indicating that the data does not follow the specified distribution. If the p-value is greater than the chosen significance level, there is not enough evidence to reject the null hypothesis, suggesting that the data could follow the specified distribution. When comparing the fit of several distributions, the distribution with the lowest Anderson-Darling statistic is considered the best fit, provided the difference is substantial. If the statistics are close, additional criteria, such as probability plots, should be used to choose between them. By using the Anderson-Darling test, researchers and practitioners can make informed decisions about the appropriateness of parametric tests and ensure the validity of their statistical analyses. Citations:[1] 2] 3] 4] 5] 6] 7] Anderson-Darling Test in R, the Anderson-Darling Test is a statistical test used to determine whether a given dataset is drawn from a particular distribution, such as the normal distribution. In this article, we will demonstrate how to conduct an Anderson-Darling Test in R using inbuilt datasets. Formulation of the Hypothesis: Before conducting the Anderson-Darling Test, it is necessary to formulate the null and alternative hypotheses. The null hypothesis (H0) assumes that the dataset is drawn from a specific distribution. It is usually written as: H0: The dataset is drawn from the specific distribution. The alternative hypothesis (H1) assumes that the dataset is not drawn from the specific distribution. It is usually written as: H1: The dataset is not drawn from the specific distribution. In the following sections, we will provide examples of how to conduct an Anderson-Darling Test in R. The pheatmap function in R » Data Science Tutorials In this example, we will use the inbuilt mtcars dataset to test whether the dataset is normally distributed. First, we will load the mtcars dataset: data(mtcars) Next, we can extract a specific variable from the dataset: mpg