

40 years of boxplots

Hadley Wickham and Lisa Stryjewski

November 29, 2011

Abstract

The boxplot plot has been around for over 40 years. This paper summarises the improvements, extensions and variations since Tukey first introduced his “schematic plot” in 1970. We focus particularly on richer displays of density and extensions to 2d.

1 Introduction

John Tukey introduced the box and whiskers plot as part of his toolkit for exploratory data analysis (Tukey, 1970), but it did not become widely known until formal publication (Tukey, 1977). The boxplot is a compact distributional summary, displaying less detail than a histogram or kernel density, but also taking up less space. Boxplots use robust summary statistics that are always located at actual data points, are quickly computable (originally by hand), and have no tuning parameters. They are particularly useful for comparing distributions across groups.

Today, over 40 years later, the boxplot has become one of the most frequently used statistical graphics, and is one of the few plot types invented in the 20th century that has found widespread adoption. Due to their elegance and practicality, boxplots have spawned a wealth of variations and enhancement. This paper pulls these together in one place, showing how the boxplot has evolved.

We begin with a review of Tukey’s definition and an overview of minor variations to both the underlying summary statistics and their visual representation. Section 3 describes the richer displays of density facili-

tated by widespread desktop computing, and Section 4 explores how the boxplot has been extended to deal with 2d data. We conclude with some comments on the state of boxplot research and describe where future contributions are most needed.

The online supplementary materials include all R code (R Development Core Team, 2011) used to create plots in this paper, and features original code for four boxplots (vase plot, `quelp` plot, rotational boxplot, and bivariate clockwise boxplot) that previously lacked publicly available implementation.

2 Tukey's boxplot

The basic graphic form of the boxplot, the range-bar, was established in the early 1950's Spear (1952, pg. 164). Tukey's contribution was to think deeply about appropriate summary statistics that worked for a wide range of data and to connect those to the visual components of the range bar. Today, what we call a boxplot is more closely related to what Tukey called a schematic plot, a box and whiskers plot with some special restrictions on the summary statistics used.

The boxplot is made up of five components, carefully chosen to give a robust summary of the distribution of a dataset:

- the *median*,
- two *hinges*, the upper and lower fourths (quartiles),
- the data values adjacent to the upper and lower *fences*, which lie 1.5 times the inter-fourth range from the median,
- two *whiskers* that connect the hinges to the fences, and
- (potential) *out-liers*, individual points further away from the median than the extremes.

These elements are summarised in Figure 1. Our notation follows Tukey's, except where we can be more precise or where common usage has changed over the last 40 years.

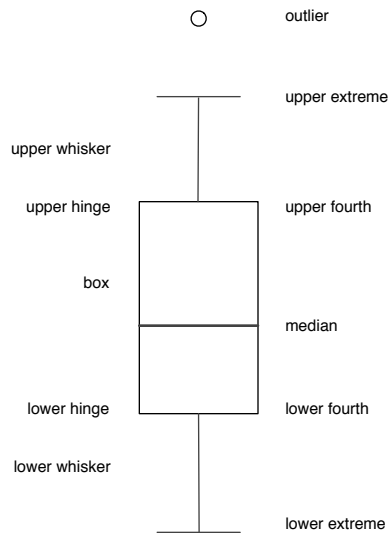


Figure 1: Construction of a boxplot. Labels on the left give names for graphic elements, labels on the right give the corresponding summary statistics.

There are a number of variations of these basic definitions. As well as variations in the definition of a quantile (Hyndman and Fan, 1996), some boxplots replace the extremes with fixed quantiles (e.g. min and max, 2% and 98%) or use multipliers other than 1.5 for the whiskers (Frigge et al., 1989). Others use the semi-interquartile ranges (e.g. $Q_1 - Q_2$) for asymmetric whiskers (Rousseuw et al., 1999), explicit adjustments to the extremes to account for skewness (Hubert and Vandervieren, 2008), alternative definitions of fences (Dümbgen and Riedwyl, 2007) or alternative definitions of outliers (Carter et al., 2009; Schwertman et al., 2004). Others have used additional graphical elements to display distributional features like kurtosis (Aslam and Khurshid, 1991), skewness and multimodality (Choonpradub and McNeil, 2005), and mean and standard error (Marmolejo-Ramos and Tian, 2010).

One of the appealing attributes of the boxplot is that if you have a rank function for the type of data you are dealing with, you can generate a boxplot. This makes it easy extend to the boxplot to work with weighted data, as described by Korn and Graubard (1998); Lumley (2011) for survey weights, by Willmott et al. (2007) for spatial area weights, and by Dykes and Brunson (2007) for distance weights.

In an effort to improve the data-ink ratio of the boxplot, (Tuft, 2001) proposed the midgap plot. As

shown in Figure 2, the box is removed and the median line replaced with a dot. No information is lost, and the boxplot becomes substantially more compact. However, perceptual studies (Stock and Behrens, 1991) have found Tufte's variation to be substantially less accurate than the original. Carr (1994) proposed a colourful variation, also shown in Figure 2. This variation is designed to be tightly perceptually linked, so that each boxplot appears a single object, not a collection of lines. No perceptual testing has been performed on this variant.

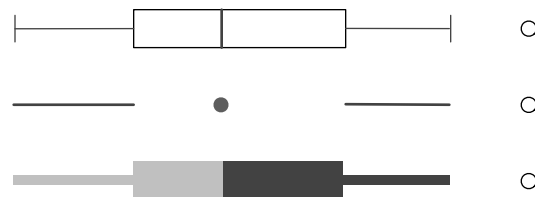


Figure 2: Tukey's original boxplot (top) compared to Tufte's box-less (middle) and Carr's colourful (bottom) variations. When colour is available, Carr suggests using red for components above the median and blue for colours below.

Another variation aims to overcome an important problem with the boxplot: there is visual display of group size, and hence no way of assessing if the differences are significant. The variable-width and notched boxplots (McGill and Larsen, 1978) add inferential detail. As the name suggests, the box widths of the variable-width boxplot vary according to the number of points in the group. The notched boxplot goes one step further by displaying confidence intervals around the medians, supporting visual assessment of statistical significance. The length of the confidence interval is determined heuristically so that non-overlapping intervals imply (approximately) a difference at the 5% level, regardless of the underlying distribution.

Other more unusual variations are an adaption for circular variables (Abuzaid et al., In press), and an adaption to make boxplots more suitable for display as glyphs Carr et al. (1998), particularly when overlaid on maps to display how data distribution varies in space.

There have been some perceptual studies on boxplots. Behrens et al. (1990) found evidence of significant bias when reading the length of the whiskers: whisker length was overestimated when whiskers were shorter than boxes and underestimated when whiskers were longer than boxes. There is a similar bias for reading the

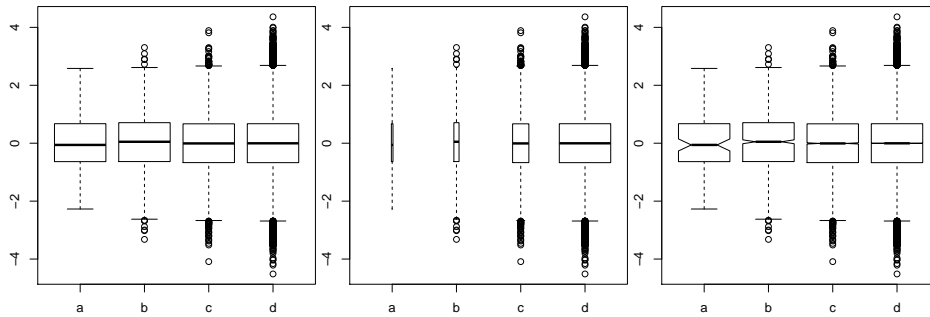


Figure 3: Boxplot variations showing 100, 1000, 10000, and 100000 numbers drawn from a standard normal distribution. (Left) In a regular boxplot the only hint that the groups are different sizes is the number of outliers. (Middle) A variable-width boxplot shows the differences in group size. (Right) The notched boxplots displays an inferentially meaningful quantity: the error associated with the estimate of the median.

length of boxes: box length is overestimated when boxes are shorter than whiskers and vice-versa. Notched plots appear to suffer from similar problems (Wells and Layne, 1996).

3 Richer displays of density

One of the original constraints on the boxplot was that it was designed to be computed and drawn by hand. As every statistician now has a computer on their desk, this constraint can be relaxed, allowing variations of the boxplot that are substantially more complex. These variations attempt to display more information about the distribution, maintaining the compact size of the boxplot, but bringing in the richer distributional summary of the histogram or density plot. These plots can overcome problems in the original such as the failure to display multi-modality, or the excessive number of “outliers” when n is large.

The first variation to display a density estimate was the *vase plot* (Benjamini, 1988), where the box is replaced with a symmetrical display of estimated density. *Violin plots* (Hintze and Nelson, 1998) are very similar, but display the density for all data points, not just the middle half. The *bean plot* (Kampstra, 2008) is a further enhancement that adds a rug that showing every value and a line that shows the mean. The name is inspired by the appearance of the plot: the shape of the density looks like the outside of a bean

pod, and the rug plot looks like the seeds within. Kampstra (2008) also suggests a way of comparing two groups more easily: use the left and right sides of the bean to display different distributions. A related idea is the raindrop plot (Barrowman and Myers, 2003), but its focus is on the display of error distributions from complex models.

Figure 4 demonstrates these density boxplots applied to 100 numbers drawn from each of four distributions with mean 0 and standard deviation 1: a standard normal, a skew-right distribution (Johnson distribution with skewness 2.2 and kurtosis 13), a leptikurtic distribution (Johnson distribution with skewness 0 and kurtosis 20) and a bimodal distribution (two normals with mean -0.95 and 0.95 and standard deviation 0.31). Richer displays of density make it much easier to see important variations in the distribution: multi-modality is particularly important, and yet completely invisible with the boxplot.

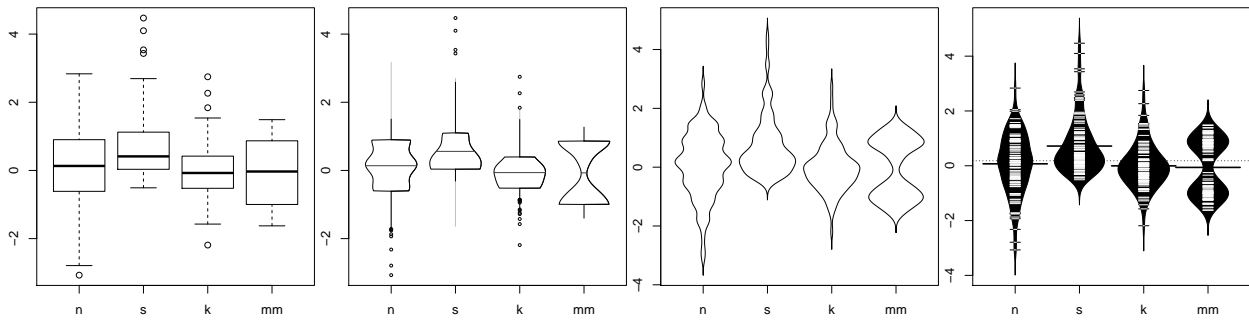


Figure 4: From left to right: box plot, vase plot, violin plot and bean plot. Within each plot, the distributions from left to right are: standard normal (n), right-skewed (s), leptikurtic (k), and bimodal (mm). A normal kernel and bandwidth of 0.2 are used in all plots for all groups.

A more sophisticated display is the sectioned density plot (Cohen and Cohen, 2006), which uses both colour and space to stack a density estimate into a smaller area, hopefully without losing any information (not formally verified with a perceptual study). The sectioned density plot is similar in spirit to horizon graphs for time series (Reijner, 2008), which have been found to be just as readable as regular line graphs despite taking up much less space (Heer et al., 2009). The density strips of Jackson (2008) provide a similar compact display that uses colour instead of width to display density. These methods are shown in Figure 5.

The summary plot (Potter et al., 2010) is a similar idea. It combines a minimal boxplot with glyphs representing the first five moments (mean, standard deviation, skewness, kurtosis and tailings), and a sectioned density plot crossed with a violin plot (both colour and width are mapped to estimated density), and an overlay of a reference distribution. It is a rather busy display.

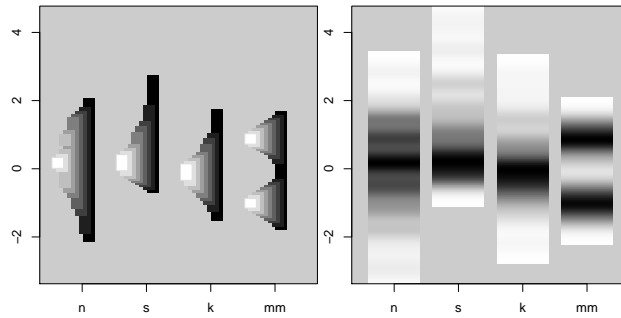


Figure 5: (Left) sectioned density plot and (right) density strips, same four distributions as Figure 4. A normal kernel and bandwidth of 0.2 are used in all plots for all groups.

The highest density region (HDR) boxplot (Hyndman, 1996) is a compromise between a boxplot and a density boxplot. It uses a density estimate but shows only two regions of highest density: the top 50% and 99%. These regions do not need to be contiguous and make it easy to spot multi-modality. The disadvantage of HDR boxplots is a less-sophisticated definition of extremes, making the outliers less useful for non-normal data. Figure 6 shows the HDR boxplot for the four distributions previously described.

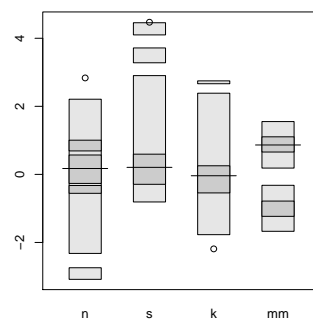


Figure 6: The highest density region boxplot for the same four distributions as Figure 4. The multimodality in the fourth distribution is easy to spot.

Each author has suggested a different density estimate to use in conjunction with their new display, but

there is no reason not to use any desired estimator. This is the price of density boxplots: the explosion of choices. Which density estimate should you use? Which choice of bandwidth or bin width is best? Bandwidth estimation is particularly challenging: if multiple groups are displayed, should each group get its own bandwidth, or should one bandwidth be used for all? Kampstra (2008) suggests using the average of the per-group bandwidth estimates. The following two methods attempt a richer display of density without the cost of additional tuning parameters.

The box-percentile plot (Esty and Banfield, 2003) displays a modified empirical cumulative density function (ECDF). The width of each box is proportional to the percentile, up to the 50th percentile, after which the width is proportional to one minus the the percentile. Lines mark the median and upper and lower quartiles. While this display of the ECDF contains all information about the distribution, it is not always easy to parse this data into an informative mental model. This is illustrated in Figure 7: without training, it is very difficult to tell that the fourth distribution is bimodal.

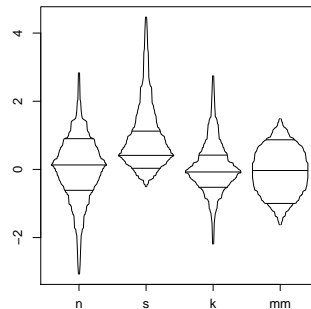


Figure 7: Box-percentile plots for the same four distributions used in Figure 4. The multimodality in the fourth distribution is hard to spot.

The letter-value boxplot (Hofmann et al., 2006) was designed to overcome the shortcomings of the boxplot for large data. For large datasets ($n \gtrsim 10,000$), the boxplot displays many outliers, and doesn't take advantage of the more reliable estimates of tail behaviour. The letter-value boxplot extends the boxplot with additional letter-values apart from the median (M) and fourths (F): eighths (E), sixteenths (D), ..., until the estimation error becomes too large. Each additional letter-value is displayed with a slightly smaller box, as

shown in Figure 8. Because this display just adds extra letter values, it suffers from the same problems as the original boxplot, and multimodality is almost impossible to spot.

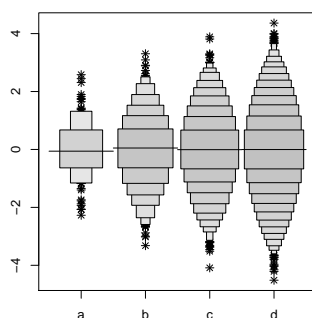


Figure 8: Letter value plots of 100, 1000, 10,000, and 100,000 points drawn from a standard normal distribution. The letter-value boxplot automatically displays additional letter values as the sample size increases.

4 Extensions to 2d

Extending the boxplot to work in 2d is challenging because of the difficulty of defining order statistics, depth and quantiles in 2d. There is no unique definition of rank in 2d dimensions, and hence the 2d analogues of medians, fourths and extremes becomes more complex (mathematically and computationally). Perhaps due to the increased complexity of creating just a single plot, there has been little development of effective methods for comparing multiple groups. Compared to 1d, it is less obvious that 2d boxplots provide significant advantages over contours of density estimates or heatmaps of binned counts.

The first attempt at generalising the boxplot to 2d was the *rangefinder plot* (Beckett and Gould, 1987). This generalisation is simple, treating the two variables as independent and drawing two independent 1d boxplots. The *relplot* (Goldberg and Iglewicz, 1992) relaxes the assumption of independence by robustly fitting a bivariate Gaussian to the data, and drawing 50% (corresponding to the box) and 99% (corresponding to the whiskers) confidence ellipses. The *quelplot* (Goldberg and Iglewicz, 1992) relaxes the assumption of normality by adding two degrees of asymmetry, accounting for residuals on both the major and minor axes of the ellipse. Another variation is the 2d boxplot of Tongkumchum (2005), which effectively constructs a

1d boxplot parallel to a (robust) line of best fit to the data.

The *bagplot*, (Rousseeuw et al., 1999), is a 2d analog of the boxplot with a *bag* and *fence* (cf. box) containing the middle 50% of the data and a *loop* (cf. whiskers) to separate outliers. The definitions of these regions follows naturally from the boxplot (although the numerical constants differ somewhat, adjusted by simulation to get behaviour that matches the boxplot), using the half-space depth in lieu of rank. The half-space depth, the smallest number of points contained within a half-plane of any orientation at a point, was informally defined early by Tukey (1975), but an efficient algorithm for computing it was not available until over 20 years later (Rousseeuw and Ruts, 1996). A similar approach is the robust bivariate boxplot (Zani et al., 1998), which uses convex hull peeling to find the central regions, and then displays the hulls smoothed with a b-spline.

Figure 9 shows the range finder plot, relplot, quelplot and bagplot. The data is generated from a mixture of two highly correlated bivariate normals.

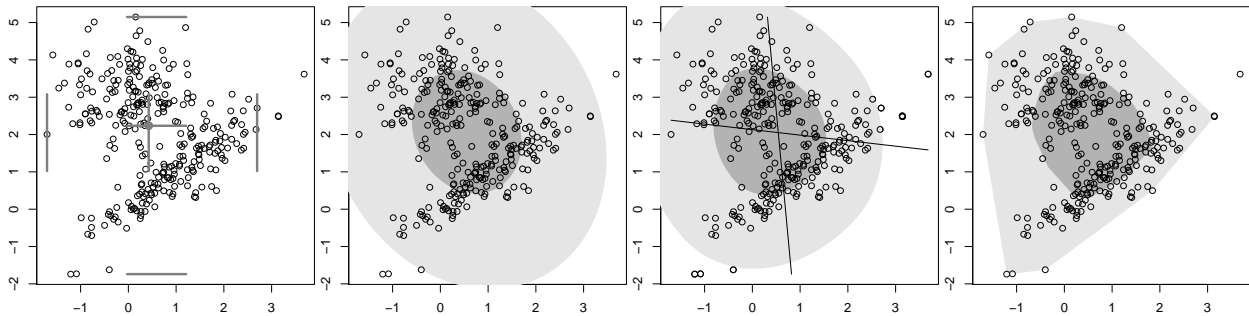


Figure 9: From left to right: a range finder plot, a relplot, a quelplot and a bagplot.

The 1d HDR boxplot extends in a straightforward manner to 2d (Hyndman, 1996). Unlike the methods described above, this 2d boxplot relies on a density estimate, and hence the selection of bandwidth is critically important. Figure 10, shows the 2d HDR region boxplots with the 50% and 95% highest density regions, for three different bandwidths.

Two other approaches take a circular approach. The rotational boxplot of Muth et al. (2000) bins the data into multiple overlapping circular sections around the centroid then computes the boxplot summary

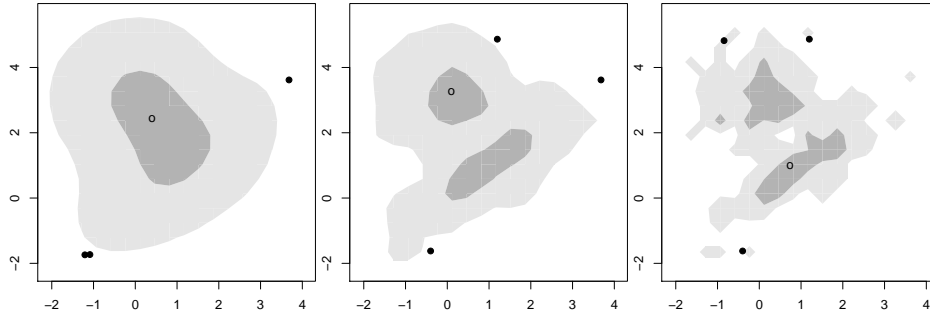


Figure 10: The 2d HDR boxplot, with both bandwidths set to (from left to right), 5, 2.5 and 1. Choice of bandwidth makes a critical difference in the appearance of the plot.

statistics for each section. The whiskers, hinges and medians are connected with lines. The clockwise bivariate boxplots of Corbellini (2002) are a similar idea, but based on projections rather than sections: the full dataset is projected onto multiple lines through the origin and the boxplot statistics computed for each. Both plots are shown in Figure 11. These graphics must be read carefully because while they look similar to topological maps, but there is no guarantee that density is highest in the central region. They are most useful for displaying the distribution of distances from a location.

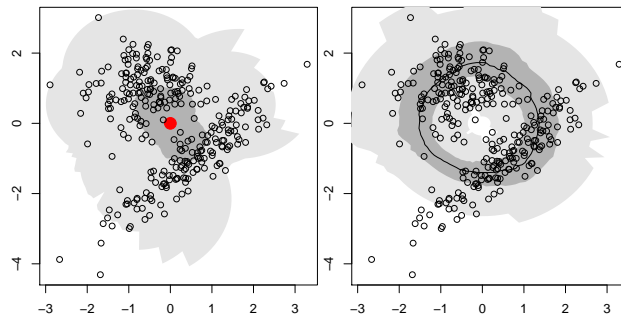


Figure 11: (Left) The clockwise bivariate boxplot displays hinges and whiskers, and (right) the rotational boxplot displays hinges, median (thicker) and whiskers (in grey). The rotational boxplot has a bin width of 90° .

Functional boxplots (Hyndman and Shang, 2010; Sun and Genton, 2011) take the ideas of 2d boxplots and extend them to the infinite-dimensional functional case. The extensions are similar in spirit to the bagplot, first finding a definition of rank in a functional space, and then extending the definition of the boxplot to use that definition, and finally finding an effective display technique.

5 Discussion

Boxplots were created to provide a succinct distributional summary that could easily be created by hand, and supported comparison across groups. As computers have become more prevalent and more powerful, it has become easier to produce compact summaries that display more data. This has led to an explosion of boxplot variations that stay true to the original goals to various extents, while supporting much richer display of the underlying distributions.

There are a few places where the existing literature is weak. We wonder if the proposed boxplot variations do actually make it easier to compare distributions. Cleveland's hierarchy (Cleveland and McGill, 1984) provides some guidance suggesting that techniques that use position, rather than colour, should be more effective, but there has been little empirical testing. Additionally, despite the number of 2d boxplots that have been developed, none attack one of the original motivating problems of boxplots: comparing the distributions of multiple groups.

Boxplots have a rich history, and a promising future. We hope that this paper has illuminated the past, and suggested new avenues for the future.

6 Acknowledgements

One of the biggest challenges of this paper was recreating all of the plots with consistent formatting. We're indebted to the authors who have provided implementations in R packages: `hdcrcde` (with contributions from Jochen Einbeck and Wand, 2010), `UsingR` (Verzani, 2011), `beanplot` (Kampstra, 2008), `denstrip` (Jackson, 2008), `aplpack` (Wolf and Bielefeld, 2010) and `lvp1ot`. Code for the `relplot` came from the website of Rand Wilcox. We wrote our own implementations of the vase plot and the `quelplot` (but only using non-robust estimators). Tufte's mid gap plot and Carr's colourful boxplot were drawn by hand.

References

- A. Abuzaid, I. Mohamed, and A. Hussin. Boxplot for circular variables. *Computational Statistics*, pages 1–12, In press. doi: 10.1007/s00180-011-0261-5.
- M. Aslam and A. Khurshid. Shape-finder box plots. *ASQC Statistics Division Newsletter*, pages 9–11, Fall 1991. URL <http://mail.asq.org/statistics/2011/03/statistics/shape-finder-box-plots.pdf>.
- N. J. Barrowman and R. A. Myers. Raindrop plots: A new way to display collections of likelihoods and distributions. *The American Statistician*, 57(4), 2003.
- S. Beckett and W. Gould. Rangefinder box plots: A note. *American Statistician*, pages 149–149, 1987.
- J. Behrens, W. Stock, and C. Sedgwick. Judgment errors in elementary box-plot displays. *Communications in Statistics-Simulation and Computation*, 19(1):245–262, 1990.
- Y. Benjamini. Opening the box of a boxplot. *The American Statistician*, 42(4):257–262, 1988.
- D. Carr. A colorful variation on boxplots. *Statistical Computing & Graphics Newsletter*, 5(3):19–23, 1994.
- D. Carr, A. Olsen, S. Pierson, and J. Courbois. Boxplot variations in a spatial context: An omernik ecoregion and weather example. *Statistical Computing & Statistical Graphics Newsletter*, 9(2):4–15, 1998.
- N. J. Carter, N. C. Schwertman, and T. L. Kiser. A comparison of two boxplot methods for detecting univariate outliers which adjust for sample size and asymmetry. *Statistical Methodology*, 6(6):604–621, 2009. doi: 10.1016/j.stamet.2009.07.001.
- C. Choonpradub and D. McNeil. Can the box plot be improved? *Songklanakarin Journal of Science and Technology*, 27(3):649–657, 2005.
- W. S. Cleveland and R. McGill. Graphical perception: Theory, experimentation and application to the development of graphical methods. *Journal of the American Statistical Association*, 79(387):531–554, 1984.

- D. J. Cohen and J. D. Cohen. The sectioned density plot. *The American Statistician*, 60:167–174, 2006.
- A. Corbellini. Clockwise bivariate boxplots. In *Compstat: Proceedings in Computational Statistics: 15th Symposium Held in Berlin, Germany, 2002*, pages 231–236. Physica Verlag, 2002. URL <http://books.google.com/books?id=YDD95HD-KE8C&pg=PA231>.
- L. Dümbgen and H. Riedwyl. On fences and asymmetry in box-and-whiskers plots. *The American Statistician*, 61(4):356–359, 2007.
- J. Dykes and C. Brunsdon. Geographically weighted visualization - interactive graphics for scale-varying exploratory analysis. In *IEEE Transactions on Visualization and Computer Graphics*, volume 13, pages 1161–1168, 2007.
- W. W. Esty and J. D. Banfield. The box-percentile plot. *Journal of Statistical Software*, 8(17):1–14, 2003.
- M. Frigge, D. Hoaglin, and B. Iglewicz. Some implementations of the boxplot. *American Statistician*, 43(1): 50–54, 1989.
- K. M. Goldberg and B. Iglewicz. Bivariate extensions of the boxplot. *Technometrics*, 34(3):307–320, 1992.
- J. Heer, N. Kong, and M. Agrawala. Sizing the horizon: The effects of chart size and layering on the graphical perception of time series visualizations. In *CHI 2009*, pages 1303–1312. ACM, 2009.
- J. L. Hintze and R. D. Nelson. Violin plots: A box plot-density trace synergism. *The American Statistician*, 52: 181–184, 1998.
- H. Hofmann, K. Kafadar, and H. Wickham. Letter-value box plots - adjusting box plots for large data sets. Technical Report 10, Department of Statistics, Iowa State University, 2006. URL <http://www.stat.iastate.edu/preprint/articles/2006-10.pdf>.
- M. Hubert and E. Vandervieren. An adjusted boxplot for skewed distributions. *Computational Statistics & Data Analysis*, 52(12):5186–5201, 2008.

- R. Hyndman. Computing and graphing highest density regions. *American Statistician*, 50(2):120–126, 1996.
- R. Hyndman and Y. Fan. Sample quantiles in statistical packages. *American Statistician*, 50(4):361–365, 1996.
- R. J. Hyndman and H. L. Shang. Rainbow plots, bagplots, and boxplots for functional data. *Journal of Computational and Graphical Statistics*, 19(1):29–45, 2010. doi: 10.1198/jcgs.2009.08158.
- C. H. Jackson. Displaying uncertainty with shading. *The American Statistician*, 62(4), 2008.
- P. Kampstra. Beanplot: A boxplot alternative for visual comparison of distributions. *Journal of Statistical Software, Code Snippets*, 28(1):1–9, 11 2008. ISSN 1548-7660. URL <http://www.jstatsoft.org/v28/c01>.
- E. L. Korn and B. I. Graubard. Scatterplots with survey data. *The American Statistician*, 52(1):pp. 58–69, 1998. ISSN 00031305. URL <http://www.jstor.org/stable/2685570>.
- T. Lumley. survey: analysis of complex survey samples, 2011. R package version 3.26.
- F. Marmolejo-Ramos and T. Tian. The shifting boxplot. *International Journal of Psychological Research*, 3(1): 37–45, 2010.
- T. J. W. McGill, R. and W. Larsen. Variations of box plots. *The American Statistician*, 32:12–16, 1978.
- S. Muth, J. Potterat, and R. Rothenberg. Birds of a feather: using a rotational box plot to assess ascertainment bias. *International Journal of Epidemiology*, 29(5):899–904, 2000.
- K. Potter, J. Kniss², R. Riesenfeld, and C. Johnson. Visualizing summary statistics and uncertainty. In *Eurographics/IEEE-VGTC Symposium on Visualization 2010*, volume 29, page 3, 2010.
- R Development Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2011. URL <http://www.R-project.org/>. ISBN 3-900051-07-0.

- H. Reijner. The development of the horizon graph. In *IEEE Visualization Workshop: From Theory to Practice: Design, Vision and Visualization*, volume 3, 2008.
- P. Rousseeuw and I. Ruts. Algorithm as 307: Bivariate location depth. *Journal of the Royal Statistical Society. Series C (Applied Statistics)*, 45(4):516–526, 1996.
- P. J. Rousseeuw, I. Ruts, and J. W. Tukey. The bagplot: A bivariate boxplot. *The American Statistician*, 53: 382–387, 1999.
- N. Schwertman, M. Owens, and R. Adnan. A simple more general boxplot method for identifying outliers. *Computational Statistics & Data Analysis*, 47(1):165–174, 2004.
- M. E. Spear. *Charting Statistics*. McGraw-Hill Book Company, Inc., 1952.
- W. A. Stock and J. T. Behrens. Box, line, and midgap plots: Effects of display characteristics on the accuracy and bias of estimates of whisker length. *Journal of Educational Statistics*, 16(1):1–20, 1991.
- Y. Sun and M. G. Genton. Functional boxplots. *Journal of Computational and Graphical Statistics*, 20(2): 316–334, 2011. doi: 10.1198/jcgs.2011.09224.
- P. Tongkumchum. Two-dimensional box plot. *Songklanakarinn Journal of Science and Technology*, 27(4): 859–866, 2005.
- E. R. Tufte. *The Visual Display of Quantitative Information*. Graphics Press, 2001.
- J. W. Tukey. *Exploratory Data Analysis*. Addison–Wesley, preliminary edition, 1970.
- J. W. Tukey. Mathematics and the picturing of data. In *Proceedings of the International Congress of Mathematicians*, volume 2, pages 523–531, 1975.
- J. W. Tukey. *Exploratory Data Analysis*. Addison–Wesley, 1977.
- J. Verzani. *UsingR: Data sets for the text "Using R for Introductory Statistics"*, 2011. URL <http://CRAN.R-project.org/package=UsingR>. R package version 0.1-17.

- J. Wells and B. Layne. The effects of display characteristics on the bias of estimates of whisker length of regular and notched boxplots. *Journal of Educational and Behavioral Statistics*, 21(3):247, 1996.
- C. Willmott, S. Robeson, and K. Matsuura. Geographic box plots. *Physical Geography*, 28(4):331–344, 2007.
- R. J. H. with contributions from Jochen Einbeck and M. Wand. *hdrcde: Highest density regions and conditional density estimation*, 2010. URL <http://CRAN.R-project.org/package=hdrcde>. R package version 2.15.
- P. Wolf and U. Bielefeld. *aplpack: Another Plot PACKage: stem.leaf, bagplot, faces, spin3R, and some slider functions*, 2010. URL <http://CRAN.R-project.org/package=aplpack>. R package version 1.2.3.
- S. Zani, M. Riani, and A. Corbellini. Robust bivariate boxplots and multiple outlier detection. *Computational Statistics & Data Analysis*, 28(3):257–270, 1998.