

Diagnostic Plots for One-Dimensional Data¹

G. Sawitzki
StatLab Heidelberg
Im Neuenheimer Feld 294
D 6900 Heidelberg

Summary: How do we draw a distribution on the line ? We give a survey of some well known and some recent proposals to present such a distribution, based on sample data. We claim: a diagnostic plot is only as good as the hard statistical theory that is supporting it. We try to illustrate this point of view for some examples.

Though the general contribution of diagnostic plots to statistics is accepted, sometimes diagnostic plots seem more of a fashion than a tool. There are uncountable possibilities to design diagnostic plots, not all being of equal use. Diagnostic plots can and should be judged the same way as any other statistical method. We have to ask: What is their power ? What is their reliability ? While we may have to stay with accidental notes or examples for some time now, in the end a diagnostic plot is only as good as the hard statistical theory that is supporting it.

For many diagnostic plots, we are still far from having this theory. For some plots, we have to ask: What precisely are they trying to diagnose ? How do we judge their reliability or confidence ? For other plots, we know at least the statistical methods they are related to. In this paper, we consider plots as *views* on probability measures: We relate plots to functionals, operating on probability measures. If we have a functional defining a plot, we can proceed in three steps. We can ask which features are exhibited by the functional, and which are collapsed. As a second step, we can analyze what is retained by the empirical version and which stochastic fluctuation is to be expected. Third, we can optimize the functional and its empirical version to gain maximal power.

Where possible, we try to indicate classical tests related to the plot. If these tests meet the core of the plot, the power of the plot may be identified with and judged by the power of these tests. The associated functional may even indicate a notion of distance, or a metric, associated with the plot. We can use this to find natural neighbourhoods of a given empirical plot, leading to confidence sets of compatible models.

We restrict ourselves to a very modest case: Assuming a continuous distribution on the real line, we look at diagnostic plots based on a sample from this (unknown) distribution. We exclude some of the more difficult questions: we assume independent sample points with identical distribution. So we do not look at plots for the diagnosis of dependency, trend, heteroskedasticity or other inhomogeneities. We give a survey of some diagnostic plots, pointing to their related statistical methods. As is to be expected, the chance is taken to advocate some new plots: the silhouette, the densitogram, and the shorth plot.

Diagnostic plots, what do we need them for ?

We use diagnostic plots to investigate a data set by itself (a descriptive problem), or in comparison to a model distribution, or family (the one-sample-problem), or to compare two data sets (the two-sample-problem).

In a classical framework, we may want to apply a certain method, like regression or analysis of variance. This method will depend on certain conditions, often on conditions which in principle cannot be verified. The role of a diagnostic plot is that of a detector: Since we are unable to verify the preconditions, we may use diagnostic plots to have at least a warning instrument.

¹Presented at the 24th meeting of the Arbeitsgruppe "Computational Statistics", Internationale Biometrische Gesellschaft (DR), Reisenburg 1992 (21.-24. Juni 1992)

In other situations, we may still be exploring. We have not settled on a specific model or method, but are looking at what the data are telling us. In the next step we select a certain model or decide to apply a certain method. Here diagnostic plots are a means to navigate through the models or methods at our disposition, and should be considered a model selection tool.

In any case diagnostic plots could be considered in a decision framework, either as filtering out bad situations after applying a model/method, or as selecting a model/method to be applied afterwards. It would be most appropriate to judge diagnostic plots as one step in an analysis process. But still too little is known about the interplay between use of diagnostic plots and application of formal models/methods.

The data we feed into diagnostic plots are rarely raw data. Often we use diagnostic plots on residuals. Of course the conditions we have to check refer to the error terms. The residuals are only some (model dependent) estimators of the errors. So the true story will be more complex than the i.i.d. simplification told here. Where no model dependence is included initially, we still have had some choice how to measure the data: what we consider to be the data is a result of our choice of a measurement process. This can be a practical choice, or this may be culture dependent. Even in simple examples it may be more than just a linear change of scale (for example, energy consumption in a car is measured in miles per gallon in the U.S.A., and as litre per 100 km in Europe). Sometimes detection of the "proper" scale is the major achievement. For example the Weber-Fechner law in psychology tells us that the amount of energy which must be added to a stimulus to produce a detectable difference is proportional to the energy level of the stimulus. Hence using a logarithmic scale may be more appropriate for perception experiments than a linear energy scale. Choice of (nonlinear) scale may be a major application of diagnostic plots. Identifying the shape of a distribution is equivalent to finding a way to transform it to some model distribution.

What do we look for in diagnostic plots ?

We use diagnostic plots to check for special features revealed by or inherent in the data. Of course, these checks are useful only if we know how presence or absence of these features affects the statistical methods we are going to apply. But then, if we do not know this, it would be wise not to apply these methods at all. Usually, for ordinary statistical applications, there are only few features we have to check (remember that we assume an i.i.d situation, so we do not look at plots for the diagnosis of dependency, trend, heteroscedasticity or other inhomogeneities). Here is a short check list:

- Missing or censored data.
Contrary to what classical statistics would like to see, real data sets usually contain registrations meaning "below detection level", "not recorded", "too large". In survival analysis, respecting missing or censored data is a mark of the trade. Although missing or censored data are a pending challenge in practical statistics, we will not deal with this problem here.
- Discretization.
Usually, all data we record are discretized (truncated or rounded to some finite precision, for example). For methods based on ranks, this may lead to ties, with appropriate corrections being well known. For other methods, these effects are often grossly ignored, although it would be easy to take them into account in tests of t or F type.
- Multimodality.
Sometimes, multimodality is a hint to a factor which separates the modes and should be included in the analysis. In other cases, as for instance in psychological preferences and choices, multimodality may be an inherent feature. Classical methods have notorious pitfalls if multi-modal distributions are involved.
- Symmetry and skewness.
In best cases, skewness is an indicator for power transformations which might bring the data to a simpler model.

- Tail behaviour.
Many classical methods are strongly affected by the tail behaviour of the distribution. Sometimes, tail problems may be avoided by going to more robust methods.

Diagnostic plots, what are they, anyway ?

The general aim of data analysis is to find interesting features in data, and to bring them to human perception. In doing this, data analysis has to avoid artifacts coming from random fluctuation, and from perception (Sawitzki, 1990). Diagnostic plots are plots tuned to serve these purposes. It depends on the context and on our intentions to say what is an interesting feature. For a general discussion, we have to ask which features can be brought to perception by a certain plot.

Full information is contained in the graph of, say, the probability density. We can easily get information about the relative location of the means, or about the standard deviations, or many other details from a plot of the probability density (Figure 1).

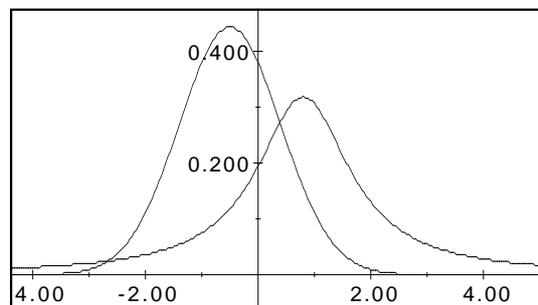


Figure 1: Densities. Differences of the means or standard deviations can be read easily from a plot of the density, but the density plot may as well be grossly misleading.

Although the full information is contained in this plot, it may not always come to our perception. In Figure 1, for example, one density is that of a Cauchy distribution. But the plot does not call our attention to the tail behaviour. It does not tell us that estimating the difference of the means may not be a good idea here, and it does not tell us that attempts such as studentization will run astray.

Perception is one side of the problem, and many discussions between specialists in this field and statisticians still may be necessary. There is standard literature addressing these problems: Bertin (1967, 1977, 1981) and Tufte (1990) are rich sources of possibilities. Tufte (1983) and Chambers et al. (1983) is the basic literature from a statistical point of view. Wainer, H. (1984) is a classical article on pitfalls to be avoided.

Some of elementary lessons we have learned:

- perception is background dependent. Avoid chart junk and background/foreground interaction.
- visual discrimination is powerful for linear or regular structures, but weak for general curves. If you have a model case, try to represent it by a straight line or a regular structure.
- perception knows quality and quantity. Avoid encoding quantitative information by qualitative features (such as colour).
- perception has more dimension than one. Make sure the information you are presenting is encoded in appropriate dimensions. In particular: avoid using 3d-effects, unless you know exactly how they are perceived.

It is left as an exercise to look for examples in which these elementary lessons are disregarded.

Here we will concentrate on the statistical side: we ask for the features (or functionals) of the distribution being represented, and for the fluctuation involved. The Cauchy example given above should be a warning:

even in the absence of fluctuation, a plot may not tell the whole story. Reducing fluctuation, or only minimizing it (by some optimal choice of parameters such as a bandwidth) by itself is not a guarantee for the usability for a certain purpose.

Notation and conventions

We assume a distribution F on the real line with density f and look at diagnostic plots based on a sample X_1, \dots, X_n from the (unknown distribution) F . We assume that X_1, \dots, X_n are independent sample points.

By $X_{(i:n)}$ or $X_{(i)}$ for short, we denote the i th order statistics, and F_n is the empirical distribution function with $F_n(X_{(i:n)}) = i/n$. For any plot expressed in terms of F , the empirical version is the corresponding plot with F_n replacing F . For simplicity, we identify distribution function and measure, allowing us to write $F(a, b] = F(b) - F(a)$.

For any plot we try to follow this sequence: We give a rough sketch how to generate the plot. Then we try to give a functional definition of the plot. Which features are preserved and which are lost by the functional? After that we study the fine points: what needs to be corrected in the rough plot? Then we turn to related tests and discuss optional choices. After each plot, we ask: how far have we got? What is the information we can gain so far, and what is still missing?

Histogram

Recipe: Chose histogram bins. For any bin, mark the hit count of data points hitting this bin.

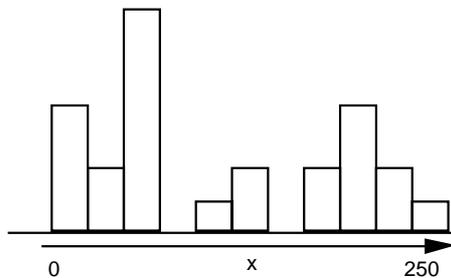


Figure 2: Histogram

The underlying functional

Histograms are "the" classical way to present a distribution. Its historical advantage is the ease of calculation - it can be reduced to putting registration notes into bins. The functional corresponding to a histogram is a discretization of the density: Given a decomposition of the real line into disjoint intervals $\{A_j, 0 \leq j \leq k\}$, we can define a histogram as the distribution, discretized to these intervals. The discretization gives a probability $p_j = F(A_j)$ for bin A_j . The vector of observed bin counts $n_j := \#\{i: X_i \in A_j\}$ has a multinomial distribution. Using Pearson's approximation (Pearson 1900) $P(n_1, \dots, n_k) \approx (2\pi n)^{-1/2} (\prod p_j)^{-1/2} \exp[-1/2 \sum (n_j - np_j)^2 / np_j + \dots]$ we see that the χ^2 test statistics is controlling the bin hit frequency for sufficiently large expected bin counts np_j .

Knowing the associated functional and its stochastic behaviour, we can tell what is to be expected from a histogram. A histogram can only show features which are preserved by the underlying functional, the discretization of the distribution. We lose all smoothness properties and local details of the distribution. The metric associated most naturally to the histogram is a χ^2 -metric. So if we want have an impression about the distributions which are compatible with our data, we should consider χ^2 confidence bands. To obtain information about the power of the histogram, we can look at the χ^2 -test as a corresponding goodness of fit test.

Practical situations are slightly more complex. One complication may arise from the sampling scheme. A common case to sample for a certain time, instead of taking a fixed sized sample. This makes the total sample size a random variable. Under independence assumptions the vector of bin counts has a multivariate poisson distribution instead of a multinomial. A similar limit applies, but we gain one degree of freedom in the limit. Another complication may arise if we define our bins in a data dependent way. If the number of bins is small compared to the number of observations, the approximation still holds with good quality even if we use the

data first to estimate location and scale, and use bins based on these estimators. We still have to correct for the degrees of freedom in the goodness of fit test.

If we have fixed reference distributions, we can head for optimal choice of bins. Common strategies are to take bins of equal probability with respect to the reference distribution, or to take bins of constant width with cut points $a_j = a_0 + j \cdot h$, $0 \leq j \leq k$, and bins $A_0 = \{x \leq a_0\}$, $A_j = \{a_{j-1} < x \leq a_j\}$ for some bin width h . This still leaves us with the problem to decide about the number of bins. The classical recommendation for equal probability bins is to let the number of bins increase as $n^{2/5}$ (Mann and Wald 1942). A more detailed study of optimal choices of the bins with respect to the χ^2 statistics for various types of reference distributions is presented in Kallenberg et al. (1985).

Walking through our short check list, it is an easy exercise to see how histograms can be modified to compensate for known censoring or discretization. If the data recording is discretized (for example by cutting down to a small number of reported digits) and the histogram is not adapted, this discretization may interfere with the histogram discretization and lead to patterns – an effect known as "aliasing". Tuning a histogram to look for multimodality is a problem. Tuning it to identify asymmetry or tail behaviour again is comparably simple.

Perception of the histogram is a different matter. Prominent features, such as local maxima or general skewness, are perceived first: the eye does not perform the rescaling which would be necessary from a statistical point of view. Choice of bin width and bin locations are most critical to histograms (Figure 3). The rule is: if you use a histogram, don't use one. Use several.

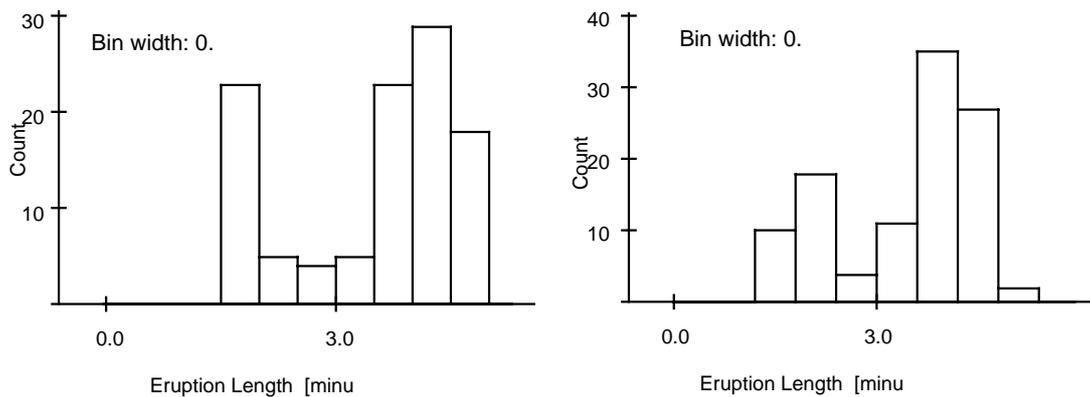


Figure 3: Histogram of eruption lengths of Old Faithful geyser. After Silverman (1986).

Using histograms to analyze for discretization effects or multimodality is inherently unsafe. A fairly safe use of histograms is for restricted purposes. Using a small number of cells to check for symmetry or tail behaviour is rather stable. But these purposes can be followed more effectively using Tukey's Box&Whisker plot, as discussed below.

Histogram density estimators

For the empirical version with bin counts $n_j := \#\{i: a_{j-1} < X_i \leq a_j\}$, the empirical histogram takes value $\hat{p}_j = \hat{\Pi}_j/n$ on A_j . We can view a histogram as an estimator for the density, the histogram density estimator being defined by $\hat{f}(x) = \hat{p}_j$ for $x \in A_j$. Judging the quality of this estimator by the integrated mean square error, $IMSE = \int E\{\hat{f}(x) - f(x)\}^2 dx$, at least for the case of constant bin width h asymptotic results are available [Scott 1979]. For the histogram density estimator \hat{f} , $IMSE = 1/(nh) + 1/12 h^2 \int f''(x)^2 dx + O(1/n + h^3)$, that is the error goes to zero in order $o(n^{-2/3})$. The integrated mean square error is minimized for $h^* = n^{-1/3} \{6/\int f''(x)^2 dx\}^{1/3}$. In particular, for the Gaussian distribution, $h^* \approx 3.49 \sigma n^{-1/3}$.

Choosing a suboptimal bin width $h = ch^*$ gives an error $IMSE \approx IMSE_{h^*} (c^3 + 2)/30$. For example: choosing a bin width as 50% of the optimal bin width increases the error by 42%.

Under regularity assumption, the optimal bin width needs approximately at least $(2n)^{1/3}$ bins. An upper bound for the bin width is $3.55 \sigma n^{-1/3}$. The regularity assumptions are: $\int f''(x)^2 dx > 0$, $\int f'''(x)^2 dx < \infty$.

But judging histograms by the integrated mean square error of the histogram density estimator is not fair. First, in constructing the histogram we deliberately restricted possible estimators to stepwise constant functions. Using the same information as contained in histograms, we can generate better density estimators by allowing piecewise linear estimators. For example we can join the bin centers of a histogram. The frequency polygon, the resulting density estimator, can achieve better error rates (Terrell and Scott 1985). For their optimal bin width $h^* = 2[15/49 n^{-1} \int \{f''(x)\}^2 dx]^{1/5}$ the error decreases as $n^{-4/5}$ under regularity condition: f'' absolutely continuous, $\int f''(x)^2 dx > 0$, $\int f'''(x)^2 dx < \infty$.

Under these regularity conditions, to achieve optimal bin width, we need at least $(147n/2)^{1/5}$ bins. An upper bound for the optimal bin width is $2.24 \sigma n^{-1/5}$. In particular, for the Gaussian distribution, $h^* \approx 2.15 \sigma n^{-1/5}$.

Thus the quality of the histogram, interpreted as a density estimator, can be greatly improved by allowing a slightly modified class of derived density estimators. But IMSE still is not an adequate criterion to judge histograms. Histograms are discretized versions of the distribution. The empirical histogram should be judged in comparison to this discretization, and competitors should be ranked on how they perform at (possible discretized) data.

Scatter plots

Recipe: Mark the data points.



Figure 4: Scatter plot. Same data as in figure 2.

The underlying functional

The functional corresponding to the scatter plot is the density; the empirical version is $\sum \delta_{X_i}$. Problems may arise from discretizations or drawing resolutions. These effects may lead to ties in the empirical version. A simple technical solution is to use a gray pen, instead of a simple black pen. If the pen has gray level $K(y)$ at distance y from its center, we get gray level $\sum K(x-X_i)$ at a point x : the plot corresponds to a kernel density estimator with kernel K , or to the smoothed density as a functional. Any pen (or kernel) gives rise to a whole scaled family K_h , with $K_h(y) = 1/h K(y/h)$, where h is the pen size, or bandwidth. For the functional, this means replacing the density f by the smoothed density $f * K_h$, the convolution of f and K_h . Of course you gain smoothness, and you lose details by using a larger bandwidth.

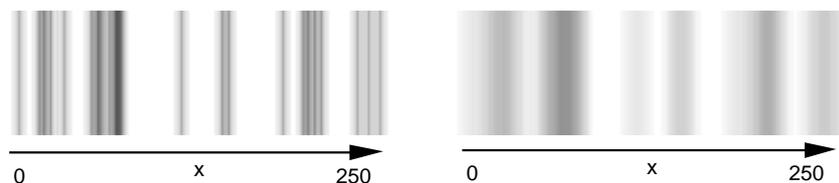


Figure 5: Scatter plot, using two different gray pens. Same data as in figure 4

Scatter plots as such can be barely considered diagnostic tools, but they may provide helpful orientation when combined with other plots: we continue to use them, but it is hard to pin down what we gain. To use them for diagnostic purposes, they must be enhanced. For example to check the symmetry behaviour, we can look at

symmetry centers. An appropriate plot suggested by J. Tukey (after Wilk and Gnanadeskian, 1968) is to show $(X_{(n-i+1)}+X_{(i)})$, plotted against the distance $X_{(n-i+1)}-X_{(i)}$. To check for discretization effects, we can look at the plot of the differences $X_{(i+1)}-X_{(i)}$ against $X_{(i)}$.

Smoothed scatterplots and kernel density estimators

In principle, the complete information of a sample is represented in a scatter plot. Perception however is easily trapped by sample size effects: small sample sizes will give the impression of pattern and inhomogeneities even for uniform samples; large sample sizes will hide non-uniformities for any distribution. This problem is even more complicated for kernel density estimators: we have the choice of a pen (or a kernel, if you like). Conventionally, this problem is split into two: choosing the pen shape (of kernel type) and the pen size (or bandwidth). We meet the same problem we have encountered with histograms: what we see depends critically on these choices. But we do not know how to judge these choices. There is a mathematical hideaway. If we accept that the density is our target functional, any distance measure between the (normalized) kernel density estimate and the true density can be used as a measure of fit, and of course L_2 distance is the easiest to deal with. Call $\hat{f}_h(x) = n^{-1}h^{-1}\sum K((x-X_i)/h)$ the kernel density estimator for kernel K and bandwidth h . Let $\underline{h}_0 = \underline{h}_0(f, X)$ be the smallest minimizer of the integrated square error $\Delta(h) = \int (\hat{f}_h - f)^2$ and h_0 the smallest minimizer of the mean integrated square error $M(h) = \int E(\hat{f}_h - f)^2$. Under regularity conditions, for any (empirical) bandwidth \underline{h} we have $\Delta(\underline{h}) - \Delta(h_0) = 1/2 (1 + o_p(1))(\underline{h} - h_0)^2 M''(h_0)$ (Hall and Marron 1987). While this does not help to estimate the error, it says that minimizing the integrated square error is essentially equivalent to optimizing the bandwidth for the data at hand. But \underline{h}_0 can be represented as $\underline{h}_0 = \underline{A}_1 + n^{1/5} \underline{A}_2 \int f^2 + o_p(n^{-3/10})$, where \underline{A}_1 and \underline{A}_2 are functions of the data, not depending on f (Hall and Johnstone 1992). \underline{A}_2 does not vanish asymptotically. So determining an optimal bandwidth is related to estimating $\int f^2$. An optimal rate of $n^{-1/2}$ for the estimation of $\int f^2$ makes the relative error of approximating \underline{h}_0 at best of order $n^{-1/10}$ (Hall and Johnstone 1992). These results tell us why optimal bandwidth selection is a hard problem even for very large sample sizes and continuous distributions, let alone for real data, that is for finite sample sizes and data truncated or rounded to some finite precision.

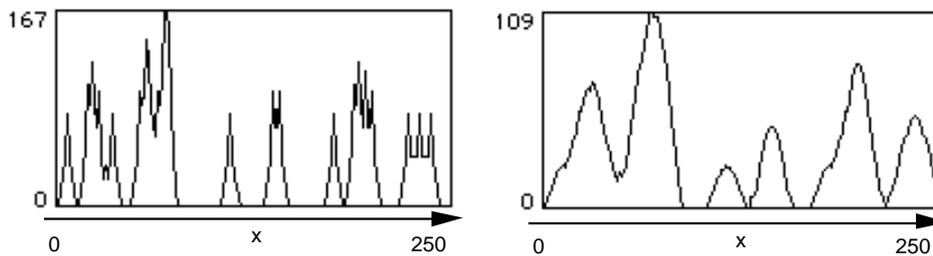


Figure 6: Pixel intensity for gray level plots of figure 5

It is possible to base goodness of fit tests on kernel density estimators (Mammen 1992, Ch. 3). But the stochastic behaviour of kernel density estimators is difficult. There is no clear notion of distance or variation associated to kernel density estimators. There are candidates, among them distances based on (penalized) square errors, possibly. These are a treatable mathematical concept, but L_2 confidence bands are not too helpful from a data analytical point of view. The information gained from scatter plots, including kernel density estimators, is doubtful. Checking the list of critical features given above, it is hard to spot a feature that is reliably detected and reported by a scatterplot. Silverman (1981) made an attempt to exploit kernel density estimators as a diagnostic tool to analyze for multimodality. A simpler approach, the densitogram (related to the excess mass test), is given below.

Distribution function and related plots

Recipe: Sort the data points. For any point, mark the proportion covered (the frequency of data points not exceeding this point).

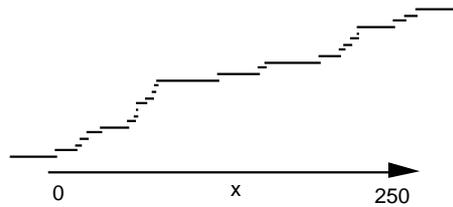


Figure 7: Distribution function

The underlying functional

The distribution function gives the probability of half-lines $F: x \in \rightarrow F(x) = P\{X \leq x\}$. It can be estimated by its empirical version, F_n where $F_n(x_{(i:n)}) = i/n$. The stochastic behaviour is described by the Glivenko-Cantelli Lemma: we have $\sup_x |F_n(x) - F(x)| \xrightarrow{P} 0$. The error has a Brownian bridge asymptotics: $\sqrt{n}(F_n - F) \rightarrow Z$. Viewed as an estimator for F , F_n has a certain general optimality: For any loss function of supremum type, the empirical distribution function is asymptotically a minimax estimator (Dvoretzky-Kiefer-Wolfowitz-Theorem). For continuous distributions F , the distance $D_n = \sup_x |F_n(x) - F(x)|$ has a distribution which does not depend on F . This allows for simultaneous confidence bands: if c denotes the α -quantile of the Kolmogoroff-Smirnov statistics, we have $F(x) \in [F_n(x) - c, F_n(x) + c]$ for all x with probability of at least $1 - \alpha$. So the distribution function is easy to reconstruct, and its statistics is well understood. Interpreting it needs some education.

Plots related to the distribution function

Comparing two distribution functions visually is quite difficult. We have to compare two graphs, both piecewise constant and monotonous. Most interesting features are hidden in details. We can help perception by using a transformation which gives a near-to linear graph for corresponding distributions. If we have a given reference distribution, our choices are to align quantiles by transforming the probability scale (the quantile-quantile-plot), or to align probabilities by transforming the data scale (the percentage-percentage plot).

Quantile-quantile-plot (Q-Q-plot)

Recipe: Choose a reference distribution. Sort the data points. For any data point, find the proportion of observations not exceeding this data point. Plot the data point against the corresponding quantile of the reference distribution.

Q-Q-plot details

To transform the probability scale, we transform a probability to the corresponding quantile. The Q-Q-plot compares two distributions by plotting quantile against quantile. If F and G are the distributions to be compared, $X \sim F$, $Y \sim G$, the Q-Q-plot shows the curve $\alpha \rightarrow (x_\alpha, y_\alpha)$. In terms of the probability distributions, this is the graph of $x \rightarrow G^{-1} F(x)$. Again, orientation has been chosen to give an easy empirical version $x \rightarrow G^{-1} F_n(x)$.

If F and G coincide, the Q-Q-plot is a diagonal line. If one is a linear transformation of the other, the Q-Q-plot is linear. The Q-Q-plot shows a high resolution in regions of low densities and vice versa. As a consequence, it emphasizes the tail behaviour for long-tailed distributions (Wilk and Gnanadesikan, 1968), and emphasis on the tails combines unluckily with high variation.

If G is the true distribution, $G=F$, the Q-Q-plot of F_n against F is given by $(X_{(i:n)}, x_{i/n})$ where $x_{i/n}$ is the i/n quantile. In particular, for $G=U[0,1]$ we have $x_{i/n}=i/n$, i.e. the Q-Q-plot coincides with the empirical distribution function.

If U_i iid $\sim U[0,1]$, $U_{(i:n)}$ is distributed as $\beta(i, n-i+1)$. Hence generally $E(F X_{(i:n)}) = E(U_{(i:n)}) = i/n+1$. We can take this into account to get an "unbiased" empirical plot by using plot positions $(X_{(i:n)}, G^{-1}(i/n+1))$ for an empirical Q-Q-plot. This is the convention used by Weibull (1939). But getting the mean behaviour right is only one part of the difficulty. Since you will not apply a diagnostic plot to a mean situation, but to a sample,

you are prone to be affected by the notorious skewness of empirical quantile distributions. This is the origin for many fine points to be considered in the actual mapping (Kimball (1960), Harter (1984)).

Direct relatives of the Q-Q-plot are goodness-of-fit tests based on the regression of order statistics on expected order statistics, like for example the Shapiro-Wilk test (Shapiro and Wilk, 1965).

Percentage-percentage plot (P-P-plot)

Recipe: Choose a reference distribution. Sort the data points. For any data point, find the proportion of observations not exceeding this data point. Plot the proportion against the corresponding proportion of the reference distribution.

P-P-plot details

To transform the data scale for linearity, we have to transform X to the corresponding probability under the reference distribution. If F and G are the distributions to be compared, $X \sim F$, $Y \sim G$, the P-P-Plot shows the curve $X \rightarrow (G(X), F(X))$. In terms of the probability distributions, this is the graph of $\alpha \rightarrow F G^{-1}(\alpha)$. We apply this to G as a reference distribution. Orientation has been chosen here to avoid the discontinuities in F_n , that is to give an empirical version $\alpha \rightarrow F_n G^{-1}(\alpha)$.

If F and G are identical, the P-P-plot will be straight line. P-P-plots are not preserved under linear transformations: they are not equivariant. So usually P-P-plots will be applied only to distributions standardized for location and scale. For the empirical version, this preferably is done using robust estimators of location and scale. As for the Q-Q-plot, the skewness of the empirical quantile function should be considered in the actual mapping. But in contrast to the Q-Q-plot, for the P-P-plot, high variability is not combined with sensitivity in the tails. So choice of the proper plotting position is a fine point for the P-P-plot, whereas it is critical choice for the Q-Q-plot.

Goodness-of-fit tests can be constructed based on the linearity of the P-P-plots (see Gan and Koehler, 1990).

Other plots related to the distribution function plot

The plots based on the distribution function suffer from the tail-orientation of the distribution function. It measures half infinite intervals, and local behaviour can be judged only by looking at differences. This is easy to compensate using a third dimension: you can localize the probability mass to intervals and define a probability mass plot $(a,b] \rightarrow F(b)-F(a)$, with the obvious empirical version. But readability and practical use are doubtful.

Box&Whisker-Plot

Recipe: Find the median and quartiles, and mark them. Connect the range of points which are not too far from the median (judged by the interquartile distance). Highlight all points which are out or far out.

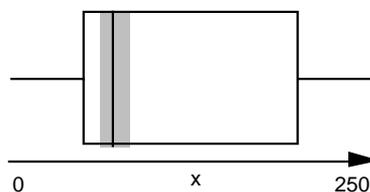


Figure 8: Box&Whisker plot

Box&Whisker plot details

In more detail, the construction is: Find the median of the data points, and mark it. Find the median of the subset above the general median, mark it, and call it the upper hinge. Find the median of the subset below the general median, mark it, and call it the lower hinge. Let Δ be the distance between the hinges. Draw a whisker from the box to the last data point not exceeding upper hinge + 1.5Δ . Mark all data points in the out area between hinge + 1.5Δ and hinge + 2.5Δ . Highlight all far out points exceeding hinge + 2.5Δ . Do the corresponding for the lower hinge.

John Tukey's Box&Whisker plot is one of the gems of data analysis. Like the histogram, the Box&Whisker plot represents a discretization of the density. But where the histogram discretizes on the observation scale, the Box&Whisker plot discretizes on the probability scale. The discretization varies, from a rough 25%-discretization in the center part, to a 1/n discretization for a sample size n in the tails.

The Box&Whisker plots achieve to present general information about the core of the data, with information hiding in this area. On the other hand, they highlight the exceptional. The exceptional data might be just tail effects, or it might be genuine outliers - they are worth a second look anyway.

The Box&Whisker plot is best understood by following its construction. Roughly, the Box&Whisker plot marks median and quartiles, and exceptional points. We will try to look at the ideas of Box&Whisker plots more carefully here. For the Box&Whisker plot, first we try to get an estimator for the location. The data median is used as the obvious (robust) candidate. The center line of the Box&Whisker plot marks the median. Now we estimate the scale. Since we have already estimated the location, we can use this information. Given an estimator for the location, estimating **the** scale would be useful in exceptional cases: it would be meaningful only for symmetric distributions. Given the location estimator, we construct two scale estimators, a lower and an upper scale estimator. In the absence of ties, we could use the differences between median and lower/upper quartile as estimators. Since we must be prepared for discretization effects, we must be more careful. We use the median of the lower or upper half instead – Tukey's hinges. Finally, using these scale estimators, we estimate "central" areas, and mark all points outside.

Tukey's Box&Whisker plot takes into account many possibilities and pitfalls of real data sets. It is very easy to miss these fine points, as can be seen from popular software packages.

The Box&Whisker plot is particularly powerful in analyzing the overall structure of a distribution, like location, scale and outliers. But it still leaves the needs to diagnose other features. Discretizations are in no way reflected in the Box&Whisker plot. The tail behaviour is made a caricature: if there are tails, outliers are identified. But if the tails are too steep, heavier tails are invented: even a uniform distribution is shown with tails. Multimodality is ruled out: the Box&Whisker plot knows about central location, but has no space for modes. It must be accompanied by other plots.

Silhouette and Densitogram

Recipe: Choose a family of sets serving as a model (e.g. sets composed of one or two intervals, if you are looking for bimodality). Choose a level λ . Mark the maximal set with average hit density exceeding level λ . Do this for a choice of levels λ .

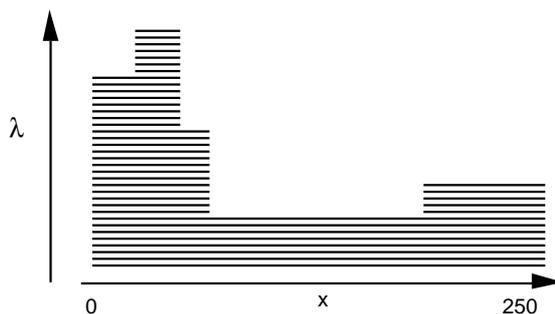


Figure 9: Silhouette. Locations of excess mass $E(l) := \int (f-\lambda)^+ dx$ for varying levels l .

The underlying functional

If you are looking for specific features in your data, it is possible to design diagnostic plots for these features. Silhouette and its accompanying plot, the densitogram, are plots tuned to inspect multi-modality (Müller and Sawitzki 1987). Both are based on the idea that a mode of a distribution is a location where the probability mass is concentrated. A corresponding functional is the excess mass, $E(\lambda) := \int (f-\lambda)^+ dx$, giving the probability mass exceeding λ . Restricting the allowed sets in an appropriate way to a family C , we define $E_C(\lambda) :=$

$\sup_{C \in \mathcal{C}} \int_C (f - \lambda) dx = \sup_{C \in \mathcal{C}} (F - \lambda \text{Leb})(C)$. The silhouette marks the maximizing sets, for any level λ . The densitogram shows the excess mass, as a function of λ . The cue lies in the freedom to choose C . For unimodal distributions, $C = C_1$ should be the family of intervals; for bimodal distributions, $C = C_2$ is made of the disjoint unions of two intervals. Given a hypothesis on the modality, silhouette and densitogram can be estimated by their empirical version.

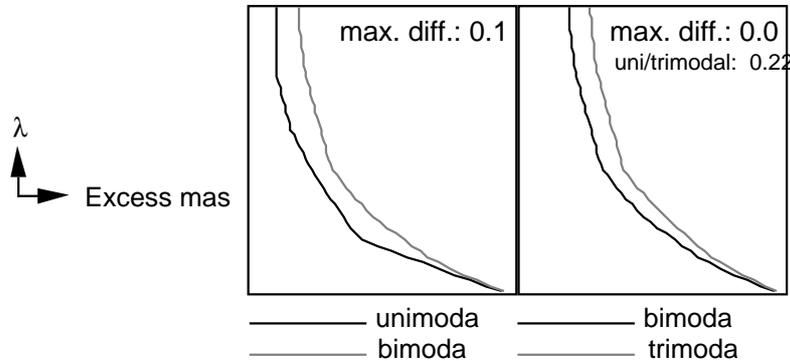


Figure 10: Densitogram, the excess mass concentration curve. Same data as figure 2. Excess mass estimated under assumption of uni- bi- and tri-modality. By assuming bi-modality, an additional excess mass of 16.4 % of the data is covered.

Silhouette and densitogram details

As an estimator for the location of the mode, the silhouette shares a poor with of order $n^{-1/3}$ with density estimation based methods. The number of modes in the silhouette however is more reliable even for small sample size. For the densitogram, the associated test is the excess mass test for multimodality (Müller and Sawitzki 1991): $\sup_{\lambda} (E_{C_2}(\lambda) - E_{C_1}(\lambda))$, the maximal difference between excess mass $E_{C_2}(\lambda)$, estimated on the assumption of bimodality, and excess mass $E_{C_1}(\lambda)$, estimated on the assumption of unimodality, can be used as a test statistic for multimodality. For a bimodal distribution, the maximal excess mass difference $\sup_{\lambda} (E_{C_2}(\lambda) - E_{C_1}(\lambda))$ is half the total variation distance between F and the closest unimodal distribution. This points to the total variation as a distance measure related to excess mass.

On the unimodal distributions, the error rate of these excess mass estimates is of order $n^{-1/2}$. In more practical terms: the difference between both excess mass curves starts providing a reliable indicator for multimodality for a sample size n in the range 20 to 50.

Shorth-Plot

Recipe: Choose a coverage a . For any point, get the length of the shortest interval containing this point and covering at least an a -fraction of the data (at least $a \cdot n$ data points). Do this for a selection of coverages a .

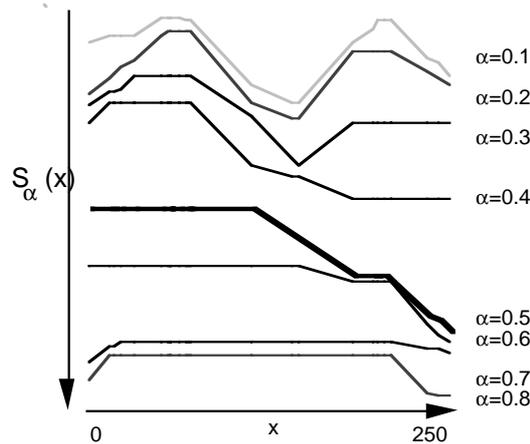


Figure 11: Shorth Plot. The shorth length axis points downwards.

The underlying functional

The shorth is the smallest interval containing at least 50% the distribution: $S = \arg \min \{ |I|: I=[a,b], P(X \in I) \geq 0.5 \}$. Here $|I|$ is the length of the interval I . More generally the α -Shorth is the smallest interval containing at least an α fraction of the distribution: $S_\alpha = \arg \min \{ |I|: I=[a,b], P(X \in I) \geq \alpha \}$.

For data analysis, we can localize the shorth. We define the α -shorth at x as the smallest interval at x containing at least a proportion α of the distribution $S_\alpha(x) = \arg \min \{ |I|: I=[a,b], x \in I, P(X \in I) \geq \alpha \}$. In particular, the shorth at x is defined as $S(x) := S_{0.5}(x)$.

More about shorth-plots

Andrews et al. (1972) use the (not localized) shorth to construct a robust estimator of location. The shorth procedure takes the center of the empirical shorth as location estimator. Unfortunately this estimator of location has an asymptotic rate of only $n^{-1/3}$, with non-trivial limiting distribution. However Grübel (1988), shows that the length of the empirical shorth is a reasonable estimator of scale, converging with a rate of $n^{-1/2}$ to a Gaussian limit. This result can be carried over to the localized shorth; Grübel's proof goes through with the obvious changes (Sawitzki 1992).

Mass concentration now can be represented by the graph of $x \rightarrow |S_\alpha(x)|$. A small length of the shorth signals a large mass concentration. To facilitate optical interpretation, we plot the negative of the lengths.

Summary

How far have we got? The general purpose plots (histogram, scatter plot, distribution function) can be applied, but provide doubtful information per se. They can be sufficiently restricted to provide reliable information on questions as rough symmetry or tail behaviour. But the rough information seems to be read off more readily from constructions as the Box&Whisker plot. The general purpose plots may have an advantage if we move to the one-sample problem or the two sample problem, where no immediate generalization of the Box&Whisker is available.

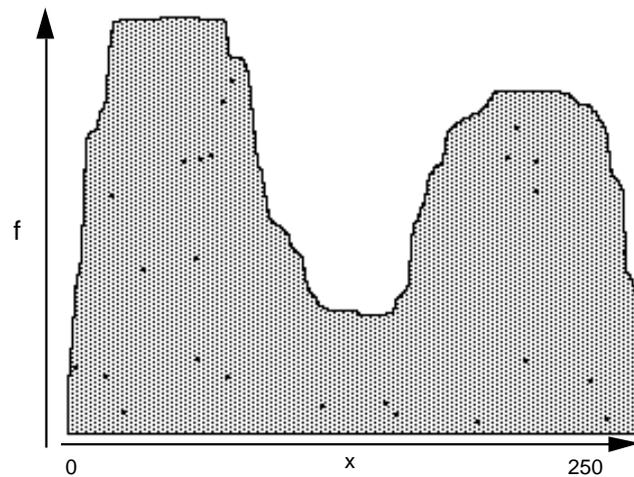
Multimodality stays a critical feature. The classical general purpose plots tend to be misleading: random fluctuation may appear as modes, and no controlled measure of significance is available. The general purpose plots are not likely to oversee modes, but are prone to show more than should be shown. The Box&Whisker plot does not address the problem of modes at all.

We can construct special plots for the detection of modes, such as silhouette and densitogram. We lose information on density and tails in these plot.

The shorth plot tries to make a compromise, allowing for information about modality as well as on local density, but avoiding the fluctuation affecting (smoothed) scatter plots and other classical plots. It may be a candidate for a general purpose plot. But practical evaluation and analysis is still necessary.

The true distribution is usually hidden from our eyes. Since we were using simulated data here, we are able to look at the true distribution. In the examples shown here, we used a bimodal distribution with two strong

modes. Sample size for the illustrations was 25 data points.



Literature:

- Bertin, J. (1967). *Semiologie Graphique*. Gauthier-Villars, Paris.
- Bertin, J. (1977). *La Graphique et le Traitement Graphique de l'Information*. Flammarion, Paris.
- Bertin, J. (1981). *Graphics and Graphic Information Processing*. De Gruyter, Berlin.
- Chambers, J.M.; Cleveland, W.S.; Kleiner, B.; Tukey, P.A. (1983). *Graphical Methods for Data Analysis*. Wadsworth Statistics/Probability Series, Wadsworth, Belmont.
- Gan, F.F.; Koehler, K.J. (1990). Goodness-of-Fit Tests Based on P-P Probability Plots. *Technometrics* 32, 289 - 303.
- Grübel, R. (1988). The Length of the Shorth. *Annals of Statistics* 16, 2:619-628.
- Hall, P.; Marron, S. (1987). Extent to which Least Squares Cross Validation Minimises Integrated Squared Errors in Nonparametric Density Estimation. *J. Probab. Theory and Related Fields* 74, 567-581
- Hall, P.; Johnstone, I. (1992). Empirical Functionals and Efficient Smoothing Parameter Selection. *Journal of the Royal Statistical Society, Series B*, 54, 475-530.
- Harter, H.L. (1984). Another Look at Plotting Positions. *Communications in Statistics—Theory and Methods* 13, 1613-1633.
- Kallenberg, W.C.M; Oosterhoff, J.; Schriever, B.F. (1985). The Number of Classes in Chi-Squared Goodness-of-Fit Tests. *Journal of the American Statistical Association* 80, 959 - 968.
- Kimball, B.F. (1960). On the Choice of Plotting Positions on Probability Paper. *Journal of the American Statistical Association* 55, 546-550.
- Mammen, E. (1992). *When does Bootstrap Work ? Lecture Notes in Statistics*. Springer, Heidelberg.
- Mann, H.B.; Wald, A. (1942). On the Choice of the Number of Intervals in the Application of the Chi-Squared Test. *Annals of Mathematical Statistics* 13, 306-317.
- McGill, R.; Tukey, J.W.; Larsen, W.A. (1978) Variations of Box Plots. *American Statistician* 32, 12-16.
- Müller, D.W.; Sawitzki, G. (1987). Using Excess Mass Estimates to Investigate the Modality of a Distribution. Universität Heidelberg, Sonderforschungsbereich 123 (Stochastische Mathematische Modelle). Reprinted in: *Proceedings of the ICOSCO-I Conference, (First International Conference on Statistical Computing, Çesme, Izmir 1987) Vol II*. American Science Press, Syracuse 1990.
- Müller, D.W.; Sawitzki, G. (1991). Excess Mass Estimates and Tests for Multimodality. *Journal of the American Statistical Association* 86, 738-746.
- Pearson, K. (1900) On a Criterion that a Given System of Deviations from the Probable in the Case of a Correlated System of Variables is such that it can be Reasonably Supposed to Have Arisen from Random Samples. *Philosophical Magazine (5th Series)* 50, 157-175.
- Sawitzki, G. (1990). Tools and Concepts in Data Analysis. In: F. Faulbaum, R. Haux, K.-H. Jöckel (eds.) *SoftStat '89 Fortschritte der Statistik-Software 2*. Gustav Fischer, Stuttgart. 237-248.

- Sawitzki, G. (1992). The Shorth Plot. Technical Note. Heidelberg 1992.
- Shapiro, S.S.; Wilk, M.B. (1965). An Analysis of Variance Test for Normality (Complete Samples). *Biometrika* 52, 591-611.
- Scott, D.W. (1979). On Optimal and Data-based Histograms. *Biometrika* 66, 605 - 610.
- Silverman, B.W. (1981). Using Kernel Density Estimates to Investigate Multimodality. *Journal of the Royal Statistical Society, Ser. B.*, 43, 97-99.
- Silverman, B.W. (1986) *Density Estimation for Statistics and Data Analysis*. London: Chapman and Hall
- Terell, G.R.; Scott, D.W. (1985). Oversmoothed Nonparametric Density Estimators. *Journal of the American Statistical Association* 80, 209 - 214.
- Tufte, E.R. (1983). *The Visual Display of Quantitative Information*. Grapics Press, Cheshire, Connecticut.
- Tufte, E.R. (1990). *Envisioning Information*. Grapics Press, Cheshire, Connecticut.
- Tukey, J.W. (1962). The Future of Data Analysis. *Annals of Mathematical Statistics* 33, 1-67.
- Wainer, H. (1984). How to Display Data Badly. *The American Statistician* 38, 137-147
- Weibull, B.A. (1939). The Phenomenon of Rupture in Solids. *Ingeniors Vetenskaps Akademien Handlingar* 153, 7.
- Wilk, M.B.; Gnanadeskian, R. (1968). Probability Plotting Methods for the Analysis of Data. *Biometrika* 55, 1 - 17.