

Interface Foundation of America

Variable Resolution Bivariate Plots Author(s): Chisheng Huang, John Alan McDonald, Werner Stuetzle Source: Journal of Computational and Graphical Statistics, Vol. 6, No. 4 (Dec., 1997), pp. 383-396 Published by: American Statistical Association, Institute of Mathematical Statistics, and Interface Foundation of America Stable URL: <u>http://www.jstor.org/stable/1390742</u> Accessed: 24/09/2009 18:09

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at http://www.jstor.org/page/info/about/policies/terms.jsp. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at http://www.jstor.org/action/showPublisher?publisherCode=astata.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is a not-for-profit organization founded in 1995 to build trusted digital archives for scholarship. We work with the scholarly community to preserve their work and the materials they rely upon, and to build a common research platform that promotes the discovery and use of these resources. For more information about JSTOR, please contact support@jstor.org.



American Statistical Association, Institute of Mathematical Statistics, Interface Foundation of America are collaborating with JSTOR to digitize, preserve and extend access to Journal of Computational and Graphical Statistics.

Variable Resolution Bivariate Plots

Chisheng HUANG, John Alan MCDONALD, and Werner STUETZLE

Scatterplots are the method of choice for displaying the distribution of points in two dimensions. They are used to discover patterns such as holes, outliers, modes, and association between the two variables. A common problem is overstriking—the overlap on the plotting surface of glyphs representing individual observations. Overstriking can create a misleading impression of the data distribution. The variable resolution bivariate plots (*Varebi plots*) proposed in this article deal with the problem of overstriking by mixing display of a density estimate and display of individual observations. The idea is to determine the display format by analyzing the actual amount of overstriking on the screen. Thus, the display format will depend on the sample size, the distribution of the observations, the size and shape of individual icons, and the size of the window. It may change automatically when the window is resized. Varebi plots reveal detail wherever possible, and show the overall trend when displaying detail is not feasible.

Key Words: Histogram; Scatterplot; Visualization.

1. INTRODUCTION AND MOTIVATION

Scatterplots are the method of choice for visualizing the distribution of points in two dimensions. They are used to discover patterns such as holes (areas with few data points), outliers, modes, association between the two variables, and so on.

A common problem with scatterplots is overstriking—the overlap on the plotting surface of glyphs representing individual observations. Overstriking can create a misleading impression of the data distribution. As an example, consider Figure 1. This scatterplot was produced during the analysis of data on a colony of Magellanic penguins in Punta Tombo, Argentina. The observations are penguin nest sites. The variables are ground composition at a site (percentage of sand) on the vertical axis and vegetation coverage (percentage of bare ground) on the horizontal axis.

The *displayed* point pattern suggests the mode to be within the rectangle defined by x = (30, 80) and y = (0, 15). Figure 2, however, was obtained by binning the points into a 20×20 grid and encoding the counts into gray levels, revealing that the visual impression is misleading and that there is a mode in the lower right corner of the plot.

Chisheng Huang is Assistant Research Fellow, Institute of Statistical Science, Academia Sinica, Taipei, Taiwan 115, R.O.C.; e-mail: cph@delphis.stat.sinica.edu.tw. John McDonald is Research Associate Professor, and Werner Stuetzle is Professor, Department of Statistics, Box 354322, Seattle, WA 98195-4322; e-mail: jam@stat.washington.edu; wxs@stat.washington.edu.

^{©1997} American Statistical Association, Institute of Mathematical Statistics, and Interface Foundation of North America

Journal of Computational and Graphical Statistics, Volume 6, Number 4, Pages 383-396



Figure 1. Where is the Mode?

Overstriking is obvious if the data set is so large or the plotting surface so small that we just see a solid blob of ink. (We use the term "ink" although most plots today are drawn on a computer screen). However, as our example illustrates, it can be a problem even with small data sets (hundreds or thousands of observations), where it is difficult to notice and therefore more insidious.

A way of coping with the problem of overstriking is to abandon the idea of drawing individual points and instead draw *agglomerative glyphs* representing collections of points. We can, for example, bin the drawing area into rectangular or hexagonal bins and compute a two-dimensional histogram (Carr, Littlefield, Nicholson, and Littlefield 1987). The histogram can be drawn as a perspective plot, or we can encode the counts in gray scale (as in Fig. 2) or glyph size (Fig. 3).

Besides force of habit and inertia there are at least two other arguments against routinely replacing scatterplots by two-dimensional histograms:

• The discretization inherent in a histogram smears out fine structure. As an illustration, Figure 4 (a) shows a snapshot of 200 rotating 3D points whose coordinates were generated by the infamous RANDU (Knuth 1981) random number generator. See Tierney (1990) for the particular implementation that we used. Figure 4



Figure 2. Same Data as Figure 1. 2-D histogram drawn by encoding bin counts in gray scale.

(b), (c), and (d) show histograms of the same 200 points, with different numbers of bins. None of the histograms reveals RANDU's lack of randomness as clearly as the scatterplot.

• Encoding counts into gray level or glyph size requires specification of a mapping. If we want to judge on how many data points a feature of the histogram, like an apparent mode, is based, we need to mentally invert this mapping. This process is not immediate, and we want to avoid it whenever possible.

In this article we present two suggestions for dealing with the problem of overstriking in scatterplots.

- 1. We mix individual and agglomerative glyphs in the same plot.
- 2. We choose between drawing individual and agglomerative glyphs by analyzing the actual amount of overstriking on the screen. The display format thus will depend on the sample size, the distribution of points, the size and shape of the individual glyphs, and the size of the drawing area.

Adapting the type of display to the size of the drawing area is particularly helpful and effective when plots are displayed in windows on a screen, where they can be shrunk to free up space and expanded again for closer inspection.

As an illustration, Figures 5 and 6 show the same data set as Figures 1 and 2 for different sizes of the drawing area. Notice that the areas around (65, 16) and (100, 16)



Figure 3. Same Data as Figure 1. 2-D histogram drawn by encoding bin counts in rectangle size.

in Figure 5 display an agglomerative glyph and switch to displaying point glyphs in the bigger Figure 6.

The rest of this article is organized as follows. Section 2 describes in detail how Varebi plots are drawn. Section 3 contains additional examples. A discussion concludes the article.

2. DESCRIPTION OF VAREBI PLOTS

We will first discuss Varebi plots on black-and-white displays and then describe a variation designed for gray level displays. The gray level version can convey information better on the data distribution when the drawing space is very small.

2.1 VAREBI PLOTS ON BLACK-AND-WHITE DISPLAYS

Drawing a Varebi plot involves a sequence of steps.



Figure 4. 200 Points Generated by RANDU. Scatterplot (a) and 2-D histograms with 5×5 (b), 10×10 (c), and 20×20 (d) bins.

2.1.1 Step 1—Binning

We bin the data points into a regular grid with a default resolution of 20×20 . This results in bins with a side length on the order of 15 pixels for a 300×300 window, a size that we frequently observed ourselves using. A bin size of 15×15 pixels is large enough to allow drawing of agglomerative glyphs in at least five visually distinguishable sizes.

Choosing a fixed number of bins, independent of the plot size on the screen, has the advantage that the plot changes in predictable ways when the window is reshaped. Enlarging the window may change some bins from showing agglomerative glyphs to showing individual observations, but never vice-versa. The analogous statement is true for reduction of the window size.



Figure 5. Same Data as Figure 1. Plot size on a computer screen is 313×334 pixels.

2.1.2 Step 2—Application of a Transfer Function

The basic tenet of binary Varebi plots is that the amount of ink actually deposited in a bin should be a non-decreasing function of the number of observations in the bin (up to some tolerance, as described in Step 3). Let s be the bin area on the screen (in pixels), n_{\max} the largest bin count, and τ the size of an individual point glyph (in pixels). (On many displays single pixel dots are hard to see, requiring $\tau > 1$.) A non-decreasing function, $T : [1 : n_{\max}] \rightarrow [1 : s]$, that assigns an amount of ink to each bin count is called a *transfer function*.

For small bin counts T has to be linear with slope τ , because we want to draw individual observations in bins without overstriking. However, if $\tau \times n_{\text{max}} > s$, the simple choice $T_l(n) = \tau \times n$ is not feasible. In this case we have to "blunt" T—that is, flatten it out for counts above some cutoff n_{crit} . There are many possible ways of doing so: any non-decreasing function T will do as long as $T(n) = \tau \times n$ for $0 \le n \le n_{\text{crit}} \le \lfloor s/\tau \rfloor$, and $T(n_{\text{max}}) = s$. We chose a simple one, shown in Figure 7. If we have to use the blunted transfer function T_b , we draw agglomerative glyphs in all bins with bin count $n_{ij} > n_{\text{crit}}$, in order to avoid violating the monotonicity condition.



Figure 6. Same Data as Figure 1. Plot size on a computer screen is 384×397 pixels.

It remains to discuss the choice of n_{crit} . Clearly, n_{crit} should increase when the size of the plot and therefore the bin area s increases—we want to use additional screen space to improve the resolution of the plot—that is, draw more individual glyphs. We set $n_{\text{crit}} = \lfloor s^2/(\tau^2 n_{\text{max}}) \rfloor$. The motivation for this choice is that, as $s \to \tau \times n_{\text{max}}$, the blunted transfer function T_b approaches T_l . This makes for a smooth transition of display format when the drawing space increases to the point where we can switch to the linear transfer function.

2.1.3 Step 3—Calculating the Overplotting Index and Drawing

The final step in producing a Varebi plot is to decide, for each bin, whether to draw individual glyphs or an agglomerative glyph. Let a_{ij} be the amount of ink (number of black pixels) in the corresponding bin on the screen if we simply drew the glyphs without paying any attention to overplotting. We measure the amount of overplotting in a bin by the overplotting index

$$o_{ij} = \frac{\tau \times n_{ij}}{a_{ij}}.$$



Figure 7. Blunted Transfer Function.



Figure 8. Same Data as Figure 1. Varebi plot (a) and scatterplot (b) of screen size 205×212 pixels. (Figure not printed to scale.)



Figure 9. Same Data as Figure 1. Varebi plot of screen size 480×485 pixels (a). Scatterplot of screen size 480×485 pixels (b). (Figure not printed to scale.)

If we can use the linear transfer function T_l , we draw agglomerative glyphs for all bins for which the overplotting index o_{ij} is greater than some threshold o_{crit} . If we have to use the blunted transfer function T_b , we draw agglomerative glyphs for all bins with $n_{ij} > n_{crit}$ or $o_{ij} > o_{crit}$. In the examples, $o_{crit} = .35$. An agglomerative glyph is drawn as closely as possible to the center of the mass within each bin while keeping the glyph completely inside the bin. This has the additional benefit of breaking the artificial regularity imposed by the grid (Carr et al. 1987).

2.2 VAREBI PLOTS ON GRAY LEVEL DISPLAYS

Varebi plots will lose their ability to convey an accurate image of the data distribution when the size of the plotting area gets so small that there are only a few possible values for the size of an agglomerative glyph. If the display can show gray levels we can, however, represent bin count by gray scale instead of glyph size.

Let n_{\min} be the smallest bin count among all the bins with significant overstriking that is, with $o_{ij} > o_{crit}$. Bins with $n_{ij} \ge n_{\min}$ are drawn in gray. The gray level is determined by a nondecreasing function $\mathcal{D} : [n_{\min}, n_{\max}] \rightarrow [d_1, d_2]$, where d_1 is lightest and d_2 is darkest. All bins with $n_{ij} < n_{\min}$ display individual observations. Therefore, gray level bins always have a higher count than bins showing individual observations. This is necessary because it is impossible to visually establish an ordering between amount of (black) ink in an area, and gray level.

In the examples we switch to encoding counts by gray level instead of glyph size whenever either bin width or bin height is less than 10 pixels. We do not use gray scale encoding otherwise because we want as many bins as possible to display point glyphs, because perception of gray level is affected by the surrounding area (Foley, van Dam, Feiner, and Hughes 1990), and because comparing gray levels is not easy, especially for objects positioned far apart.



Figure 10. 400 points generated by RANDU. Varebi plot (a) and 2-D histogram (b).



Figure 11. 1,600 Points Generated by RANDU. Varebi plot (a) and 2-D histogram (b).

2.3 HEXAGONAL VERSUS RECTANGULAR BINS

The Varebi plots were implemented using rectangular bins, rather than hexagonal bins as advocated by Carr et al. (1987). Hexagonal bins give slightly better density estimates and result in more eye pleasing displays because they de-emphasize horizontal and vertical directions. We chose rectangular bins for reasons of speed and ease of implementation.

Rectangular bins also have an advantage when Varebi plots are used in conjunction with scatterplot brushing, or when we wish to encode an additional categorical variable in color. In this situation it is not clear how to divide a hexagonal glyph into colored parts, whereas we can convert rectangles into divided color bars.



Figure 12. 6,400 Points Generated by RANDU. Varebi plot (a) and 2-D histogram (b).



Figure 13. Scatterplot of 65,536 Points from a 256×256 MRI Image.

3. VAREBI PLOT EXAMPLES

We present three sets of examples:

- Figures 8a, 9a and Figures 8b, 9b show Varebi plots and conventional scatterplots, respectively, of the same data displayed in Figure 1, for different window sizes. Overstriking in this example is quite serious, so that even a large plot like the one shown in Figure 9, which would take up a quarter of a $1,000 \times 1,000$ screen, displays many agglomerative glyphs.
- Figures 10–12 show RANDU samples of sizes ranging from 400 to 6,400. Varebi plots are on the left; two-dimensional histograms on the right. Note that, in contrast to the bivariate histograms, all Varebi plots clearly reveal the lack of randomness.
- Figures 13–15 show 65,536 pairs of measurements from a 256 × 256 MRI (magnetic resonance imaging) image. The MRI data set has an enormous peak (38,455 points) at (0,0), corresponding to the background of the image. The second highest peak corresponds to 70 points. The scatterplot in Figure 13 fails to reveal the peak at the origin. The Varebi plot in Figure 14 shows the enormous peak, but has not much resolution anywhere else. This is hard to avoid because the dynamic range is so large. Figure 15 shows a Varebi plot with the background pixels removed. This plot gives a much clearer picture of the rest of the data.



Figure 14. Varebi Plot of 65,536 Points from a 256×256 MRI Image.

4. DISCUSSION

There are a number of existing techniques attempting to deal with the problem of overstriking in scatterplots. They tend to fall into one of two categories—those that display individual data points and those that display a density estimate.

Examples for techniques in the first category are *jittering* (Chambers, Kleiner, and Tukey 1983) and use of unfilled circles as plotting symbols (Cleveland 1985).

Jittering was originally proposed to alleviate overstriking in plots of a euclidean variable against a categorical variable. In scatterplots, points obscured by overplotting could be offset by a small random displacement. (Of course, it would be necessary to remind viewers that jittering was employed in producing the plot.)

Using unfilled circles as plotting symbols helps as long as there are not too many exact ties. The intersection of unfilled circles gives a geometric shape distinctly different from a circle. Intersections of axis-parallel rectangles, on the other hand, again are rectangles, which makes it hard to differentiate observations.

Those techniques clearly break down if the density of points in the drawing area gets too large.

Techniques displaying a density estimate typically use a two-dimensional histogram with rectangular or hexagonal bins. The histogram can be displayed as a three-dimensional perspective plot or a contour plot. Alternatively, bin counts can be encoded into size or



Figure 15. Varebi Plot of 27,081 Non-Background Points from a 256 × 256 MRI Image.

gray level of glyphs representing the bins. Cleveland and McGill (1984) proposed representing bins by sunflowers. The number of petals in a sunflower encodes the bin count. If the bin count is 1, the sunflower degenerates to a point. Sunflowers can also be used with unbinned data when exact overplotting occurs.

These techniques result in an unnecessary loss of detail in regions of the drawing area with low point density, which could be the entire area if it is large, or if the sample size is small.

Carr et al. (1987) proposed a technique using both individual and agglomerative glyphs. They compute a two-dimensional histogram density estimate with hexagonal bins, encode bin count into the size of hexagons drawn into bins with four or more observations, and draw individual glyphs in bins with three or fewer observations. This technique falls in between the two categories.

The Varebi plots proposed and illustrated in this article mix agglomerative and individual glyphs, display of a density estimate, and display of individual observations. The novel idea is to determine the display format by analyzing the actual amount of overstriking on the screen. Thus, the display format will depend on the sample size, the distribution of the observations, the size and shape of the individual glyphs, and the size of the window. It may automatically change when the window is resized. Varebi plots do not suffer from the shortcomings of techniques using a fixed display format, and they can be used on binary displays. They reveal detail wherever possible, and show the overall trend when displaying detail is not feasible.

[Received November 1995. Revised March 1997.]

ACKNOWLEDGMENTS

This work was supported by Department of Energy grant DE-FG06-85-ER25006 and National Science Foundation grant DMS-9114027.

REFERENCES

- Carr, D.B., Littlefield, R.J., Nicholson, W.L., and Littlefield, J.-S. (1987), "Scatterplot Matrix Techniques for Large n," Journal of the American Statistical Association, 82, 424–436.
- Chambers, J.M., Cleveland, W.S., Kleiner, B., and Tukey, P.A. (1983), *Graphical Methods for Data Analysis*, Monterey, CA: Wadsworth.

Cleveland, W.S. (1985), The Elements of Graphing Data, Monterey, CA: Wadsworth.

- Cleveland, W.S., and McGill, R. (1984), "The Many Faces of a Scatterplot," *Journal of the American Statistical Association*, 79, 807–822.
- Foley, J.D., van Dam, A., Feiner, S.K., and Hughes, J.F. (1990), *Computer Graphics: Principles and Practice*, Reading, MA: Addison-Wesley.
- Knuth, D.E. (1981), Seminumerical Algorithms, volume 2 of The Art of Computer Programming, Reading, MA: Addison-Wesley.

Tierney, L. (1990), Lisp-Stat, New York: John Wiley.