Computing and Graphing Highest Density Regions

Rob J. HYNDMAN

Many statistical methods involve summarizing a probability distribution by a region of the sample space covering a specified probability. One method of selecting such a region is to require it to contain points of relatively high density. Highest density regions are particularly useful for displaying multimodal distributions and, in such cases, may consist of several disjoint subsets—one for each local mode. In this paper, I propose a simple method for computing a highest density region from any given (possibly multivariate) density $f(x)$ that is bounded and continuous in x. Several examples of the use of highest density regions in statistical graphics are also given. A new form of boxplot is proposed based on highest density regions; versions in one and two dimensions are given. Highest density regions in higher dimensions are also discussed and plotted.

KEY WORDS: Boxplots; Density estimation; Graphical summary; Highest density regions.

1. INTRODUCTION

Many statistical methods involve summarizing a probability distribution by a region of the sample space covering a specified probability. For example, reporting a prediction interval for a future observation involves an implicit summary of the underlying distribution. A boxplot is an alternative summary where the various components mark regions of the empirical distribution function corresponding approximately to regions of the underlying probability distribution.

In summarizing a probability distribution by a region, it is not always clear which region should be used. Suppose you wish to give a two-sided 95% prediction interval from a given distribution. Should you use the interval symmetric about the mean, the interval symmetric about the median, the interval defined between the 2.5% and 97.5% quantiles, the interval of shortest length, or the interval that minimizes the probability of covering a given set? Each of these intervals has 95% coverage, and they may all be different. In higher dimensions the difficulty of selecting an appropriate region is even greater.

This article investigates one approach to this problem by suggesting that highest density regions are often the most appropriate subset to use to summarize a probability distribution. As a motivating example Figure 1 shows five possible 75% regions for a mixture density comprising an $N(0,1)$ density and an $N(4,1)$ density with weights .7 and .3, respectively. Only the highest density region shows the bimodality.

The usual purpose in summarizing a probability distribution by a region of the sample space is to delineate a comparatively small set which contains most of the probability, although the density may be nonzero over infinite regions of the sample space. This is the underlying philosophy of prediction regions and boxplots. Clearly, there are an infinite number of ways to choose a region with given coverage probability. In some cases the nature of the problem will suggest a specific region. (For instance, "one-sided" regions of given coverage are uniquely defined.) Often, though, which region to use is not clear, and it is necessary to decide what properties we wish the region to possess. The following criteria seem intuitively sensible:

1. The region should occupy the smallest possible volume in the sample space;
2. Every point inside the region should have probability density at least as large as every point outside the region.

In fact, the criteria are equivalent (Box and Tiao 1973) and lead to regions called highest density regions or HDR's. I shall define them more precisely using the second criterion.

Definition. Let $f(x)$ be the density function of a random variable X. Then the $100(1-\alpha)\%$ HDR is the subset $R(f_\alpha)$ of the sample space of X such that

$$R(f_\alpha) = \{x\colon f(x) \geq f_\alpha\}$$

where f_α is the largest constant such that $\Pr(X \in R(f_\alpha)) \geq 1-\alpha$.

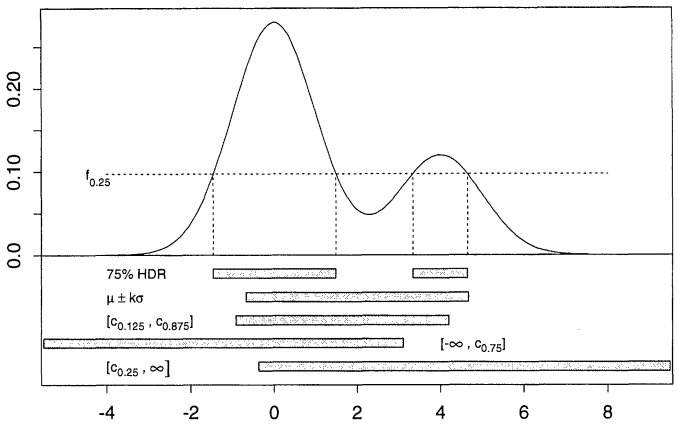

Figure 1. Five Different 75% Probability Regions From a Normal Mixture Density. Here, c_q denotes the qth quantile, μ denotes the mean, and σ denotes the standard deviation of the density.

Rob Hyndman is Lecturer, Department of Mathematics, Monash University, Clayton, Vict., Australia 3168. The author thanks Gary Grunwald (University of Melbourne) for several constructive discussions in developing the ideas contained here, a referee for several suggestions that led to improvements in the paper, and David Scott (Rice University, Houston, TX) for generously providing the software to produce Figure 6.

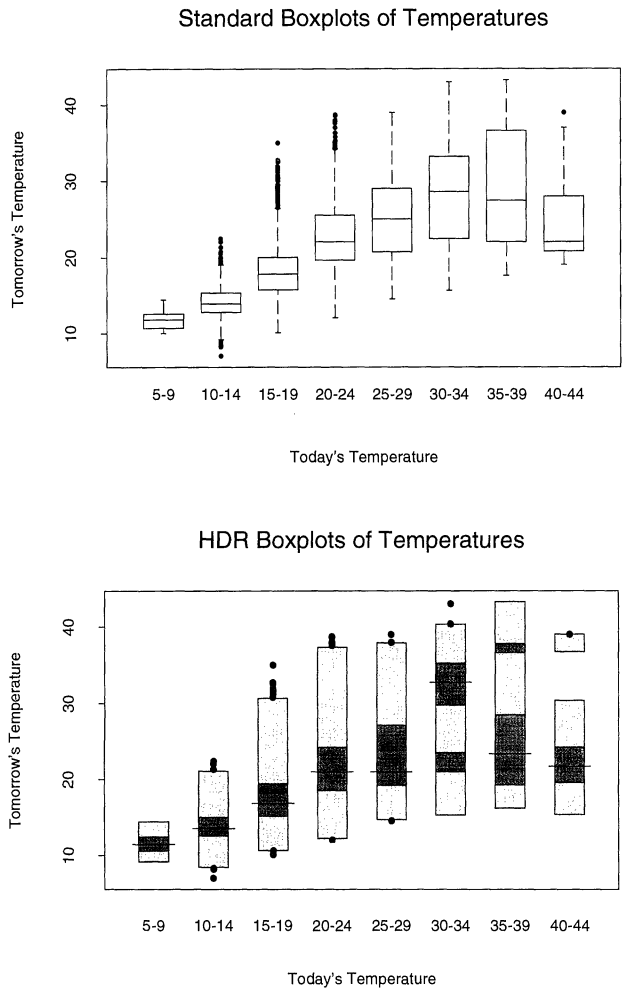

Figure 2. Boxplots for the Daily Maximum Temperature in Melbourne, Australia Between 1981 and 1990. The bimodality of the distribution following a hot day, which is due to the possible onset of a cool change, is only visible with the HDR boxplots.

The same definition also applies for discrete variables with the density function replaced by the probability mass function.

Figure 1 shows the particular value of $f_{.25}$, which was used to construct the 75% region shown.

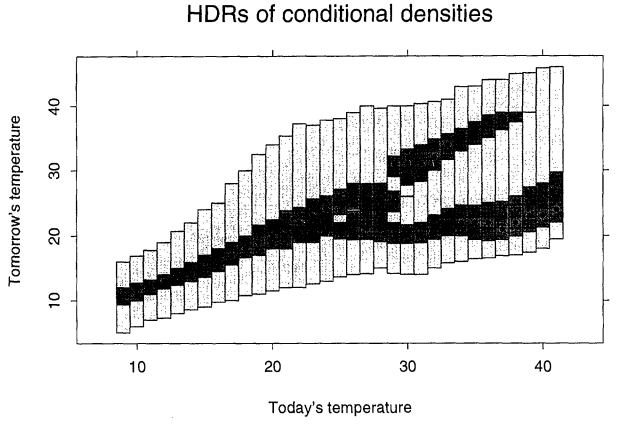

Figure 3. An Extension of Figure 2. Highest density regions (50% and 99%) for tomorrow's temperature conditional on today's temperature are displayed. The conditional modes are marked by • for each conditioning value.

It follows immediately from the definition that of all regions of probability coverage $1-\alpha$, the HDR has the smallest possible volume in the sample space of X. Furthermore, the mode is contained in every HDR. Unlike other probability regions, HDR's are easily summarized, even for high dimensions, as they are defined by a single number f_α.

Such regions are common in Bayesian analysis where they are applied to a posterior distribution (e.g., Box and Tiao 1973). In that context they are also called "credible sets," "plausible sets," "highest posterior density regions," or "Bayesian confidence sets."

In the case of a normal distribution an HDR coincides with the usual probability region symmetric about the mean, spanning the $\alpha/2$ and $1-\alpha/2$ quantiles. The same is true for any unimodal, symmetric distribution. However, in the case of a multimodal distribution, an HDR often consists of several disjoint subregions. This provides useful information which is "masked" by other probability regions.

We shall first demonstrate the value of highest density regions by considering several graphical display methods in Section 2. The computational algorithms that were used to construct these graphs will be discussed in Section 3.

2. GRAPHICAL DISPLAY

2.1 HDR Boxplots

Boxplots, introduced by Tukey (1977), are a common method for summarizing univariate samples. Several variations of boxplots are discussed by McGill, Tukey, and Larsen (1978), Benjamini (1988), and Esty and Banfield (1992). All of these boxplot variants include a central box bounded by Q_1 and Q_3 where Q_1 and Q_3 denote sample quartiles or approximate sample quartiles such as fourths. Hence the box contains approximately 50% of the observations. The most common form of boxplot contains "whiskers" extending to $1.5(Q_3 - Q_1)$ beyond the ends of the central box. For a standard normal distribution the quartiles are $\pm.6745$, so for a large data set, the whiskers lie at approximately $\pm[.6745 + 1.5(1.349)] = \pm 2.698$. Hence the probability of an observation falling inside the whiskers is approximately $1 - 2[1 - \Phi(2.698)] = 99.30\%$ for large samples. (Hoaglin, Iglewicz, and Tukey (1986) show that for small to moderate samples the probability of observations inside the whiskers is smaller than this asymptotic rate.)

Here, I propose a new form of boxplot based on HDR's which summarizes the distribution in a similar way, but allows the display of multimodality. An HDR boxplot replaces the box bounded by the interquartile range with the 50% HDR. In both cases the coverage probability of the region is 50%, but only the HDR will display multimodality. Similarly, the region bounded by the whiskers is replaced by the 99% HDR. This is chosen to roughly reflect the probability coverage of the whiskers on a standard boxplot for a normal distribution. Data beyond the 99% HDR are displayed as points. To emphasize the different densities of the two regions, shaded boxes are used with higher density shading in the 50% HDR. Finally, in keeping with the emphasis on highest density, the mode rather than the median is marked by a horizontal line.

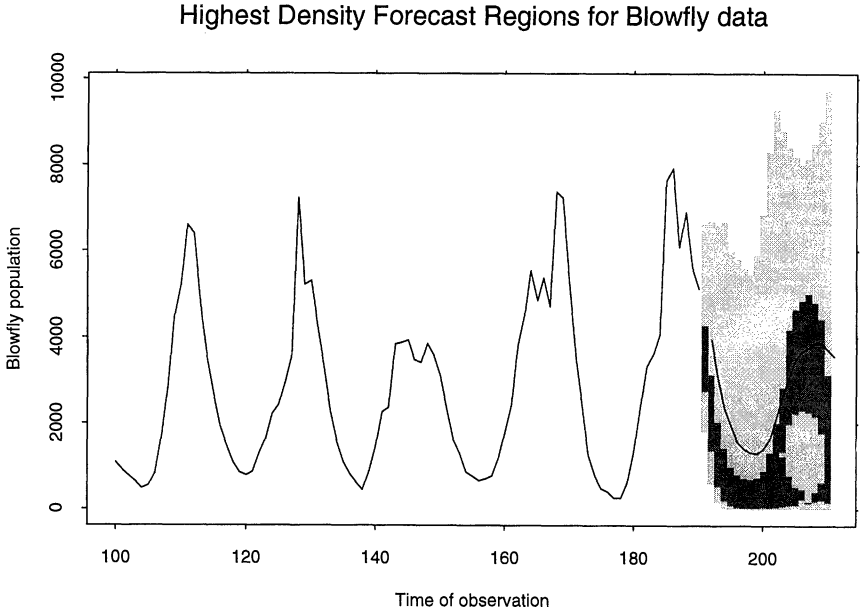

Figure 4. Forecast Means and 50% and 95% Forecast Regions for Nicholson's Australian Blowfly Data Based on a Threshold Model for the Logged Population.

Example. Figure 2 shows both standard and HDR boxplots for the daily maximum temperature in Melbourne, Australia between 1981 and 1990. The data are divided according to the temperature of the previous day, and the density for each group was estimated using a kernel density estimator (e.g., Scott 1992). The HDR boxplots were calculated using the density quantile approach of Section 3. Both displays demonstrate that the mean and variance of tomorrow's temperature increase as today's temperature increases, except on very hot days (over 40°C) which tend to be followed by cooler days. However, only the HDR boxplots reveal the bimodality of the distribution of the temperature following a warm to hot day. The 50% HDR's consist of two disjoint intervals showing that days of 30–39°C tend to be followed by days of similar temperature or of much lower temperature; they are not generally followed by days with maximum temperature in the high 20s. This occurs because temperatures slowly increase as high-pressure systems pass over the city from west to east. At the tail end of a high-pressure system a strong north wind often blows (from off the Australian mainland), bringing high temperatures. A high-pressure system is often followed by a cold front, causing a rapid drop in temperature. Hence hot days are generally followed by days of similar or greater temperature or by much cooler days. The hotter the day, the more likely it is to be followed by a cool day.

The bimodality of these distributions is perhaps the most interesting feature of these data. Yet standard boxplots give no hint that such a feature exists.

The vaseplots of Benjamini (1988) can show some forms of multimodality in that they allow the shape of the central box to be proportional to a density estimate of the data. However, there are several drawbacks to vaseplots compared with HDR boxplots. If the modes are close to or outside the quartiles, as in the 30–34°C group of temperatures in Figure 2, the vaseplots will not adequately display the multimodality. Also, for distributions with regions of very low density between modes, HDR boxplots allow outliers to be present between the modes, unlike vaseplots or any other boxplot variant. Finally, unlike HDR boxplots, vaseplots do not extend naturally to allow continuous conditioning or to higher dimensions.

2.2 Conditional HDR's

The preceding example is easily extended to allow a more continuous display of the density of the maximum daily temperature conditional on the previous day's temperature. The conditional densities were calculated using a kernel approach as described in Hyndman, Bashtannyk, and Grunwald (in press), and the HDR's conditional on the previous day's temperature being 7, 8, ..., 43 were computed. The results are displayed in Figure 3. This display avoids the artificial grouping of Figure 2 and allows the gradual change in the shape of the conditional densities to be seen more clearly.

In forecasting it is common to compute the density of future values conditional on observed values of the series and the time horizon of the forecast. In this case the HDR's may be plotted against time.

Example. A famous dataset in nonlinear time series analysis is A. J. Nicholson's blowfly data. These consist of the number of Australian blowflies recorded every two days in a caged population on a strictly controlled diet (Nicholson 1957). The data are also given in Brillinger, Guckenheimer, Guttorp, and Oster (1980). Of particular interest are the aperiodic population cycles. Tong (1988) fitted a threshold autoregressive model to the log of the population. We are interested in forecasting the process for the next 20 observation times. Hyndman (1995) used a Monte Carlo technique to obtain forecast densities of the logged population at each time. These densities were then transformed back

to the original scale. The highest density forecast regions were obtained from the densities on the original scale.

Figure 4 shows the 50% and 95% HDR's for these distributions. The means of the forecast densities are plotted as a solid line. The HDR's clearly show positive skewness and bimodality from time 203. The left hump becomes narrower until time 207 and then widens again, so that by time 210, the humps corresponding to each mode have become sufficiently small that the 50% HDR contains both modes. The bimodality indicates uncertainty in the period of the next population cycle: the lower mode corresponds to the possibility of a minimum in the cycle, and the larger mode corresponds to a possible maximum in the cycle. Standard forecast regions, usually obtained from quantiles, are unable to show this bimodality.

2.3 HDR's for Multivariate Densities

It follows immediately from the definition that the boundary of an HDR consists of those values of the sample space with equal density. Hence a plot of a bivariate HDR is simply a form of contour plot. A bivariate HDR boxplot may be constructed using the 50% HDR and 99% HDR, with points lying outside the 99% HDR displayed as in a scatterplot. The mode is marked by a circle ($\circ$).

Example. Azzalini and Bowman (1990) examine duration and waiting times for eruptions from the Old Faithful geyser in Yellowstone National Park, Wyoming. The data were collected continuously from August 1 until August 15, 1985. Several similar data sets have also been analyzed: Weisberg (1985) and Denby and Pregibon (1987) consider data collected in 1978 and 1979, and Cook and Weisberg (1982) consider data collected in 1980. These authors mainly consider the regression relationship between the waiting time between eruptions and the previous eruption duration.

The focus in this example is the bivariate density of the duration of each eruption and the duration of the previous eruption. There are 299 observations; the times are measured in minutes. Figure 5 shows a bivariate HDR boxplot of the data based on a kernel estimate of the density. This is similar to a plot in Scott (1991) showing the contours of an estimate of the analogous density for the 1978 data.

Figure 5 shows that eruptions tend to be either long (around 4 minutes) or short (around 2 minutes), but rarely of medium length (around 3 minutes). Furthermore, there are never consecutive short eruptions. These observations are not new; previous studies of Old Faithful data have noted similar phenomena, and Azzalini and Bowman present a tentative physical model for these eruption patterns.

The striking features of the Old Faithful data are much more obvious from the HDR boxplot than from other bivariate boxplots that have been proposed. Becketti and Gould (1987) proposed a form of bivariate boxplot which simply superimposes the boxplots of the marginal variables on the scatterplot. As such, it fails to capture any correlation, let alone the unusual structure seen in Figure 5. Goldberg and Iglewicz (1992) propose bivariate boxplots in which both the inner region and outer region are constructed using four quarter-ellipses. Although their boxplot will show the correlation and possibly the lack of data in the lower left of the plot, it cannot reveal the presence of the three modes because both regions must be convex. If the data are from a unimodal density, then the boxplots of Goldberg and Iglewicz (1992) will yield very similar results to HDR boxplots. Hence HDR bivariate boxplots are suitable for a greater variety of data than the alternative bivariate boxplots.

In three dimensions a highest density region is a shell in three-dimensional space. If the density is trivariate normal, the HDR's are nested hyperellipsoids. David Scott (Scott 1991, 1992) has proposed a similar concept that he calls an α-shell defined by the surfaces $S_\alpha(x, y, z) := \{(x, y, z) : f(x, y, z) = \alpha f(\text{mode})\}$. In fact, these are HDR's under a different parameterization to that used in this article.

For d-dimensional densities where $d > 3$ Scott proposes plotting a three-dimensional shell after conditioning on $d - 3$ of the variables. Computer animation is also possible where the probability coverage is changed through time or, for $d > 3$, one of the conditioned variables is changed through time.

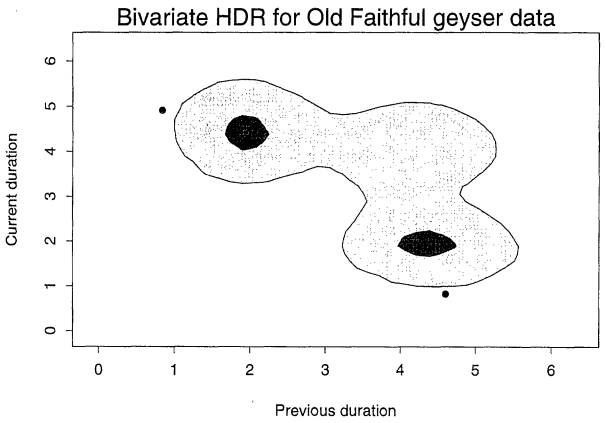

Figure 5. A Bivariate HDR Boxplot for the Old Faithful Geyser Data. The duration of each eruption is on the vertical axis; the duration of the previous eruption is on the horizontal axis. Times are measured in minutes.

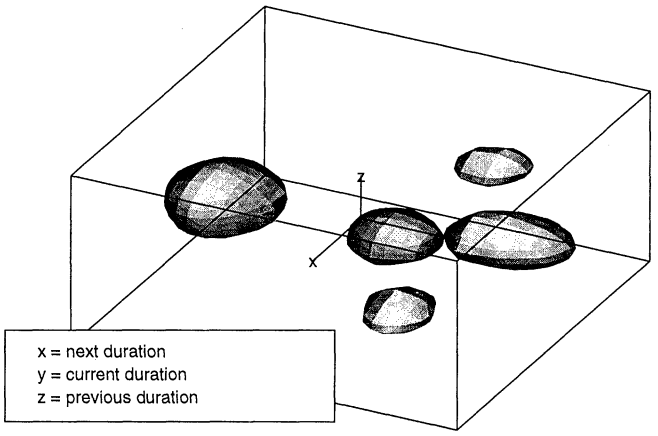

Figure 6. A 70% Trivariate HDR for the Old Faithful Geyser Data. The density was computed from the duration and its two lagged variables.

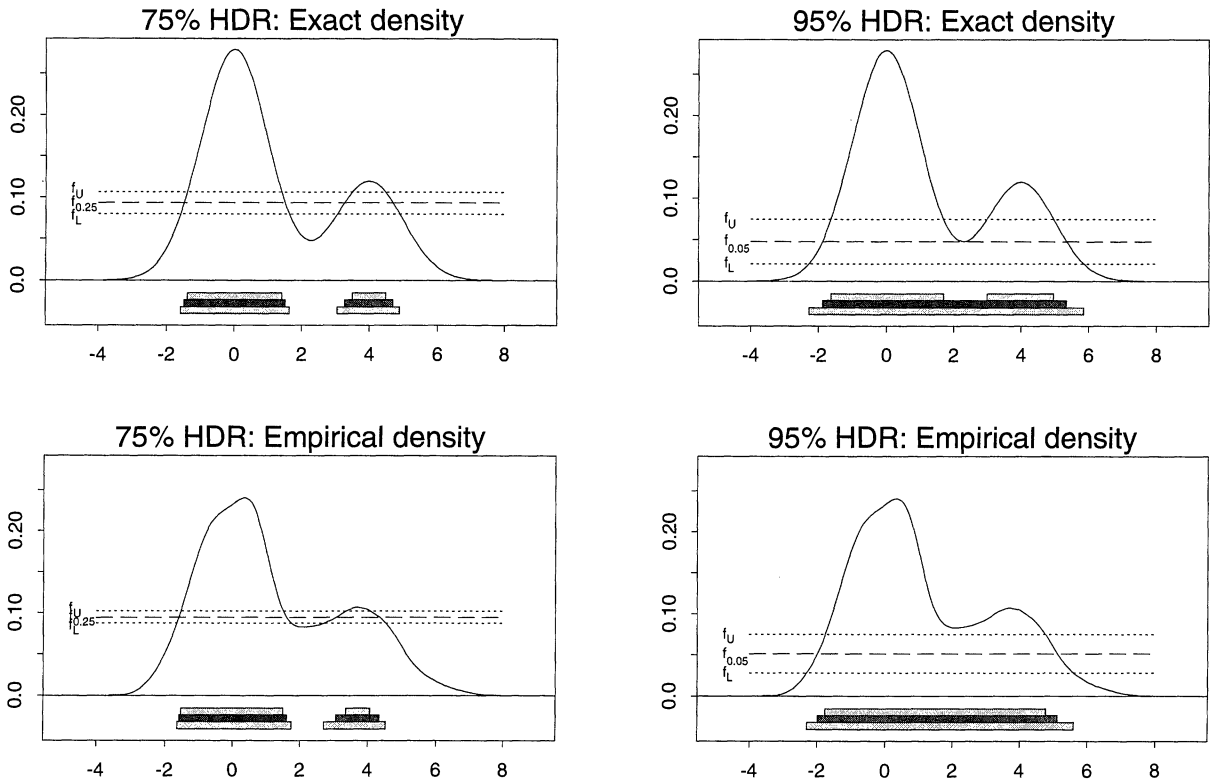

Figure 7. 50% and 75% HDR's for the Mixture Density of Figure 1 With 95% Confidence Regions.

Example. Figure 6 shows the 70% HDR from the trivariate density computed from the duration of the Old Faithful geyser data and its first two lagged variables. If L denotes a long duration and S a short duration, then the five shells correspond to the situations SLS, SLL, LSL, LLS, and LLL. The situations SSL, LSS, and SSS do not occur as two short eruptions are never consecutive.

This example does not reveal any new information about the geyser data that was not apparent in Figure 5. However, it shows clearly the value of HDR plots in identifying clusters in three-dimensional data. Using a three-dimensional scatterplot, even with spinning, it is very difficult to spot the five clusters because they tend to overlap when projected onto a two-dimensional plane.

3. CALCULATION OF HDR'S

For discrete-valued distributions, HDR's simply consist of those elements of the sample space with highest probability. Therefore computation is simple.

3.1 Numerical Integration Approach

For continuous distributions there have been several suggestions for constructing an HDR from a general univariate density $f(x)$, which is a bounded, continuous function of x. Wright (1986) proposed an algorithm involving numerical integration of $f(x)$, but he assumed the density was unimodal and so restricted the HDR to a single interval. Hyndman (1990) developed a more general algorithm that computed an HDR for any given density where $f(x)$ is a bounded, continuous function of x and the inverse of $f(x)$ is uniquely defined in the neighborhood of the boundary of the HDR. Both Wright and Hyndman use various numerical methods to improve the speed of computation. However, neither of these algorithms is easily generalized to multivariate densities. One major problem is the computational difficulty in numerically integrating over a general region in high-dimensional space. An alternative approach is required that avoids explicit integration.

3.2 Density Quantile Approach

Let $\mathbf{X}$ be a (possibly multivariate) random variable with density $f(\mathbf{x})$. Rather than use numerical integration under $f(\mathbf{x})$, we may obtain information about the probability coverage of a given region of the sample space using a Monte Carlo technique.

Consider the random variable, $Y = f(\mathbf{X})$, obtained by transforming $\mathbf{X}$ by its own density function. Now f_α is such that $\Pr(f(\mathbf{X}) \geq f_\alpha) = 1 - \alpha$. So f_α is the α quantile of Y. Therefore f_α can be estimated as a sample quantile from a set of iid random variables with the same distribution as Y.

Let $\{\mathbf{x}_1, \ldots, \mathbf{x}_n\}$ denote a set of independent observations from the density $f(\mathbf{x})$. Then $\{f(\mathbf{x}_1), \ldots, f(\mathbf{x}_n)\}$ is a set of independent observations from the distribution of Y. Let $f_{(j)}$ be the jth largest of $\{f(\mathbf{x}_i)\}$ so that $f_{(j)}$ is the (j/n) sample quantile of Y. We shall use $f_{(j)}$ as an estimate of f_α. Specifically, we choose $\hat{f}_\alpha = f_{(j)}$ where $j = \lfloor \alpha n \rfloor$. Then $\hat{f}_\alpha \to f_\alpha$ as $n \to \infty$, and so $R(\hat{f}_\alpha) \to R(f_\alpha)$ as $n \to \infty$.

If $f(x)$ is a known function, the observations will usually be generated pseudorandomly. The approximation can be made arbitrarily accurate by increasing n.

Often, however, we will not know the density $f(\mathbf{x})$, but will have a set of iid observations $\{\mathbf{y}_1, \ldots, \mathbf{y}_m\}$ from an unknown density. In this case $f(\mathbf{x})$ may be estimated empiri-

```
hdrboxplot <- function(x, y, a, b, ...)
# Input: x and y are independent observations on a density f(x,y).
# Call: density2d(x,y,x0,y0,a,b) returns a vector containing the
#       estimated density of (x,y) at the points x0,y0 using the
#       smoothing parameters a and b.
# Output: bivariate HDR boxplot
{
    fxy <- density2d(x, y, x, y, a, b)
    falpha <- quantile(fxy, c(0.01, 0.5))
    range.x <- diff(range(x))
    range.y <- diff(range(x))
    grid <- expand.grid(list(
          x = seq(min(x)-0.2*range.x, max(x)+0.2*range.x, length=20),
          y = seq(min(y)-0.2*range.y, max(y)+0.2*range.y, length=20)))
    fxy.grid <- density2d(x, y, grid$x, grid$y, a, b)
    junk <- contour(interp(grid$x,grid$y,fxy.grid),levels=falpha,
             xlab=deparse(substitute(x)), ylab=deparse(substitute(y)),
             labex=0, save=T, plotit=T, ...)
    polygon(junk[[1]]$x, junk[[1]]$y, col=4)
    polygon(junk[[2]]$x, junk[[2]]$y, col=2)
    index <- fxy < 0.999*falpha[1]
    points(x[index], y[index])
    index <- (1:length(x))[fxy==max(fxy)]
    points(x[index],y[index],pch="o")
}
```

Figure 8. Splus Code for Bivariate Data: Computes $\hat{f}_\alpha$ Using the Density Quantile Algorithm and Plots an HDR Boxplot.

cally from $\{\mathbf{y}_1, \ldots, \mathbf{y}_m\}$ (see, for example, Scott 1992). If m is large, there is no need to generate the sample $\{\mathbf{x}_1, \ldots \mathbf{x}_n\}$ as the actual observations $\{\mathbf{y}_1, \ldots, \mathbf{y}_m\}$ may be used directly by setting $n = m$ and $\mathbf{x}_i = \mathbf{y}_i$. If m is moderate, it may be preferable to generate "observations" $\{\mathbf{x}_1, \ldots, \mathbf{x}_n\}$ pseudorandomly from the density estimate. For small m it may not be possible to get a reasonable density estimate. Besides, with few observations and no prior knowledge of the underlying density, there seems little point in attempting to summarize the sample space. See Wand and Jones (1995) for some discussion on the number of observations needed for a reasonable density estimate. Note that the problem here is not with the density quantile algorithm (that will give results to an arbitrary degree of accuracy given a density), but with estimating the density from insufficient data.

A crude form of the density quantile algorithm can also be found in Tanner (1993, pp. 70–71) in the context of a posterior density from a data-augmentation algorithm.

The density quantile algorithm was used to estimate the HDR's displayed in Section 2.

3.3 Confidence Regions for Estimated HDR's

Because the density quantile algorithm involves a statistical approximation, it is helpful to compute some uncertainty limits on the estimated regions. Here I only consider the case where $\mathbf{X}$ is univariate.

We first obtain the distribution of $Y = f(X)$. Let $\{z_i\}$ denote those points in the sample space of X such that $f(z_i) = y, i = 1, 2, \ldots, n(y)$. That is, $\{z_i\}$ denote the endpoints of the subintervals that make up $R(y)$. Also define $A(y) = \int_{R(y)} f(u)\,du$ so that $A(f_\alpha) = 1 - \alpha$. Then for small δ

$$A(y + \delta) = A(y) - \delta y \sum_{i=1}^{n(y)} |f'(z_i)|^{-1} + O(\delta^2)$$

and so

$$\frac{dA(y)}{dy} = -y \sum_{i=1}^{n(y)} |f'(z_i)|^{-1}.$$

Now $\Pr(Y \le y) = 1 - A(y)$. Therefore the density of Y is simply

$$g(y) = -\frac{dA(y)}{dy} = y \sum_{i=1}^{n(y)} |f'(z_i)|^{-1}.$$

Then, using the standard asymptotic results for a sample quantile (e.g., Cox and Hinkley 1974), we have that $\hat{f}_\alpha$ is asymptotically normally distributed with mean f_α and variance $\alpha(1-\alpha)/n[g(f_\alpha)]^2$. We can use this result to obtain a confidence interval for f_α, say $[f_L, f_U]$. Then $R(f_L)$ and $R(f_U)$ represent lower and upper confidence regions for $R(f_\alpha)$.

The top two plots of Figure 7 show 50% and 90% HDR's for the mixture density of figure 1. The estimate of f_α and a 95% confidence interval for $\hat{f}_\alpha$ are shown as horizontal lines. The HDR estimate is shown as the dark shaded region, and the values of $R(f_L)$ and $R(f_U)$ are shown with lighter shading. Here, $n = 200$ has been used to compute the HDR.

To assess the effect of using an empirical density estimate rather than the actual density, $f(x)$, analogous plots are shown in the bottom half of Figure 7. Here, the density was estimated using a kernel estimator from the 200 observations. The additional variability in the HDR estimate due to the density estimation seems to be relatively small.

3.4 Computational Details and Timing

The density quantile algorithm and the computation of confidence regions has been implemented in Splus. For example, Figure 8 gives code for producing a bivariate boxplot as displayed in Figure 5. Here, the data are used directly in the algorithm by setting $n = m$ and $\mathbf{x}_i = \mathbf{y}_i$. This function can be easily modified to produce other HDR's or to compute the density from a known function rather than use an empirical estimate.

Note that two calls must be made to the function `density2d`, one to compute the density at each of the observation points, and one to compute the density on a 20×20 grid. Linear interpolation is used between the grid points when computing the contour.

Most of the computational work in this function is in estimating the density using `density2d`. Hence the speed of the function depends largely on the speed of `density2d`. Using a relatively slow implementation of bivariate kernel density estimation, Figure 5 was produced using this function in 6.1 seconds of processor time on Splus 3.1 running on a DECstation 5000/25. With a faster implementation of density estimation, using the ideas of Fan and Marron (1994), this time should be able to be reduced substantially.

Ironically, it is more difficult to plot a univariate HDR than a bivariate HDR because there is no equivalent of the contour function in one dimension. Therefore a new function has been written for this purpose. Again, the density is computed (or estimated) on a grid of values; between the grid points, spline interpolation is used. Table 1 shows the processor times (in seconds) to compute an HDR and its associated 95% confidence regions using a DECstation 5000/25 running Splus 3.1. The mixture density of Figure

Table 1. Processor Times (in seconds) for Splus to Compute an HDR and Its Associated 95% Confidence Regions for the Mixture Density of Figure 1

	α			
n	.5	.25	.05	.01
10,000	1.38	1.64	1.45	1.36
1,000	1.01	1.28	1.19	1.02
200	1.09	1.25	.85	.76

1 was used and measurements were made for different values of n and α. Because the density is known in this case, kernel density estimation was not necessary.

The Splus code used to find HDR's and produce HDR boxplots in one and two dimensions can be obtained from the S archive of statlib@lib.stat.cmu.edu.

[Received April 1993. Revised October 1994.]

REFERENCES

Azzalini, A., and Bowman, A. W. (1990), "A Look at Some Data on the Old Faithful Geyser," *Applied Statistics*, 39, 357–365.

Becketti, S., and Gould, W. (1987), "Rangefinder Box Plots: A Note," *The American Statistician*, 41, 149.

Benjamini, Y. (1988), "Opening the Box of a Boxplot," *The American Statistician*, 42, 257–262.

Box, G. E. P., and Tiao, G. C. (1973), *Bayesian Inference in Statistical Analysis*, Reading, MA: Addison-Wesley.

Brillinger, D. R., Guckenheimer, J., Guttorp, P. E., and Oster, G. (1980), "Empirical Modelling of Population Time Series Data: The Case of Age and Density Dependent Vital Rates" (Vol. 13 of Lectures on Mathematics in the Life Sciences), Providence, RI: American Mathematical Society, pp. 65–90.

Cook, R. D., and Weisberg, S. (1982), *Residuals and Influence in Regression*, London: Chapman and Hall.

Cox, D. R., and Hinkley, C. V. (1974), *Theoretical Statistics*, London: Chapman and Hall.

Denby, L., and Pregibon, D. (1987), "An Example of the Use of Graphics in Regression, *The American Statistician*, 41, 33–38.

Esty, W. W., and Banfield, J. D. (1992), "The Box-Percentile Plot," Technical Report, Montana State University, Dept. of Mathematical Sciences.

Fan, J., and Marron, J. S. (1994), "Fast Implementation of Nonparametric Curve Estimators," *Journal of Computer Graphics and Statistics*, 3, 35–56.

Goldberg, K. M., and Iglewicz, B. (1992), "Bivariate Extensions of the Boxplot," *Technometrics*, 34, 307–320.

Hoaglin, D. C., Iglewicz, B., and Tukey, J. W. (1986), "Performance of Some Resistance Rules for Outlier Labeling," *Journal of the American Statistical Association*, 81, 991–999.

Hyndman, R. J. (1990), "An Algorithm for Constructing Highest Density Regions," unpublished manuscript available from the author.

——— (1995), "Highest Density Forecast Regions for Non-Linear and Non-Normal Time Series Models," *Journal of Forecasting*, 14, 431–441.

Hyndman, R. J., Bashtannyk, D. M., and Grunwald, G. K. (in press), "Estimating and Visualizing Conditional Densities," *Journal of Computational Graphics and Statistics*, 5.

McGill, R., Tukey, J. W., and Larsen, W. A. (1978), "Variations of Box Plots," *The American Statistician*, 32, 12–16.

Nicholson, A. J. (1957), "The Self-Adjustment of Populations to Change," in *Cold Spring Harbor Symposia on Quantitative Biology* (Vol. 22), pp. 153–173.

Scott, D. W. (1991), "On Estimation and Visualization of Higher Dimensional Surfaces," in *Computing and Graphics in Statistics*, eds. A. Buja and P. A. Tukey (IMA Vol. 36 in Mathematics and Its Applications), New York: Springer-Verlag.

——— (1992), *Multivariate Density Estimation: Theory, Practice, and Visualization*, New York: John Wiley.

Tanner, M. A. (1993), *Tools for Statistical Inference: Methods for the Exploration of Posterior Distributions and Likelihood Functions* (2nd ed.), New York: Springer-Verlag.

Tong, H. (1988), "Non-Linear Time Series Modelling in Population Biology: A Preliminary Case Study," in *Nonlinear Time Series and Signal Processing* (Vol. 106 of Lecture Notes in Control and Information Sciences), New York: Springer-Verlag, pp. 75–87.

Tukey, J. W. (1977), *Exploratory Data Analysis*, Reading, MA: Addison-Wesley.

Wand, M. P., and Jones, M. C. (1995), *Kernel Smoothing*, London: Chapman and Hall.

Weisberg, S. (1985), *Applied Linear Regression* (2nd ed.), New York: John Wiley.

Wright, D. E. (1986), "A Note on the Construction of Highest Posterior Density Intervals," *Applied Statistics*, 35, 49–53.