

Learning to Calculate Cohen's d Effect Size in R with Examples

Authored by
Mohammed loot

October 29, 2025

RECOMMENDED CITATION

Mohammed loot (2025). *Learning to Calculate Cohen's d Effect Size in R with Examples*. PSYCHOLOGICAL STATISTICS. Retrieved from <https://statistics.arabpsychology.com/?p=5663>

Understanding the Role of Effect Size in Statistical Analysis

In applied **statistics**, researchers frequently employ hypothesis tests, such as the independent samples t-test, to determine if there is a statistically significant difference between the means of two distinct groups. These tests rely heavily on the computation of a [p-value](#), which helps assess the evidence against a null hypothesis. While crucial for hypothesis testing, the p-value only provides information regarding the probability of observing the data given that the null hypothesis is true; it does not convey the practical importance or magnitude of the observed difference.

The mere existence of a statistically significant difference, indicated by a low p-value, can often be misleading, especially when dealing with very large sample sizes. A small, practically insignificant difference might register as statistically significant. This is where the concept of [effect size](#) becomes indispensable. An effect size metric moves beyond simple significance testing to quantify the actual magnitude of the difference between groups or the strength of a relationship, providing essential context for interpreting research findings.

Among the numerous measures available for quantifying effect size, **Cohen's d** stands out as one of the most widely accepted and utilized methods, particularly when comparing the means of two groups. It is an extremely valuable tool for meta-analysis and for reporting results in fields ranging from psychology to experimental biology. Cohen's d standardizes the mean difference by expressing it in terms of pooled standard deviations, making the measure comparable across different studies and contexts.

Defining and Interpreting Cohen's d

The mathematical formulation of [Cohen's d](#) is designed to measure the separation between two means. Conceptually, it represents the standardized mean difference, effectively telling us how many pooled **standard deviations** the means are apart. This standardization is key because it removes the dependence on the original units of measurement, allowing for standardized comparisons.

The standard formula for calculating Cohen's d, often referred to as the pooled standard deviation method, is given by:

$$\text{Cohen's } d = (x_1 - x_2) / \sqrt{(s_1^2 + s_2^2) / 2}$$

In this formula, the numerator represents the difference between the sample means, and the denominator represents the pooled standard deviation of the two samples. The terms are defined as follows:

x_1 , x_2 : Represents the respective **mean** of sample 1 and sample 2.

s_1^2 , s_2^2 : Represents the respective **variance** of sample 1 and sample 2.

Interpreting the raw value of Cohen's d is straightforward: it indicates the difference between the group means in units of standard deviation. For instance, if you calculate a Cohen's d of 0.5, it signifies that the mean of the first group is half a [standard deviation](#) above the mean of the second group. This provides a clear, scalable metric for understanding the practical significance of the findings, regardless of the original scale (e.g., inches, kilograms, or test scores).

We can illustrate this interpretation with a simple scale:

A d value of **0.5** means the two group means are separated by exactly 0.5 standard deviations.

A d value of **1** signifies that the group means differ by a full 1 standard deviation.

A d value of **2** implies a very large difference, where the group means are 2 **standard deviations** apart.

Standard Rules of Thumb for Interpreting Cohen's d Magnitude

While the direct interpretation of Cohen's d in terms of standard deviations is useful, Jacob Cohen provided conventional benchmarks to help researchers categorize the magnitude of the effect size observed. These "rules of thumb" are widely adopted across many scientific disciplines, though it is critical to remember that context matters; what constitutes a large effect in one field (e.g., medical trials) might be considered small in another (e.g., sociology).

These commonly cited benchmarks, established by Cohen in 1988, facilitate communication and comparison between research findings:

A value of **0.2** generally represents a **small effect size**. This indicates a minimal, often barely noticeable, difference between the groups.

A value of **0.5** typically represents a **medium effect size**. This difference is large enough to be visible and is often considered a standard reference point.

A value of **0.8** or higher represents a **large effect size**. This signifies a substantial and readily apparent difference between the two populations being compared.

For a more intuitive understanding, consider the percentile equivalent. A medium effect size ($d = 0.5$) means that the average person in the superior group performs better than 69% of the people in the inferior group. Conversely, a small effect size ($d = 0.2635$, as we will see in our example) implies that the overlap between the two distributions is extremely high, suggesting that the intervention or group difference has only a marginal practical impact.

Preparing to Calculate Cohen's d in R

The **R programming language** is an essential tool for statistical analysis, offering robust libraries

that simplify complex calculations, including the computation of effect sizes. While one could manually input the means and variances into the mathematical formula, utilizing specialized packages within [R](#) ensures accuracy and efficiency, especially when handling large datasets.

To calculate Cohen's d in R, we must first ensure that the necessary statistical libraries are installed and loaded into the current session. Two of the most reliable and popular packages for this purpose are the **lsr** package (Learning Statistics with R) and the **effsize** package. Both packages offer dedicated functions for calculating effect size measures for two independent samples.

Before running the calculation, the data must be organized into two distinct vectors representing the measurements from each group. The following practical example will demonstrate the implementation of both methods using a simple experimental dataset, highlighting that both methods converge on the same crucial result.

Case Study Example: Comparing Plant Growth Fertilizers

To illustrate the calculation of Cohen's d, consider a common scenario in experimental biology. Suppose a botanist is conducting an experiment to evaluate the efficacy of two distinct fertilizers (Fertilizer #1 and Fertilizer #2) on plant growth. The goal is to determine if there is a substantive difference in the average plant height (measured in inches) after a standard period of one month.

The data collected consists of 12 measurements for each fertilizer group. We are not just interested in whether one fertilizer is statistically superior (a question for a t-test and [p-value](#)), but precisely by how much the mean growth differs in standardized terms (the function of Cohen's d).

We will proceed using the two primary R methodologies available for calculating the standardized mean difference.

Method 1: Calculating Cohen's d Using the lsr Package

The **lsr** package provides a straightforward function, `cohensD()`, which is optimized for calculating the standardized mean difference. This method requires the user to define the data vectors for both groups and pass them directly to the function. This is often the preferred method for its simplicity and directness in providing the Cohen's d estimate.

The following R code first loads the necessary library, defines the observed plant growth data for Group 1 (Fertilizer #1) and Group 2 (Fertilizer #2), and then executes the calculation:

```
library(lsr)
```

```
#define plant growth values for each group
```

```
group1 <- c(8, 9, 11, 11, 12, 14, 15, 16, 16, 18, 20, 21)
group2 <- c(7, 9, 10, 10, 11, 11, 12, 14, 14, 16, 20, 23)

#calculate Cohen's d
cohensD(group1, group2)

0.2635333
```

The result, 0.2635333, represents the standardized difference between the two groups. Although this function is efficient, it provides only the point estimate of Cohen's d without additional information such as confidence intervals or predefined interpretation categories.

Method 2: Calculating Cohen's d Using the `effsize` Package

Alternatively, the **effsize** package offers a more comprehensive approach to calculating Cohen's d. The primary function, `cohen.d()`, is highly valuable because it not only provides the point estimate but also calculates the 95% confidence interval (CI) for the effect size. Including the CI is crucial for rigorous statistical reporting, as it quantifies the uncertainty surrounding the estimate.

The implementation is similar to the first method, requiring the loading of the library and the definition of the two independent data vectors:

library(effsize)

```
#define plant growth values for each group
group1 <- c(8, 9, 11, 11, 12, 14, 15, 16, 16, 18, 20, 21)
group2 <- c(7, 9, 10, 10, 11, 11, 12, 14, 14, 16, 20, 23)

#calculate Cohen's d
cohen.d(group1, group2)

Cohen's d

d estimate: 0.2635333 (small)
95 percent confidence interval:
lower upper
-0.5867889 1.1138555
```

As expected, both methods yield the identical **Cohen's d** estimate of **0.2635**. This consistency confirms the reliability of the calculation across different R packages. Crucially, the `cohen.d()` output automatically categorizes this estimate as "small," aligning perfectly with the established

rules of thumb (where 0.2 is the threshold for a small effect).

This result means that the average height of plants treated with Fertilizer #1 is **0.2635 standard deviations** greater than the average height of plants treated with Fertilizer #2. Although Fertilizer #1 resulted in slightly taller plants, the effect size is quite minimal. In practical terms, even if a t-test were to yield a statistically significant result, the actual, meaningful difference between the fertilizers is trivial.

Conclusion and Additional Resources

Calculating Cohen's d is a vital step in moving beyond the binary interpretation of p-values to truly understand the practical implications of experimental results. By standardizing the mean difference, [Cohen's d](#) provides a context-independent measure of the effect magnitude. The availability of efficient functions within [R](#), such as those found in the **lsr** and **effsize** packages, makes this calculation accessible for any researcher utilizing the platform for data analysis.

Remember that proper interpretation requires not only knowing the numerical value of d but also understanding the context of the study and the established conventions for small, medium, and large effects within the relevant scientific domain.

For further study and expanded knowledge on related topics, the following tutorials offer additional information on effect size and Cohen's d: