

How many bins should I use in my histogram?

Authored by
stats writer

December 12, 2025

RECOMMENDED CITATION

stats writer (2025). *How many bins should I use in my histogram?*. PSYCHOLOGICAL SCALES. Retrieved from <https://scales.arabpsychology.com/?p=107260>

Determining the optimal number of bins for a histogram is one of the most fundamental yet challenging tasks in data visualization. The choice critically impacts how the underlying distribution of the data set is presented, influencing statistical inference and pattern recognition. A poor choice can either mask vital features or introduce misleading noise, making the distribution appear jagged or overly smoothed.

The ideal number of bins is not fixed but is intrinsically linked to the inherent properties of your data, specifically its overall sample size (n) and its underlying shape, including volatility and skewness. Analysts often recommend a starting point, typically ranging between 5 and 15 bins, as a general rule of thumb for many common data sets. However, relying solely on intuition is insufficient for rigorous analysis; hence, statistical formulas provide objective methods for calculating an appropriate bin count, moving beyond simple visual estimation.

```
@import url('https://fonts.googleapis.com/css?family=Droid+Serif|Raleway');
```

```
.axis--y .domain {  
display: none;  
}
```

```
h1 {  
color: black;  
text-align: center;  
margin-top: 15px;  
margin-bottom: 0px;  
font-family: 'Raleway', sans-serif;  
}
```

```
h2 {  
color: black;  
font-size: 20px;  
text-align: center;  
margin-bottom: 15px;  
margin-top: 15px;  
font-family: 'Raleway', sans-serif;  
}
```

```
p {  
color: black;  
text-align: center;  
margin-bottom: 15px;
```

```
margin-top: 15px;  
font-family: 'Raleway', sans-serif;  
}
```

```
#words_intro {  
color: black;  
font-family: Raleway;  
max-width: 550px;  
margin: 25px auto;  
line-height: 1.75;  
}
```

```
#words_intro_center {  
text-align: center;  
color: black;  
font-family: Raleway;  
max-width: 550px;  
margin: 25px auto;  
line-height: 1.75;  
}
```

```
#words_outro {  
color: black;  
font-family: Raleway;  
max-width: 550px;  
margin: 25px auto;  
line-height: 1.75;  
}
```

```
#words {  
color: black;  
font-family: Raleway;  
max-width: 550px;  
margin: 25px auto;  
line-height: 1.75;  
padding-left: 100px;  
}
```

```
#calcTitle {  
text-align: center;  
font-size: 20px;
```

```
margin-bottom: 0px;
font-family: 'Raleway', serif;
}
```

```
#hr_top {
width: 30%;
margin-bottom: 0px;
margin-top: 10px;
border: none;
height: 2px;
color: black;
background-color: black;
}
```

```
#hr_bottom {
width: 30%;
margin-top: 15px;
border: none;
height: 2px;
color: black;
background-color: black;
}
```

```
.input_label_calc {
display: inline-block;
vertical-align: baseline;
width: 350px;
}
```

```
#button_calc {
border: 1px solid;
border-radius: 10px;
margin-top: 20px;
padding: 10px 10px;
cursor: pointer;
outline: none;
background-color: white;
color: black;
font-family: 'Work Sans', sans-serif;
border: 1px solid grey;
}
```

```
/* Green */  
}  
  
#button_calc:hover {  
background-color: #f6f6f6;  
border: 1px solid black;  
}  
  
.label_radio {  
text-align: center;  
}
```

The Core Function of Histograms and the Binning Process

A histogram serves as an indispensable visual tool used in descriptive statistics to represent the distribution of numerical data. It groups data into user-defined ranges, known as bins or classes, and then plots the frequency of observations falling into each range. The resulting bar-like chart reveals crucial characteristics such as the central tendency, spread, and overall shape of the distribution, making it invaluable for initial data exploration and hypothesis generation. Understanding the frequency distribution is essential for subsequent statistical modeling, whether you are examining error rates, biological measurements, or financial volatility.

The process of binning inherently involves a trade-off between reducing sampling variation and preserving the true shape of the population distribution. If too few bins are chosen, the visualization oversimplifies the data, potentially merging distinct peaks (multimodal data) into a single, misleading peak. Conversely, choosing too many bins can introduce excessive granularity, resulting in a jagged plot where each bar contains very few observations, making it difficult to discern underlying patterns or separating signal from noise.

Effective bin selection ensures that the chosen visualization accurately reflects the reality of the data set without sacrificing interpretability. Statisticians often emphasize that the goal is to choose a number of bins that minimizes the mean integrated squared error (MISE), which balances the bias (due to smoothing) and the variance (due to noise). Since calculating the true MISE is often impractical without knowing the underlying distribution, various established rules derived from statistical theory are employed to approximate this optimal bin width, providing objective starting points for visualization.

The Critical Impact of Bin Width on Data Interpretation

The width of the bins--which is mathematically inverse to the number of bins for a fixed data range--is the primary determinant of the resulting histogram's appearance. Wide bins, resulting from a

small number of bins, lead to a high degree of smoothing. While smoothing can be useful for very noisy data or small sample sizes, excessive smoothing can mask important structural details, such as outliers, gaps, or bimodality. For instance, if you are analyzing test scores, over-smoothing might hide the fact that the population consists of two distinct groups performing at high and low levels.

Conversely, narrow bins, resulting from a large number of bins, offer high resolution but suffer from high variance. Each bar becomes highly sensitive to individual data points, leading to a "picket-fence" appearance. In this scenario, the histogram might falsely suggest peaks and valleys that are merely artifacts of random sampling variability rather than true population characteristics. This high variance makes the plot unstable; adding or removing just a few data points could drastically change the visual distribution, making reliable interpretation challenging.

Therefore, the selection process is fundamentally iterative and requires balancing these visual and statistical concerns. Starting with an optimal, statistically derived number of bins provides a strong foundation. From that point, slight adjustments--perhaps increasing or decreasing the number by one or two--can be made to visually test how robust the key features (peaks, symmetry, tails) are, ensuring the final visual representation is both informative and stable. The ability to critically assess the visual outcome against the statistical recommendation separates competent data visualization from casual plotting.

Rule 1: Sturges' Formula--A Classic Starting Point

Sturges' Rule, introduced by Herbert Sturges in 1926, is perhaps the oldest and most widely adopted method for calculating the number of histogram bins. This formula assumes that the data follows an approximately normal distribution and is particularly suited for smaller to moderately sized data sets. It is based on the idea that the number of class intervals should be such that the frequency distribution resembles the binomial coefficients, which are related to the normal distribution approximation.

The primary advantage of Sturges' Rule lies in its simplicity and reliance solely on the total number of observations, or the sample size (n). However, this simplicity is also its major drawback. Because it ignores the variance and spread of the data--two factors critical to determining bin width--it often generates too few bins for large data sets or data that is highly skewed or widely dispersed. For data sets with tens of thousands of points, Sturges' Rule might fail to capture fine details, leading to over-smoothing and loss of crucial information about the distribution tails.

Despite these limitations, Sturges' Rule remains a valuable initial guide, especially when analyzing new or unknown data distributions, providing a quick and easy calculation for a preliminary histogram view. Many statistical software packages utilize this rule as a default setting due to its historical prevalence and ease of implementation. The formula mathematically connects the

number of bins (k) to the sample size (n) using a logarithmic scale, ensuring that the number of bins grows slowly as the data size increases.

The application of Sturges' Rule uses the following concise formula to determine the optimal number of bins required for a given data sample:

Number of bins = $\lceil \log_2 n + 1 \rceil$?

To apply Sturges' Rule to your data set, simply enter the total sample size (n) of the observations in the input box below and then click the "Calculate" button to obtain the suggested bin count.

n (the sample size)

Calculate

Calculated Number of Bins to Use (Sturges'): 7

```
function calc() {  
  //get input values  
  var n = document.getElementById('n').value*1;  
  
  //find number of bins  
  var bins = Math.ceil( Math.log2(n) - (-1) );  
  
  //output  
  document.getElementById('n_out').innerHTML = bins;  
}
```

Rule 2: Scott's Normal Reference Rule

Scott's Normal Reference Rule, developed by David W. Scott, represents a statistically more sophisticated approach compared to Sturges' Rule. This method is specifically designed to produce an optimally smooth histogram when the data is known to be distributed according to a normal reference distribution. Unlike Sturges' Rule, Scott's Rule determines the optimal bin width (h) first, and the number of bins is then derived from the data range divided by h .

The calculation for the bin width (h) under Scott's Rule involves the sample standard deviation (s) and the sample size (n), using the formula: $h = 3.5s / n^{(1/3)}$. This introduction of the standard deviation is crucial because it accounts for the spread and variability of the data, thereby adapting the bin width dynamically based on how scattered the observations are. For data with low variance, the bins will be narrower, providing more detail; for highly dispersed data, the bins will be wider, providing more generalization.

While Scott's Rule offers a statistically robust solution for normally distributed data, its performance

can degrade when applied to heavily skewed or multimodal data sets. Because it relies on the standard deviation, it is susceptible to the influence of outliers, which can inflate the standard deviation and result in overly wide bins, potentially leading to over-smoothing. Statisticians prefer Scott's Rule over Sturges' Rule when they have a reasonable assurance that the underlying process generating the data is approximately Gaussian (normal).

Rule 3: The Freedman-Diaconis Rule for Robustness

For data that is known to be non-normal, skewed, or contaminated by outliers, the Freedman-Diaconis Rule is often considered the superior choice for histogram binning. Developed by David Freedman and Persi Diaconis, this rule addresses the shortcomings of Scott's Rule by utilizing the interquartile range (IQR) instead of the standard deviation to determine the bin width (h). The IQR is the difference between the 75th percentile and the 25th percentile of the data, making it far more resistant to the influence of extreme values.

The formula for the bin width (h) is defined as: $h = 2 * IQR / n^{(1/3)}$. By incorporating the IQR, the rule effectively focuses on the central 50% of the data distribution, leading to a bin width calculation that is robust against outliers and heavily skewed tails. This robustness is especially valuable in fields like finance or environmental science, where data distributions are frequently non-normal and heavily influenced by rare, extreme events.

Although the Freedman-Diaconis Rule often yields a larger number of bins compared to Sturges' Rule, it is generally praised for producing histograms that reveal the true underlying structure of non-normal data with high fidelity. It is widely recommended as the most reliable default choice when no prior knowledge about the distribution of the data set exists, as it balances the need for detail with resistance to sampling noise effectively across diverse data types.

Practical Considerations and Visual Optimization

While statistical rules provide necessary objective benchmarks, the final decision on the number of bins must always incorporate human judgment and visual verification. Statistical rules are often based on asymptotic theory (the assumption of infinitely large data sets) or specific distribution assumptions (like normality). Therefore, a pragmatic approach involves calculating the recommended bin counts using two or three different methods (e.g., Sturges, Scott, and Freedman-Diaconis) and then comparing the resulting visualizations.

When comparing multiple histograms, ask critical questions: Do the major peaks identified remain consistent across different bin counts? Does the chosen bin count adequately reflect any known domain expertise? For example, if you know the underlying process should produce a bimodal distribution, but the Sturges' Rule result shows only one peak, you should opt for a method (like Freedman-Diaconis) that produces narrower bins to reveal that structure. Visual inspection is key

to ensuring that the histogram is fit for its intended purpose--communicating the distribution accurately.

Furthermore, attention must be paid to how the bin edges are defined. Poorly chosen bin boundaries can sometimes lead to visual artifacts, especially if data points cluster exactly on the boundary points. It is generally advisable to use bins that start and end on 'clean', easily interpretable values, even if it means slightly adjusting the width suggested by the formula. Modern statistical software handles boundary conditions automatically, typically using half-open intervals (e.g.,), ensuring every data point falls unambiguously into one bin.

Conclusion: Finding the Balance Between Precision and Clarity

The quest for the perfect number of bins in a histogram is fundamentally a balancing act between bias and variance. Too few bins lead to high bias (over-smoothing), obscuring important details. Too many bins lead to high variance (jaggedness), obscuring the true underlying pattern through excessive noise. The expert approach necessitates moving beyond arbitrary selection and employing statistically grounded techniques.

For analysts, the recommended strategy involves using a suite of rules based on the characteristics of the data set. If the data is small or assumed normal, Sturges' Rule or Scott's Rule provide excellent starting points. However, for large data sets, non-normal distributions, or data containing significant outliers, the Freedman-Diaconis Rule is often the most robust choice, providing a consistent and reliable bin width calculation that minimizes visual distortion caused by extreme values.

Ultimately, while statistical formulas provide the scientific foundation for initial binning decisions, the final iteration should be visually optimized. Successful data visualization requires the editor to ensure the chosen bin count maximizes clarity and accurately conveys the data's story to the intended audience, ensuring that the derived statistical insights are both valid and easily digestible.