



The Bag-and-Whisker Plot: A New Bagplot for Bivariate Data

Shenghao Qin^a, Bowen Gang^a, Tiejun Tong^b, and Hengjian Cui^c

^aDepartment of Statistics and Data Science, Fudan University, Shanghai, China; ^bDepartment of Mathematics, Hong Kong Baptist University, Hong Kong, China; ^cSchool of Mathematical Sciences, Capital Normal University, Beijing, China

ABSTRACT

The bagplot, also known as the “bag-and-bolster plot”, is a notable extension of the boxplot from univariate to bivariate data. Although widely used, its practical application is hindered by two key limitations: the fixed inflation factor for outlier detection that does not adapt to the sample size, and the unstable convex hull used to visualize its fence. In this paper, we propose a new bagplot, namely the “bag-and-whisker plot”, as an improvement method to address these limitations. Our framework recasts outlier detection as a multiple testing problem, yielding a data-adaptive fence that controls statistical error rates and enhances the reliability of outlier identification. To further resolve graphical instability, we introduce a refined visualization that abandons the convex hull (the bolster) with a direct rendering of the statistical fence, complemented by granular whiskers that effectively illustrate the data’s spread. Extensive simulations and real-world data analyses demonstrate that our new bagplot exhibits superior adaptivity and robustness compared to the existing standard, and thus can be highly recommended for practical use. To increase the visibility of the work, a user-friendly R package named `BagWhiskerPlot` has been made publicly available on CRAN.

ARTICLE HISTORY

Received December 2025
Accepted June 2026

KEYWORDS

Data depth; Data visualization; Exploratory data analysis; Multiple testing; Outlier detection; Robust statistics

1. Introduction

Tukey’s box-and-whisker plot is a foundational tool in exploratory data analysis, celebrated for its robust and concise summary of a dataset’s key features (Tukey 1977). A defining element is its rule for flagging potential outliers, which designates any point beyond 1.5 times the interquartile range (IQR) from the central box as suspect. This $1.5 \times \text{IQR}$ heuristic, while simple and effective for its time, was not derived from formal statistical principles and its performance is known to degrade in large samples, often flagging an excessive number of points. For decades, numerous modifications to the inflation factor of 1.5 were proposed, representing significant efforts to refine the heuristic approach for outlier detection (Hoaglin, Iglewicz, and Tukey 1986; Hoaglin and Iglewicz 1987; Sim, Gan, and Chang 2005). Recent advancements have addressed this by introducing a new generation of boxplots with data-adaptive fences. These newer methods, such as the Chauvenet-type, Holm-type, and BH-type boxplots, adjust their outlier detection thresholds based on the sample size, offering a more statistically principled and reliable framework for univariate data analysis (Gang, Lin, and Tong 2026; Lin, Zhang, and Tong 2026; Tong et al. 2026).

Just as the univariate boxplot evolved, so too did efforts to generalize it to bivariate data. This task, however, is fundamentally more complex due to the lack of a natural ordering of points in two dimensions (Wickham and Stryjewski 2011). Early attempts, such as the rangefinder plot (Beckett and Gould 1987), were simple constructions based on marginal statistics, effectively treating the variables as independent. The 1990s witnessed a proliferation of more sophisticated proposals, ranging from model-based elliptical plots (Goldberg and Iglewicz 1992) to

non-parametric approaches based on convex hull peeling (Zani, Riani, and Corbellini 1998).

The key conceptual advance that led to a widely adopted solution was the use of data depth. The bagplot, introduced by Rousseeuw, Ruts, and Tukey (1999), was built specifically on the concept of halfspace depth (Tukey 1975), while the broader utility of depth functions for robust multivariate graphics and inference was concurrently formalized in the influential work of Liu, Parelius, and Singh (1999). The bagplot’s construction is analogous to the univariate boxplot: a central “bag” contains the 50% of data points with the maximum depth, and an outer “fence” is used for outlier detection.

The bagplot has since established itself as the *de facto* standard for 2D data visualization, a status attested to by over 850 citations in Google Scholar as of May 2026 and its permeation into diverse applied disciplines. This widespread adoption underscores a critical and pervasive demand for tools that can intuitively reveal bivariate relationships, non-linear associations, and outlier patterns. In ecological and environmental science, for instance, Schirpke et al. (2019) employed bagplots to unravel tradeoffs and synergies among ecosystem services, leveraging the tool to visualize data dispersion that linear methods overlook. In physical anthropology, O’Connell et al. (2012) utilized the method to clarify isotopic offsets in paleodietary studies, relying on the visualization of data concentration to make accurate interpretations. The tool’s cross-disciplinary relevance is further exemplified in environmental epidemiology, where Damialis et al. (2021) used bagplots to confirm correlations between airborne pollen and SARS-CoV-2 infection rates across 31 countries, and in neuroscience, where Baliki et al.

(2011) applied them to classify brain morphological signatures in chronic pain.

However, the contrast between this extensive cross-disciplinary demand and the bagplot's stagnant methodological development reveals a significant opportunity. While the bagplot is the prevailing standard, its original formulation relies on heuristics that have not kept pace with modern statistical rigor. Specifically, it suffers from two critical limitations. First, its outlier detection rule uses a fixed inflation factor (typically 3) that fails to adapt to the sample size or dimension—a static heuristic similar to the univariate $1.5 \times \text{IQR}$ rule that modern methods have since superseded. Second, its fence visualization, known as the “loop”, is defined as the convex hull of non-outlying points. Rousseeuw, Ruts, and Tukey (1999) noted that this loop effectively plays the role of a bolster, suggesting that their method could be more accurately termed a “bag-and-bolster plot” rather than a box-and-whisker equivalent. This construction is empirically unstable and often visually misleading, as the hull's shape is highly sensitive to the specific locations of peripheral points, creating a disconnect between the statistical rule and its graphical representation.

In this paper, we enhance the bagplot by addressing these statistical and visual limitations, effectively transforming the “bag-and-bolster plot” into a true “bag-and-whisker plot”. We first incorporate the principles of multiple testing, replacing the fixed inflation factor with a data-adaptive method. This allows for the construction of bagplots with fences that are inspired by controlling various error rates. To resolve the visual instability of the bolster (loop), we propose a new graphical approach. Our method visualizes a theoretical fence, which is a stable contour based on the chosen statistical threshold, alongside a new implementation of granular whiskers. This dual representation provides a more robust interpretation by clearly distinguishing the underlying statistical boundary from the empirical shape of the data, creating a more reliable and informative tool for bivariate exploration.

2. The Bagplot by Rousseeuw, Ruts, and Tukey (1999)

A bivariate generalization of the univariate boxplot, the bagplot, developed by Rousseeuw, Ruts, and Tukey (1999), serves as a tool for visualizing essential characteristics of a dataset, including its location, spread, correlation, and skewness. Its construction is founded on the concept of halfspace depth (Tukey 1975), a multivariate extension of the univariate notion of rank.

2.1. Construction of the Original Bagplot

For a given bivariate dataset $\mathcal{Z} = \{z_1, \dots, z_n\}$, where $z_i \in \mathbb{R}^2$, the halfspace depth of a point $\theta \in \mathbb{R}^2$ is the minimum number of data points contained in any closed half-plane with a boundary passing through θ . The bagplot is constructed through the following steps:

- (1) **The Depth Median:** The point with the maximum halfspace depth is identified as the depth median, denoted μ_{depth} . This serves as the bivariate analogue of the univariate median, providing a robust estimate of the data's center.
- (2) **The Bag:** A central polygon, called the “bag”, is constructed to contain about 50% of observations with the maximum depth. The bag is the two-dimensional counterpart to the IQR and robustly captures the data's spread and shape.
- (3) **The Fence:** A theoretical outer boundary, called the “fence”, is determined by magnifying the bag by a factor of 3 relative to the depth median. Any point lying outside this fence is flagged as a potential outlier.
- (4) **The Loop:** The region between the bag and the fence contains the remaining non-outlying points. The bagplot visualization does not typically render the fence directly. Instead, the boundary of the non-outlier region is depicted by the convex hull of all points contained within the loop (including the bag itself).

2.2. Limitations of the Original Bagplot

While the bagplot is a powerful exploratory tool, its original formulation presents two key limitations that can affect its statistical and visual interpretation.

2.2.1. Fixed Inflation Factor

The use of 3 as the inflation factor for the bagplot is a fixed heuristic. This can be contrasted with the univariate boxplot, where the $1.5 \times \text{IQR}$ rule for the two fences is a well-established convention. A notable property of this univariate rule is that, for data drawn from a normal distribution, it flags approximately 0.7% of observations as outliers. Under a bivariate normal distribution, however, the bagplot's inflation factor of 3 results in only 0.195% of observations being outside of the fence. This proportion is substantially more conservative than the 0.7% benchmark, highlighting a clear inconsistency between the methods.


The challenge of selecting an appropriate inflation factor has been previously noted in the literature. For instance, Zani, Riani, and Corbellini (1998), in their work on bivariate boxplots based on convex hull peeling and B-spline smoothing, derived an inflation factor of 1.58 specifically to achieve an outlier percentage of approximately 1% under bivariate normality. While this represents a principled attempt to standardize the outlier rate, the resulting factor is still a fixed constant. As Zani, Riani, and Corbellini (1998) further observed, for any fixed factor, the proportion of flagged points from a normal sample is a decreasing function of the sample size n . This underscores the fundamental limitation of any fixed-factor approach. Furthermore, this discrepancy is compounded by dimensionality. The need for an inflation factor to decrease as the data's dimension increases, in order to maintain a consistent outlier proportion of 0.7%, demonstrates the inadequacy of a single fixed value. The inflation factors required to achieve this 0.7% proportion, up to 10 dimensions, are shown in Table 1.

2.2.2. Fence Visualization

The second limitation of the bagplot arises from a mismatch between its statistical definition and its graphical implementation. The outlier detection rule is based on a theoretical fence, which is a simple magnification of the central bag. However, the visualization does not show this theoretical boundary. Instead, it displays the convex hull of all points not flagged as outliers. This

Table 1. Inflation factors required to maintain a constant outlier proportion of 0.7% (the benchmark from a standard univariate boxplot) as dimensionality grows.

Dimension	1	2	3	4	5	6	7	8	9	10
Inflation Factor	4.00	2.68	2.26	2.05	1.91	1.82	1.75	1.69	1.65	1.61

**Figure 1.** Visual limitations of the “bolster” (convex hull) representation in the classic bag-and-bolster plot by Rousseeuw, Ruts, and Tukey (1999). Left: the star (★) lies strictly inside the theoretical fence (dashed contour), yet it falls outside the visualized convex hull (solid polygon), leading to incorrect outlier inference. Right: the shape of the convex changes drastically with the inclusion of a single point at the star’s position, thus altering the entire shape of the visualized boundary.

practice creates a visually unstable and potentially misleading representation for two reasons.

First, it creates a visual paradox not present in the univariate boxplot. In a standard boxplot, the fence is simply the end of the whisker. Any data point lying between the median and the whisker is unambiguously inside the fence. In the bagplot, a point can be well within the theoretical fence but lie outside the visualized convex hull, as illustrated in the left panel of Figure 1. The dashed line represents the theoretical fence boundary obtained by magnifying the bag. The solid polygon is the convex hull of the loop points, which is what is typically visualized. A point at the location of the star (★) is inside the theoretical fence but outside the visualized convex hull. An observer would incorrectly infer that a point at the star’s location would be flagged as an outlier. In reality, the point’s data depth may be no smaller than that of the vertices defining the convex hull; according to the statistical rule, it is not an outlier.

Second, this visualization is empirically unstable. The shape of the convex hull is highly sensitive to the exact location of its defining vertices. The addition, removal, or slight perturbation of a single data point near the periphery can dramatically and non-locally alter the shape of the visualized boundary. As illustrated in the right panel of Figure 1, if there is a data point at the star’s position, the shape of the convex hull will change drastically. This instability means the graph is not a robust representation of the data’s underlying distribution. This disconnect between the formal rule and the visual evidence undermines the plot’s reliability. The graphic should be a faithful representation of the statistical procedure, a principle that the conventional bagplot violates.

3. New Bagplot: The Bag-and-Whisker Plot

In this section, we propose a comprehensive enhancement to the bagplot that addresses the limitations of both its fixed inflation factor and its fence visualization. Our methodology con-

sists of two primary contributions. First, we introduce a formal multiple testing framework to construct a data-adaptive fence. Second, we redesign the visualization of spread and skewness by replacing the unstable “bolster” with explicit whiskers, effectively transforming the original “bag-and-bolster” display into a genuine bivariate bag-and-whisker plot.

3.1. A Multiple Testing Framework for Adaptive Inflation Factor

We recast the bagplot’s outlier detection rule as a formal multiple testing procedure, analogous to recent work unifying the univariate boxplot (Gang, Lin, and Tong 2026). We treat the classification of each data point as a distinct hypothesis test. For each point z_i in a dataset $\mathcal{Z} \subset \mathbb{R}^2$, the null hypothesis, H_{0i} , states that the point is drawn from the same underlying distribution as the bulk of the data. Flagging z_i as an outlier is thus equivalent to rejecting H_{0i} .

The original bagplot, with its inflation factor of 3, corresponds to an unadjusted testing procedure that rejects H_{0i} if its associated p -value (under the null hypothesis that the majority of the data follow a bivariate normal distribution) is less than approximately 0.00195. This fixed threshold presents a dual limitation. In small samples, this cutoff is substantially more conservative than the 0.007 benchmark of the standard univariate boxplot, resulting in overly strict control that may fail to detect genuine outliers. Conversely, in large samples, the lack of adjustment for multiplicity causes the expected number of false positives to grow linearly with the sample size, potentially resulting in an excess of flagged outliers. Our solution is to replace this fixed-cutoff rule with a general and adaptive p -value-based pipeline. This approach allows for the application of standard multiple testing procedures to control error rates such as the Per-Family Error Rate (PFER), the Family-Wise Error Rate (FWER), or the False Discovery Rate (FDR) by Benjamini and Hochberg (1995), resulting in a bagplot with a statistically principled fence.

Our methodology is built on the pragmatic assumption that the central body of the data follows a bivariate normal distribution with mean μ and covariance matrix Σ . This approach is rooted in Winsor's principle that "all distributions are normal in the middle" (Tukey 1960), providing a robust foundation for identifying points that deviate from this central tendency. Our pipeline proceeds in the following steps:

- (1) **Robust Parameter Estimation.** The first step of our pipeline requires robust estimation of the parameters (μ, Σ) that characterize the central body of the data under a bivariate normal assumption. To ensure these estimates are resistant to observations distant from the main data cloud, we employ distinct methods for the location and scatter parameters. Specifically, we estimate the location vector μ with the depth median. In cases where the point achieving the maximum halfspace depth is not unique, we define the estimator $\hat{\mu}$ as the coordinate-wise median of the set of deepest points. For the covariance matrix Σ , we employ the highly robust Minimum Covariance Determinant (MCD) estimator (Rousseeuw 1984). The resulting estimates are denoted by $\hat{\mu}$ and $\hat{\Sigma}$, respectively.
- (2) **Bag Construction.** Following the standard definition in Rousseeuw, Ruts, and Tukey (1999), we construct the central "bag" to contain about 50% of observations with the maximum halfspace depth.
- (3) **p-value Calculation.** To obtain a p -value for each point z_i , we first measure its standardized distance from the robustly estimated center using the squared Mahalanobis distance:

$$d_i^2 = (z_i - \hat{\mu})^\top \hat{\Sigma}^{-1} (z_i - \hat{\mu}). \quad (1)$$

Under the null hypothesis H_{0i} that an observation z_i is not an outlier, its squared Mahalanobis distance can be compared to a χ^2_2 distribution. This reference is based on the standard result that, assuming the bulk of the data follows a bivariate normal distribution, the squared robust Mahalanobis distance converges asymptotically to a χ^2_2 distribution as $n \rightarrow \infty$ (Hardin and Rocke 2005). The resulting p -value is the upper-tail probability:

$$p_i = P(\chi^2_2 \geq d_i^2). \quad (2)$$

- (4) **Multiple Testing Adjustment.** With the set of p -values $\{p_1, \dots, p_n\}$, we apply a standard multiple testing procedure to control a desired error rate (e.g., FWER, FDR, PFER) at a specified level q . This yields a data-dependent significance threshold, t_{adj} . Any hypothesis H_{0i} with $p_i \leq t_{\text{adj}}$ will be rejected. We then convert this threshold back to the critical value for the squared Mahalanobis distance:

$$d_{\text{adj}}^2 = F_{\chi^2_2}^{-1}(1 - t_{\text{adj}}), \quad (3)$$

where $F_{\chi^2_2}^{-1}$ is the quantile function of the χ^2_2 distribution.

- (5) **Data-Adaptive Inflation Factor.** Finally, we derive a data-adaptive inflation factor, λ , that ensures the visualized fence is both statistically principled and geometrically consistent with the multiple testing results. This is achieved through a two-stage process.

First, we define a *statistical* inflation factor, λ_{stat} , that aligns the fence with the statistical threshold d_{adj}^2 from Equation (3). Since the bag's boundary represents the 50th percentile of data depth, it should correspond approximately to the median of the squared Mahalanobis distances. We thus define λ_{stat} as:

$$\lambda_{\text{stat}} = \sqrt{\frac{d_{\text{adj}}^2}{\text{median}(d_1^2, \dots, d_n^2)}}.$$

This factor scales the bag such that its boundary would, under ideal elliptical symmetry, align with the multiple testing cutoff.

However, the bag is a data-dependent polygon, not a perfect ellipse corresponding to an isocontour of the Mahalanobis distance. Consequently, magnifying the bag by λ_{stat} does not guarantee that all non-rejected points will lie inside the resulting fence. To resolve this potential contradiction between the statistical test and the visual representation, we introduce a second, *data-driven* inflation factor, λ_{data} , which is defined as the minimum magnification required to geometrically enclose all points z_i for which the null hypothesis H_{0i} was not rejected. More precisely, let b_i be the intersection of the bag's boundary with the ray originating from $\hat{\mu}$ and passing through z_i . We define λ_{data} as the maximum ratio of distances for all non-rejected points:

$$\lambda_{\text{data}} = \max_{i \in \mathcal{S}} \frac{\|z_i - \hat{\mu}\|}{\|b_i - \hat{\mu}\|},$$

where $\mathcal{S}$ is the set of indices for points where H_{0i} is not rejected. The final data-adaptive inflation factor λ is chosen to satisfy both the statistical criterion and the geometric constraint:

$$\lambda = \max(\lambda_{\text{stat}}, \lambda_{\text{data}}). \quad (4)$$

This factor is then used to construct the outer fence by magnifying the bag. A few remarks are in order.

Remark 1. A key strength of the proposed framework is its modularity. While our implementation is instantiated with the halfspace depth median and the MCD estimator, these are not essential components. The pipeline is compatible with any robust estimators for the location and scatter parameters. Alternative notions of data depth (e.g., Liu, Parelius, and Singh 1999) could be readily substituted for the location estimate. The MCD estimator is one of many powerful tools for robust covariance estimation, other possibilities include using the depth weighted scatter estimators by Zuo and Cui (2005) or the T-type estimator by Tian, Zhang, and Cui (2024). For a comprehensive treatment of robust statistics, readers may refer to the foundational texts such as Hampel et al. (2005), Huber and Ronchetti (2009), and Maronna et al. (2019).

Remark 2. We note that for finite samples, the distribution of squared robust Mahalanobis distances from an MCD estimator can be appropriately approximated by a scaled F -distribution when the data are from a normal distribution as shown by Hardin and Rocke (2005). However, this approximation requires Monte Carlo calibration of the scaling parameters, which introduces substantial computational overhead. Given the computational cost we adopt the asymptotic χ^2 reference as the default.

For illustrative purposes, Sections 5 and 6 include bagplots constructed under both p -value constructions. Our accompanying R package `BagWhiskerPlot` implements both distributions for p -value constructions, allowing users to select the scaled F -distribution if they prefer.

Remark 3. The necessity of our two-stage approach for deriving the inflation factor, $\lambda = \max(\lambda_{\text{stat}}, \lambda_{\text{data}})$, is a feature unique to the multivariate setting. In the univariate setting (Gang, Lin, and Tong 2026), a single-stage factor is sufficient. This is because a fundamental consistency exists in one dimension: the robust parameters estimates, the inner box, and the fences are all derived directly from the same set of sample quantiles. The statistical rule and the graphical construction are perfectly aligned. This direct alignment breaks down in two dimensions. Our statistical rule for outlier detection is based on the squared Mahalanobis distance, which defines perfectly elliptical contours of probability. The graphical construct, however, is the bag, which is a data-dependent convex polygon. There is no guarantee that this polygon is a perfect ellipse. Consequently, a simple magnification of the polygonal bag (the statistical target set by λ_{stat}) may not produce a fence that is geometrically consistent with the elliptical boundary defined by the multiple testing procedure. A point with a small Mahalanobis distance could lie in a direction where the bag has a “flat side”, causing it to fall outside the magnified polygon. The data-driven factor, λ_{data} , serves as an essential geometric correction, ensuring that the final visualized fence is a faithful representation of the statistical test’s outcome.

3.2. A Direct and Stable Visualization

To address the visual instability and misinterpretation inherent in the convex hull representation (see Section 2.2.2), we replace the “bolster” with explicit whiskers and a direct fence visualization. This design abandons the volatile loop to faithfully map the data’s tail structure, preserving details of skewness and correlation that a polygon often obscures. Our approach is defined by two principles.

- (1) **Direct Fence Visualization.** The primary change is to draw the fence explicitly. This fence is the exact contour obtained by magnifying the central bag by the data-adaptive inflation factor λ in (4). This modification ensures a direct, one-to-one correspondence between the statistical rule and the graphical boundary: a point is flagged as an outlier if and only if it lies outside this line.
- (2) **A Refined Whisker Implementation for Visual Clarity.** To better align with the spirit of Tukey’s original box-and-whisker plot and enhance visual interpretation, we propose a refined implementation of bivariate whiskers that integrates their design and graphical composition. Our approach draws individual lines from the bag’s boundary to *each* outer point (point that is inside the fence but outside of the bag) in the surrounding region.

Crucially, these whiskers are rendered with two key graphical properties. First, they are placed on a background layer, ensuring that all data points are plotted on top and remain fully visible.




Figure 2. Scatterplot and halfspace depth illustration for the toy example. The star-shaped (red) point at $z_1 = (5, 7)$ marks the depth median. z_2 to z_4 indicate points with larger depth (the bag’s core), z_5 to z_7 are outer points, and the point $z_8 = (19, 20)$ is the extreme outlier located farthest from the majority of the data. Shaded halfspaces in the plot show example halfplanes used to compute the halfspace depth of z_1 , z_2 , and z_5 .

Second, each whisker features a transparent gradient along its length, being most transparent at its origin on the bag’s boundary and becoming progressively more solid as it approaches the data point. This unique gradient serves a dual purpose: it prevents visual clutter near the dense central bag while simultaneously drawing the eye outward, visually emphasizing the most extreme observations that define the data’s spread. Together, these design choices create a clear visual hierarchy: the solid bag represents the data’s robust core, the gradient whiskers illustrate the extent of the non-outlying data, and the distinct fence marks the formal decision boundary.

4. A Step-by-Step Toy Example

To illustrate the computation of the new bagplot, we present a simple toy example. Consider a small dataset of $n = 8$ observations as follows:

$$\begin{aligned} Z &= \{z_i, i = 1, \dots, 8\} \\ &= \{(7, 5), (7, 7), (9, 4), (5, 4), (14, 9), (0, 9), (7, -3), (19, 20)\}. \end{aligned}$$

The scatterplot of these points is shown in Figure 2. In this dataset, the observation z_8 is well-separated from the bulk of the data. We will walk through the pipeline to see how the procedure identifies outliers and constructs the bagplot.

Step 1: Robust Parameter Estimation. First, we estimate the location using the depth median. For each point z_i , we compute its halfspace depth. Consider the point z_5 as an illustrative example. The halfplane defined by $y \leq 2x - 19$ contains only the single point z_5 . Since this count of 1 is minimal across all possible halfplanes through z_5 , its halfspace depth is 1. For the point z_2 , any closed halfplane passing through it contains at least two points, for example, the half plane defined by $y \geq 7(x - 1)/6$ contains z_2 and z_6 . Thus,

its depth is 2. Similarly, we can conclude that z_6, z_7 and z_8 have the depth of 1, and z_3 and z_4 have the depth of 2. The depth median is z_1 , which attains the maximum depth of 3. To see this, the halfplane defined by $y \leq (-5x + 80)/9$ passes through z_1 but excludes z_3 and z_6 . The MCD estimator yields the covariance matrix $\begin{pmatrix} 53/3 & 0 \\ 0 & 17 \end{pmatrix}$.

Step 2: Bag Construction. Next, we construct the central bag. In this toy example, we identify the set of the four deepest observations: z_1, z_2, z_3 , and z_4 . The bag is defined as the convex hull of these points, which forms the triangle depicted in Figure 2.

Step 3: p -value Calculation. Using the robust parameter estimates from step 1, we calculate the squared Mahalanobis distance, d_i^2 , for each of the 8 data points via (1). We then convert these distances into upper-tail probabilities using the χ_2^2 reference distribution, as defined in (2), to obtain a p -value for each point. The resulting squared Mahalanobis distances and their corresponding p -values are:

$$d^2 = \{0.00, 0.24, 0.29, 0.29, 3.71, 3.71, 3.76, 21.39\},$$

$$p\text{-value} = \{1.00, 0.89, 0.87, 0.87, 0.16, 0.16, 0.15, 2.27 \times 10^{-5}\}.$$

The smallest p -value corresponds to the point z_8 , providing strong evidence that it does not follow the central data distribution.

Step 4: Multiple Testing Adjustment. With the set of p -values, we apply different multiple testing procedures to identify which points should be flagged as outliers. We consider three common error-control strategies:

- **FWER Control:** Using the Holm-Bonferroni method (Holm 1979) at a level of $q = 0.1$, the procedure rejects any hypothesis corresponding to a p -value less than or equal to the single significant result, yielding a threshold of $t_{\text{adj}} \approx 2.27 \times 10^{-5}$.
- **FDR Control:** Applying the BH procedure (Benjamini and Hochberg 1995) at $q = 0.1$ also results in rejecting only the hypothesis for the most extreme point, giving an identical threshold of $t_{\text{adj}} \approx 2.27 \times 10^{-5}$.
- **PFER Control:** To control the Per-Family Error Rate at a target of $q = 0.5$ expected false positives, we reject any hypothesis where the p -value is less than or equal to q/n . This gives a threshold of $t_{\text{adj}} = 0.5/8 = 0.0625$.

These significance thresholds are then converted back to the scale of squared Mahalanobis distance using (3). This results in a critical value of $d_{\text{adj}}^2 \approx 21.39$ for both FWER and FDR control, and a more liberal critical value of $d_{\text{adj}}^2 \approx 5.55$ for PFER control. In this example, the rejection set is identical across all three procedures, consisting solely of the extreme observation z_8 .

Step 5: Data-Adaptive Inflation Factor. We now derive the inflation factor λ using our two-stage approach. First, we calculate λ_{stat} , based on the critical distances from step 4. The median of the squared Mahalanobis distances is $(0.29 + 3.71)/2 = 2$.

- For FWER and FDR control, $\lambda_{\text{stat}} = \sqrt{21.39/2} \approx 3.27$.
- For PFER control, $\lambda_{\text{stat}} = \sqrt{5.55/2} \approx 1.67$.

Next, we determine the *data-driven* inflation factor, λ_{data} . Given that all three procedures rejected only the point z_8 (the extreme outlier), the fence must be sufficiently large to geometrically enclose all remaining observations. The most extreme non-rejected point is z_7 ; thus, we calculate that a minimum inflation factor of $\lambda_{\text{data}} = 8$ is required to position this point on the boundary of the magnified bag. This calculation is derived from the ratio of the distance between the depth median and z_7 to the distance between the median and the bag's boundary along the ray connecting them. Specifically, the distance to z_7 is $|5 - (-3)| = 8$, while the distance to the boundary intersection point $(4, 7)$ is $|5 - 4| = 1$. Finally, to satisfy both statistical and geometric criteria, we set $\lambda = \max(\lambda_{\text{stat}}, \lambda_{\text{data}})$.

- For FWER and FDR, $\lambda = \max(3.27, 8) = 8$.
- For PFER, $\lambda = \max(1.67, 8) = 8$.

The resulting adaptive bagplots are constructed using these final inflation factors and are shown in Figure 3.

As intended, under all three procedures the point z_8 is flagged as an outlier. This result aligns with our visual intuition and showcases the method's ability to produce a principled and coherent graphical summary. It is instructive to contrast this outcome with the original bagplot proposed by Rousseeuw, Ruts, and Tukey (1999). Applying the conventional fixed inflation factor of 3 to this dataset excludes four points (z_5, z_6, z_7 and z_8) from the fence. Consequently, the set of non-outlying points is reduced strictly to the central bag itself, causing the visualized "bolster" to collapse onto the bag's boundary.

The toy example above allows for an exact, manual verification of the methodology. However, for general datasets where n is large, exact computation becomes computationally intensive. Our implementation therefore relies on established algorithmic approximations to ensure efficiency without sacrificing robust performance. The following remarks clarify the specific computational routines and software constraints relevant to the proposed method.

Remark 4. In the toy example with so few points, we were able to compute the halfspace depth by hand, but it is obvious that such a manual approach is not practical when n is large. Rousseeuw and Ruts (1996) proposed an efficient algorithm that can compute the halfspace depth of n points in $O(n \log n)$ time. In our implementation, we adopt the same approximation method used in the `aplpack` package in R, which iterates over a fixed set of directions (e.g., 180 or 360) to find the minimum number of points in the halfspaces on either side.

Remark 5. In finite datasets, the point achieving the maximum halfspace depth is not always unique. Consistent with the original definition provided by Rousseeuw, Ruts, and Tukey (1999), when the set of deepest points contains more than one observation, we define the depth median as the center of gravity of this subset. Computationally, this is obtained by calculating the arithmetic mean of the coordinate vectors of all observations that attain the global maximum depth.

Remark 6. The MCD estimator (Rousseeuw 1984) is defined by the mean and covariance of the subset of h observations (typically $h \approx n/2$) whose covariance matrix possesses the smallest




Figure 3. On the left panel is the original bagplot from Rousseeuw, Ruts, and Tukey (1999). On the right panel is the new bag-and-whisker plot. Since the inflation factor derived from the three error rates are the same, the resulting bag-and-whisker plots are identical.




Figure 4. Bagplots for a bivariate normal mixture with a weakly separated secondary component. Panel (a): original bagplot (Rousseeuw, Ruts, and Tukey 1999) with fixed inflation factor and convex-hull boundary. Panels (b), (d), (f): proposed bag-and-whisker plots using the χ^2_2 reference with FWER, FDR, and PFER control, respectively. Panels (c), (e), (g): corresponding plots using the scaled F approximation. Key features: (i) smooth elliptical fences (dashed) that reflect the statistical threshold and avoid convex-hull artifacts; (ii) gradient whiskers extending from the bag to peripheral points, indicating the direction of the contaminating cluster; (iii) adaptive fence radii across error-control procedures, reflecting varying tolerance for false discoveries. The F -based fences are slightly more conservative than their χ^2 counterparts.

possible determinant. To ensure statistical consistency under the multivariate normal model, this raw covariance estimate is multiplied by a specific scaling factor. Since the exact computation of the MCD is computationally intensive even for relative

small n as it requires the evaluation of $\binom{n}{h}$ subsets. Rousseeuw and Driessen (1999) proposed the FAST-MCD algorithm to efficiently approximate the solution. In our toy example and subsequent implementation, we utilize this algorithm, including




Figure 5. Bagplots for a correlated bivariate normal mixture ($\rho = 0.3$) with a minor contaminating component. Panel (a) displays the original bagplot, where the convex hull boundary is jagged and misaligned with the underlying elliptical geometry. Panels (b), (d), (f) show the proposed method using the χ^2_2 reference distribution, while panels (c), (e), (g) use the scaled F reference. Note the orientation of the fences relative to the data cloud. Unlike the fixed rule in panel (a), the proposed fences are smooth and elongated, accurately reflecting the covariance structure estimated by the MCD. This alignment ensures that outlier detection respects the data's correlation. The whiskers further clarify the direction of the minor component along the major axis of correlation.

the necessary consistency scaling, via the `cov.mcd` function available in the `MASS` package in R.

Remark 7. In this toy example, the subset of observations with a halfspace depth ≥ 2 contains exactly $n/2 = 4$ points, allowing the bag to be defined simply as the convex hull of this set. In general practice, however, due to the discrete nature of halfspace depth, there is rarely a depth contour that strictly encloses exactly 50% of the data. Consequently, constructing the bag requires a more sophisticated geometric interpolation strategy to approximate the 50% central region. To handle these non-trivial cases robustly, our implementation relies on the internal routines of the `aplpack` package in R to compute the bag.

5. Simulated Data Analysis

In this section, we use simulated data to demonstrate the performance of the proposed bag-and-whisker plot under controlled conditions where the ground truth is known. To provide a

clear benchmark, we compare our method against the widely used implementation of the original bagplot in the R package `aplpack`, which represents the conventional approach of a fixed inflation factor and a convex hull visualization. We showcase three variants of our adaptive bagplot, corresponding to the control of FWER, FDR, and PFER, respectively. For FWER and FDR we set the nominal levels to 0.1. For PFER we set the nominal level to 0.5 which corresponds to Chauvenet's criterion (Chauvenet 1863; Lin, Zhang, and Tong 2026). Each of the following three examples is designed to test a specific aspect of the method: its statistical adaptivity to data contamination, its robustness to correlation structure, and its visual stability in the presence of heavy tails.

Example 1: Statistical Adaptivity in Independent Normal Mixture. Our first example tests the methods on a subtle normal mixture, designed to assess how the fences adapt to a small, contaminated component. We generate $n = 5000$ points from the following model:

$$\theta_i \stackrel{iid}{\sim} \text{Bernoulli}(0.05), \quad \vec{X}_i | \theta_i \stackrel{ind}{\sim} (1 - \theta_i) \mathcal{N}(\vec{\mu}_0, \mathbf{I}_2) + \theta_i \mathcal{N}(\vec{\mu}_1, \mathbf{I}_2),$$




Figure 6. Bagplots for simulated heavy-tailed log-normal data without true contamination, illustrating visual stability in the presence of skewness and heavy tails. Panel (a) shows the original bagplot, where the convex hull produces an irregular, jagged boundary driven by peripheral points. Panels (b), (d), (f) (χ^2_2 reference) and panels (c), (e), (g) (scaled F reference) yield smooth and stable fences that faithfully reflect the underlying distribution. The gradient whiskers clarify directional spread without visual clutter.

where $\vec{\mu}_0 = (100, 300)^\top$, $\vec{\mu}_1 = (103, 298)^\top$ and $\mathbf{I}_2$ is the 2×2 identity matrix. As shown in Figure 4, the `aplpack` implementation (panel (a)) produces a boundary based on a fixed rule and a convex hull visualization that can be sensitive to the specific sample. In contrast, our bagplots (panels (b)-(g)) demonstrate clear advantages. The fences are smooth and elliptical, reflecting the underlying normality of the components. The different error control levels provide a nuanced view: the FWER fence is conservative, flagging only the most extreme points of the minor cluster, while the FDR fence adapts to the evidence of contamination to create tighter boundaries. The PFER fence at level 0.5 produces an intermediate threshold. The refined whisker construction extends toward the minor component, indicating the direction of separation while avoiding the geometric constraints of a convex hull. The fences derived from the scaled F approximation are marginally more conservative than those based on the χ^2_2 distribution.

Example 2: Robustness to Correlation Structure. The second example employs a correlated normal mixture to demonstrate

how the framework adapts to covariance structure. Data are generated from the model in the first example, but with a covariance matrix $\begin{pmatrix} 1 & 0.3 \\ 0.3 & 1 \end{pmatrix}$. The `aplpack` visualization (Figure 5(a)) produces a jagged convex hull boundary which provides a noisy and unstable representation of the data's underlying correlation structure.

The bag-and-whisker plot excels in this scenario. Because our pipeline uses the robust MCD estimator for the covariance matrix, both the central bag and the resulting fences naturally elongate along the main axis of correlation (Figure 5(b)-(g)). This yields an interpretable elliptical-like contour that correctly reflects the data's covariance structure, providing a much more accurate and stable boundary for outlier detection in correlated multivariate settings.

Example 3: Visual Stability with Heavy-Tailed Data. Our final example highlights the importance of visual stability. We generate 500 points with independent x and y coordinates drawn from $\log\text{-normal}(0, 0.5)$ distribution, which is skewed and heavy-tailed but has no true contaminant outliers.




Figure 7. Reanalysis of the data from Figure 1 of Rousseeuw, Ruts, and Tukey (1999). Panel (a) displays the original implementation, where the outer boundary is constructed as the convex hull of non-outlying points. This yields a jagged, geometrically unstable loop that can contract in sparse regions. Panels (b)–(g) present the proposed method, which replaces the hull with a directly rendered fence obtained by magnifying the central bag, supplemented by gradient whiskers.

Figure 6(a) demonstrates the primary drawback of using a convex hull for visualization in such cases. The `aplpack` boundary is highly irregular and jagged, its shape dictated by the handful of most extreme observations. This creates a visually unstable representation that does not reflect the smooth underlying data distribution.

In contrast, the fences of the bag-and-whisker plot (Figure 6(b)–(g)) remain smooth and stable. By constructing the fence directly from a magnified version of the robustly estimated bag, our method is insensitive to small perturbations of peripheral points.

6. Real-World Data Applications

This section demonstrates the practical benefits of the new bagplot using real data. For a fair comparison, we revisit several of the classic datasets originally used to introduce the bagplot in Rousseeuw, Ruts, and Tukey (1999). This allows for a direct comparison of our method’s output with the visualizations produced by the conventional approach, as implemented in the R package `aplpack`. Specifically, we reexamine the data from Figure 1,

Figure 3(a), Figure 3(b), and Figure 6 of the original paper in our Figures 7, 8, 9, and 10, respectively.

6.1. Enhancing Visual Stability and Interpretability

We revisit the automobile dataset from Figure 1 of Rousseeuw, Ruts, and Tukey (1999), which records vehicle weight (x , in pounds) and engine displacement (y , in cubic inches) for $n = 60$ cars (Chambers and Hastie 1993). The data exhibit moderate positive correlation with a distinct upper-right cluster of high-displacement vehicles. The analytic objective of this reanalysis is to assess boundary stability in a moderate-sample setting where peripheral points typically dominate convex-hull geometry. Figure 7 contrasts the original `aplpack` implementation with the proposed method. Panel (a) displays the conventional convex-hull loop, which contracts inward in the sparse upper-right region, producing an artifact that misrepresents the underlying depth geometry. Panels (b)–(g) replace this loop with a smooth, convex fence derived from the data-adaptive inflation factor. The gradient whiskers extend toward the high-displacement cluster, clarifying the direction of the tail without obscuring the central




Figure 8. Reanalysis of the data from Figure 3(a) of Rousseeuw, Ruts, and Tukey (1999). Panel (a) displays the original convex-hull loop, yielding a jagged, vertex-driven boundary. Panels (b)–(g) replace the hull with a directly rendered fence. This eliminates polygonal artifacts, producing smooth and stable fences that more accurately reflect the data's structure.

bag. This revision eliminates the visual instability inherent in the bolster representation while preserving the original identification of four extreme vehicles as outliers.

We next examine the plasma concentration data from Figure 3(a) of Rousseeuw, Ruts, and Tukey (1999), which comprises measurements of cholesterol (x , mg/dl) and triglycerides (y , mg/dl) for $n = 320$ patients with arterial narrowing (Hand et al. 1994). The distribution is right-skewed with a heavy upper tail, a structure that typically challenges polygonal boundary definitions. Figure 8 presents the results. While both our method and the original bagplot identify the same set of clear outliers, our directly-rendered fence offers a more nuanced boundary. It smoothly delineates the space between the majority and outliers, enhancing the visual representation of the data's spread, whereas the convex hull in panel (a) is again jagged and sensitive to the specific locations of a few points.

6.2. Providing Statistical Flexibility in Outlier Detection

Beyond visual clarity, our framework provides a more powerful and flexible statistical tool. The classic bagplot offers a single,

static assessment based on a fixed rule. To illustrate the benefit of our adaptive approach, we re-analyze two datasets from the original paper. Figure 9 revisits the log-transformed plasma cholesterol and triglyceride concentrations for $n = 320$ patients with evidence of narrowing arteries (Hand et al. 1994). The log-transformation was originally applied to address the pronounced right-skewness in these chemical concentration measurements. Figure 10 examines the polychlorinated biphenyl (PCB) concentration and eggshell thickness for $n = 65$ Anacapa pelican eggs (Hand et al. 1994), a dataset originally used to study the environmental impact of industrial pollutants on bird reproduction.

In both examples, our bag-and-whisker plot identifies a varying number of outliers depending on the chosen error metric, exhibiting far greater flexibility than the original bagplot. This variability is not an inconsistency but a key feature of our framework. It empowers the analyst to choose an error metric and level that matches their scientific goals. The proposed bag-and-whisker plot effectively and transparently visualizes the results of this tailored statistical decision, offering a more nuanced and powerful tool for data exploration.




Figure 9. Reanalysis of the data from Figure 3(b) of Rousseeuw, Ruts, and Tukey (1999). Different choices of error metrics yield different inflation factors and outliers. The use of scaled F -distribution results in more conservative fences.

7. Conclusion

In this paper, we introduced a new bivariate boxplot, namely the bag-and-whisker plot, to address the long-standing statistical and graphical limitations of the original bagplot. Our contribution is twofold. First, we developed a general p -value pipeline that recasts outlier detection as a multiple testing problem, generating a data-adaptive fence whose stringency is aligned with formal error control principles. Second, and perhaps more significantly for data visualization, we replaced the unstable convex hull (the bolster) with a refined visualization that renders the fence directly and employs gradient whiskers. This ensures a one-to-one correspondence between the method's statistical rule and its graphical output. The result is not merely an adaptive update, but a structural evolution of the method, offering a principled and interpretable standard for modern exploratory data analysis.

We emphasize that the bag-and-whisker plot should be understood as a tool for exploratory data analysis rather than formal inference. While its fences are constructed in alignment with multiple testing principles, the procedure does

not rigorously control error rates due to its reliance on estimated null parameters. Its main contribution is conceptual, providing a coherent framework for developing and evaluating principled outlier detection rules for graphical exploration, rather than serving as a formal hypothesis test itself.

Acknowledging this exploratory role, the framework nevertheless opens several promising avenues for future research. While this paper focused only on the bivariate case for clarity of visualization, the underlying statistical methodology is inherently general and extends directly to higher dimensions. The Mahalanobis distance, robust estimators, and the multiple testing pipeline are well-defined for any dimensions. The primary challenge in higher dimensions is visualization, and future work could explore effective renderings of 3D bags or apply the framework within scatterplot matrices. Furthermore, the modularity of our pipeline invites extensions to non-normal data by substituting more appropriate reference distributions. By transforming a classic graphic into a dynamic and statistically transparent tool, the bag-and-whisker plot is better suited for the rigor and complexity of modern data exploration.




Figure 10. Reanalysis of the data from Figure 6 of Rousseeuw, Ruts, and Tukey (1999). Similar as the previous figures, different choices of error metrics yield different inflation factors and outliers. The use of scaled F -distribution results in more conservative fences again.

Acknowledgments

The authors would like to thank the Editor, the Associate Editor, and the two anonymous reviewers for their insightful comments and constructive suggestions, which greatly improved the presentation and clarity of this paper. The authors acknowledge the use of Google Gemini 2.5 Pro for language editing and refinement of the English text in this manuscript. The tool was used exclusively to improve grammatical correctness and clarity; all scientific content, data analysis, and conclusions remain the sole responsibility of the authors.

Disclosure statement

No potential conflict of interest was reported by the authors.

Funding

Tiejun Tong's research was supported in part by the General Research Fund of Hong Kong (HKBU12300123) and the Initiation Grant for Faculty Niche Research Areas of Hong Kong Baptist University (RC-FNRA-IG/23-24/SCI/03). Hengjian Cui's research was supported in part by the National Natural Science Foundation of China (12031016, 12531012).

Data Availability Statement

The source code used to generate the numerical results and figures presented in this paper is publicly available in the GitHub repository at <https://github.com/seanq31/BagWhiskerPlot>. The R package BagWhiskerPlot that implements the new bag-and-whisker plots is available on CRAN.

References

- Baliki, M. N., Schnitzer, T. J., Bauer, W. R., and Apkarian, A. V. (2011), "Brain Morphological Signatures for Chronic Pain," *PLoS One*, 6, e26010. DOI: [10.1371/journal.pone.0026010](https://doi.org/10.1371/journal.pone.0026010). [2]
- Beckett, S., and Gould, W. (1987), "Rangefinder Box Plots: A Note," *The American Statistician*, 41, 149–149. DOI: [10.1080/00031305.1987.10475467](https://doi.org/10.1080/00031305.1987.10475467). [1]
- Benjamini, Y., and Hochberg, Y. (1995), "Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing," *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 57, 289–300. DOI: [10.1111/j.2517-6161.1995.tb02031.x](https://doi.org/10.1111/j.2517-6161.1995.tb02031.x). [3,6]
- Chambers, J. M., and Hastie, T. J. (1993), *Statistical Models in S*. London, UK: Chapman and Hall. [10]
- Chauvenet, W. (1863), *A Manual of Spherical and Practical Astronomy*. Philadelphia: J. B. Lippincott Company. [8]

- Damialis, A., Gilles, S., Sofiev, M., Sofieva, V., Kolek, F., Bayr, D., Plaza, M. P., Leier-Wirtz, V., Kaschuba, S., Ziska, L. H., Bielory, L., Makra, L., del Mar Trigo, M., Traidl-Hoffmann, C., Oliver, G., Pham-Thi, N., Thibaudon, M., Arino, A. H., Belmonte, J., Morales, P. C., de Linares, C., Fernández, D., Fernandez-Rodriguez, S., Arguisuelas, A. G., Galan, C., González-Alonso, M., Lara, B., Grau, J. M. M., Oteros, J., Pérez-Badia, R., Pérez-de-Zabala, A., Picornell, A., Recio, M., Robles, E., Rodríguez-Fernández, A., Rodríguez-Rajo, J. F., Rojo, J., Valenzuela, L. R., Bergmann, K.-C., Werchan, B., Werchan, M., Buters, J. T. M., Bastl, M., Dunker, S., Hornick, T., Roldán, N. G., Gilge, S., Clot, B., Finemann, S., Ford, L. B., Gomez, R. A., Kamboj, S., Wilhelm, W., Beggs, P. J., Burton, P., Davies, J. M., Haberle, S., Katelaris, C. H., Keaney, B., Milic, A., Miller, V., van Haeften, S., Bonini, M., Bordin, A., Ceriotti, V., Cristofolini, F., Cristofori, A., Gottardini, E., Marcer, G., Marraccini, P., Mascagni, P., Meriggi, A., Pace, L., Pini, A., Tacca, M. C., Bruffaerts, N., Hoebeke, L., Adams-Groom, B., Pashley, C. H., Satchwell, J., Skjoth, C. A., Symon, F. A., Miguel Antunes, C. M., Caeiro, E., Camacho, I. G. C., Rodrigues Costa, A. C., Deus, R. J. R. C., Ferreira, M. B., Lopes Fonseca, J. A., Galveias Jorge, A. C., Ribeiro, H., Tavares, B., Grewling, L., Grinn-Gofroń, A., Jurkiewicz, D., Kalinowska, E., Lipiec, A. A., Myszkowska, D., Piotrowska-Weryszko, K., Puc, M., Rapijko, A. A., Rapijko, P., Weryszko-Chmielewska, E., Ziemianin, M., Berman, D., Hoek, W., Manjra, A. I., Peter, J., Dahl, A. A., Ekeboom, A., Stjepanovic, B., Večenaj, A., Celenk, S., Goksel, O., Goksel, T., Guvensen, A., Pinar, N. M., Sackesen, C., Sahin, A. A., Uguz, U., Yazici, D., Kajitor-Apatini, D., Magyar, D., Szigeti, T., Sikoparija, B., Seliger, A. K., Simčič, A., Oh, J.-W., Charalampopoulos, A., Vokou, D., Rasmussen, K., Barionuevo, L. B., Ramon, G. D., de Weger, L. A., Koenders, M. M., van Vliet, A. J. H., Dušička, J., Lafférová, J., Ščevková, J., Rybníček, O., Coates, F., Jurgens, D., Šaulienė, I., Severova, E., Rodinkova, V., Bortnyk, M., Palamarchuk, O., Yasniuk, M., Louna-Korteniemi, M., Päätsi, S., Saarto, A., Toivainen, L., Sozinova, O., and Jia, P.; COVID-19/POLLEN Study Group. (2021), "Higher Airborne Pollen Concentrations Correlated with Increased SARS-CoV-2 Infection Rates, as Evidenced from 31 Countries across the Globe," *Proceedings of the National Academy of Sciences*, 118, e2019034118. DOI: [10.1073/pnas.2019034118](https://doi.org/10.1073/pnas.2019034118). [1]
- Gang, B., Lin, H., and Tong, T. (2026), "Unifying Boxplots: A Multiple Testing Perspective," *Statistical Science*. In press. [1,3,5]
- Goldberg, K. M., and Iglewicz, B. (1992), "Bivariate Extensions of the Boxplot," *Technometrics*, 34, 307–320. [1]
- Hampel, F. R., Ronchetti, E. M., Rousseeuw, P. J., and Stahel, W. A. (2005), *Robust Statistics: The Approach Based on Influence Functions*. Chichester, UK: John Wiley & Sons. [4]
- Hand, D. J., Daly, F., McConway, K., Lunnon, D., and Ostrowski, E. (1994), *A Handbook of Small Data Sets*. London, UK: Chapman and Hall. [11]
- Hardin, J., and Rocke, D. M. (2005), "The Distribution of Robust Distances," *Journal of Computational and Graphical Statistics*, 14, 928–946. DOI: [10.1198/106186005X77685](https://doi.org/10.1198/106186005X77685). [4]
- Hoaglin, D. C., and Iglewicz, B. (1987), "Fine-Tuning Some Resistant Rules for Outlier Labeling," *Journal of the American Statistical Association*, 82, 1147–1149. DOI: [10.1080/01621459.1987.10478551](https://doi.org/10.1080/01621459.1987.10478551). [1]
- Hoaglin, D. C., Iglewicz, B., and Tukey, J. W. (1986), "Performance of Some Resistant Rules for Outlier Labeling," *Journal of the American Statistical Association*, 81, 991–999. DOI: [10.1080/01621459.1986.10478363](https://doi.org/10.1080/01621459.1986.10478363). [1]
- Holm, S. (1979), "A Simple Sequentially Rejective Multiple Test Procedure," *Scandinavian Journal of Statistics*, 6, 65–70. [6]
- Huber, P. J., and Ronchetti, E. M. (2009), *Robust Statistics*, 2nd ed. Chichester, UK: John Wiley & Sons, Inc. [4]
- Lin, H., Zhang, R., and Tong, T. (2026), "When Tukey Meets Chauvenet: A New Boxplot Criterion for Outlier Detection," *Journal of Computational and Graphical Statistics*, 35, 198–211. DOI: [10.1080/10618600.2025.2520577](https://doi.org/10.1080/10618600.2025.2520577). [1,8]
- Liu, R. Y., Parelius, J. M., and Singh, K. (1999), "Multivariate Analysis by Data Depth: Descriptive Statistics, Graphics and Inference (with Discussion and a Rejoinder by Liu and Singh)," *The Annals of Statistics*, 27, 783–858. DOI: [10.1214/aos/1018031260](https://doi.org/10.1214/aos/1018031260). [1,4]
- Maronna, R. A., Martin, R. D., Yohai, V. J., and Salibián-Barrera, M. (2019), *Robust Statistics: Theory and Methods (with R)*. Chichester, UK: John Wiley & Sons. [4]
- O'Connell, T. C., Kneale, C. J., Tasevska, N., and Kuhnle, G. G. (2012), "The Diet-Body Offset in Human Nitrogen Isotopic Values: A Controlled Dietary Study," *American Journal of Physical Anthropology*, 149, 426–434. [1]
- Rousseeuw, P. J. (1984), "Least Median of Squares Regression," *Journal of the American Statistical Association*, 79, 871–880. DOI: [10.1080/01621459.1984.10477105](https://doi.org/10.1080/01621459.1984.10477105). [4,6]
- Rousseeuw, P. J., and Driessens, K. V. (1999), "A Fast Algorithm for the Minimum Covariance Determinant Estimator," *Technometrics*, 41, 212–223. DOI: [10.1080/00401706.1999.10485670](https://doi.org/10.1080/00401706.1999.10485670). [7]
- Rousseeuw, P. J., and Ruts, I. (1996), "Bivariate Location Depth," *Journal of the Royal Statistical Society: Series C*, 45, 516–526. [6]
- Rousseeuw, P. J., Ruts, I., and Tukey, J. W. (1999), "The Bagplot: A Bivariate Boxplot," *The American Statistician*, 53, 382–387. DOI: [10.1080/00031305.1999.10474494](https://doi.org/10.1080/00031305.1999.10474494). [1,2,3,4,6,7,10,11,12,13]
- Schirpke, U., Candiago, S., Vigl, L. E., Jäger, H., Labadini, A., Marsoner, T., Meisch, C., Tasser, E., and Tappeiner, U. (2019), "Integrating Supply, Flow and Demand to Enhance the Understanding of Interactions among Multiple Ecosystem Services," *The Science of the Total Environment*, 651, 928–941. DOI: [10.1016/j.scitotenv.2018.09.235](https://doi.org/10.1016/j.scitotenv.2018.09.235). [1]
- Sim, C. H., Gan, F. F., and Chang, T. C. (2005), "Outlier Labeling with Boxplot Procedures," *Journal of the American Statistical Association*, 100, 642–652. DOI: [10.1198/016214504000001466](https://doi.org/10.1198/016214504000001466). [1]
- Tian, W., Zhang, L., and Cui, H. (2024), "Abnormal Sample Detection Based on Robust Mahalanobis Distance Estimation in Adversarial Machine Learning," *Statistics and Its Interface*, 17, 91–106. DOI: [10.4310/23-SII18](https://doi.org/10.4310/23-SII18). [4]
- Tong, T., Lin, H., Gang, B., and Zhang, R. (2026), "ChauBoxplot and AdaptiveBoxplot: two R Packages for Boxplot-Based Outlier Detection," *Statistical Theory and Related Fields*, 0, 1–10. DOI: [10.1080/24754269.2026.2642439](https://doi.org/10.1080/24754269.2026.2642439). [1]
- Tukey, J. W. (1960), "A Survey of Sampling from Contaminated Distributions," in *Contributions to Probability and Statistics: Essays in Honor of Harold Hotelling*, pp. 448–485. Stanford: Stanford University Press. [4]
- Tukey, J. W. (1975), "Mathematics and the Picturing of Data," *Proceedings of the International Congress of Mathematicians*, 2, 523–531. [1,2]
- (1977), *Exploratory Data Analysis*. Reading, MA: Addison-Wesley. [1]
- Wickham, H., and Stryjewski, L. (2011), "40 years of boxplots," Technical Report. <https://vita.had.co.nz/papers/boxplots.pdf>. [1]
- Zani, S., Riani, M., and Corbellini, A. (1998), "Robust Bivariate Boxplots and Multiple Outlier Detection," *Computational Statistics & Data Analysis*, 28, 257–270. DOI: [10.1016/S0167-9473\(98\)00040-1](https://doi.org/10.1016/S0167-9473(98)00040-1). [1,2]
- Zuo, Y., and Cui, H. (2005), "Depth Weighted Scatter Estimators," *The Annals of Statistics*, 33, 381–413. DOI: [10.1214/009053604000000922](https://doi.org/10.1214/009053604000000922). [4]