

Advanced Distribution Theory for SiZer

J. Hannig

Department of Statistics
Colorado State University
Fort Collins, CO 80523-1877

J. S. Marron

Department of Statistics
University of North Carolina
Chapel Hill, NC 27599-3260

January 2, 2004

Abstract

SiZer is a powerful method for exploratory data analysis. In this paper the distributions underlying the statistical inference are investigated, and large improvements are made in the approximation. This results in improved size, and also in an improved global inference version of SiZer. The main points are illustrated with real data and simulated examples.

1 Introduction

SiZer has proven to be a valuable technique for exploratory data analysis by smoothing methods. These methods include histograms and smoother approaches to understanding the structure of one-dimensional distributions (called the “density estimation setting” here), and scatterplot smoothers (called the “regression setting” here). See for example Scott (1992), Wand and Jones (1995) and Fan and Gijbels (1996), for an introduction to this area. As noted in these monographs, many smoothing, i.e. estimation, schemes have been proposed. See Marron (1996) for an overview of the many criteria that have been used to compare different smoothing methods. Kernel based methods (definitions are given in Section 2) are considered here for their simplicity, ease of interpretation, and because they have been very widely studied.

Practical use of kernel methods, in both density estimation and regression, is profoundly affected by the choice of the window width (the tuning parameter which controls the amount of local averaging being used). When this is too small, the resulting estimated curve strongly feels sampling variation, and is wiggly, reflecting spurious artifacts of the sampling process. For too large a window width, the curve estimate smooths away important underlying features. There is a large literature on data based selection of the window width, where one tries to estimate it from the data, see Jones, Marron and Sheather (1996a,b). However, the problem is very challenging, there are limits on how well this selection can be done in practice, and there has never been a consensus on what

is “the best” method of doing this, which has appeared to slow actual use of these methods, for example through their implementation in software packages.

Scale space ideas (see Chaudhuri and Marron (2000) for broad discussion of these issues) have provided a practical means of avoiding the problem of bandwidth selection. Scale space is a theoretical model for vision, that was constructed in the computer vision community. The model is simply a family of Gaussian window smooths, indexed by the window width. It is a model for vision in the sense that large values of the window width correspond to standing back and viewing a scene macroscopically, while small values correspond to a zoomed in view. See Lindeberg (1994) and ter Haar Romeny (2001) for access to the scale space literature. A fundamental concept of scale space, that is the heart of SiZer, is that instead of trying to choose a single “best scale” (i.e. best window width), one should use all of them, i.e. study the full family of smooths. This is clear in a vision modeling context, because different levels of resolution (i.e. smooths with different window widths) of an image contain different types of useful information.

SiZer is a combination of the scale space idea of simultaneously considering a family of smooths, with the statistical inference that is needed for exploratory data analysis, in the presence of noise. In particular, SiZer addresses the question of “which features observed in a smooth are really there?”, meaning representing important underlying structure, not artifacts of the sampling noise.

For reasonable statistical inference using SiZer, care needs to be taken about the multiple comparison issue. In particular, the visual display of SiZer, can be viewed as a summary of a large number (hundreds) of hypothesis test results. Current implementations of SiZer address this issue using the fairly crude “independent blocks” idea, developed in Section 3 of Chaudhuri and Marron (1999). In this paper, a much deeper distributional investigation is done, with the goal of improving the statistical performance of SiZer.

The SiZer method, as well as potential advantages from an improved distribution theory, are illustrated in Figure 1. The underlying regression function, shown as the thick black curve in Figure 1a, is the Blocks example from Donoho and Johnstone (1994), which appears to be rather challenging to estimate by smoothing methods, because of the 11 sharp jumps. To make the problem even more challenging, a high level of Gaussian noise, $\sigma = 0.1$, (much higher than is typical in the wavelet literature) that was first used by Marron et al (1998) is used in the generation of the $n = 1024$ data points shown as green dots in Figure 1a.

The thin blue curves in Figure 1a show the scale space for this data set, i.e. the family of smooths for a wide range of different window widths. Some of these are seriously oversmoothed, showing strong rounding of the corners. Some are undersmoothed showing spurious wiggles. None of these are very good at attaining the goal of recovering the thick black curve. Wavelets, see e.g. Donoho and Johnstone (1994), are a compelling approach to the problem of recovering curves such as this with non-smooth features. However, for this data, even wavelets give poor signal recovery, because the noise level is so high.

SiZer has a somewhat different goal. Instead of trying to recover the un-

derlying black curve as well as possible, it aims instead at understanding which of its features can be distinguished from the background noise, i.e. determining which aspects observable in the blue curves are important underlying structure, and which are spurious noise-driven artifacts.

Figure 1a




Figure 1b




FIGURE 1: *Conventional SiZer analysis of the Donoho - Johnstone Blocks regression, with high noise. Shows good performance, plus a spurious feature. True regression, data and scale space shown in Figure 1a. SiZer analysis in Figure 1b.*

SiZer focuses on finding regions of statistically significant slope in the blue curves. Slope works well in the example of Figure 1, because the interesting features there are the 11 jumps (elsewhere the regression is flat). With the high noise level used in Figure 1a (making signal recovery challenging, even by the best wavelet methods), determining which jumps are statistically significant turns out to be attainable by SiZer. In other cases of data analysis using

smoothing methods, bumps are of interest. Bumps are also determined by slope, because the curve slopes up on one side, and down on the other. In general, SiZer flags features of these various types using a color map, such as the one shown in Figure 1b.

The horizontal location in the SiZer map are the same as in the scale space plot above. The vertical locations correspond to the window widths of the family of blue curves, shown on the log scale. Each pixel shows a color that essentially gives the result of a hypothesis test for the slope of the blue curve, at the point indexed by the horizontal location, and at the scale (window width) corresponding to that row. When the slope is significantly positive (negative) the pixel is colored blue (red, respectively). When the slope is not significant (as happens in regions where sampling noise is dominant), the color purple is used. There is a fourth SiZer color, that does not appear in Figure 1b, which is gray, used to show pixel locations where the data are too sparse for reasonable statistical inference.

Note that each jump in Figure 1a corresponds to a red or a blue (depending on the direction of the jump) region in the SiZer map in Figure 1b. Thus SiZer has correctly found all 11 of the jumps in the thick black curve, so for the specific goal of finding important features it substantially outperforms wavelet methods.

A very careful look at the SiZer map shows a small, unexpected feature: a tiny blue region at the finest scales (the bottom of the map) near 0.58. This is suggesting the slope is statistically significant, when in fact the underlying target curve is flat. Such features have been observed in a number of other cases as well. This has not presented a serious obstacle to data analysis by SiZer, because analysts have learned to not put too much credence into such very small features when they are flagged by SiZer. But it is still very desirable to eliminate these, to give a more precise analysis. This goal is attained in the present paper, by developing an improved distribution theory.

First, we take a deeper look at the extent of the problem of small spurious features appearing in the SiZer map, by studying some simulations. Figures 2 and 3 show some SiZer maps for simulated data from the null distribution in the case of equally spaced design regression. Since the regression function is 0, the data are simply i.i.d. standard Gaussian random variables. In this situation, the SiZer map should ideally be completely purple, except for perhaps $\alpha 100\%$ of the cases in the size α case (here α is always taken to be 0.05).

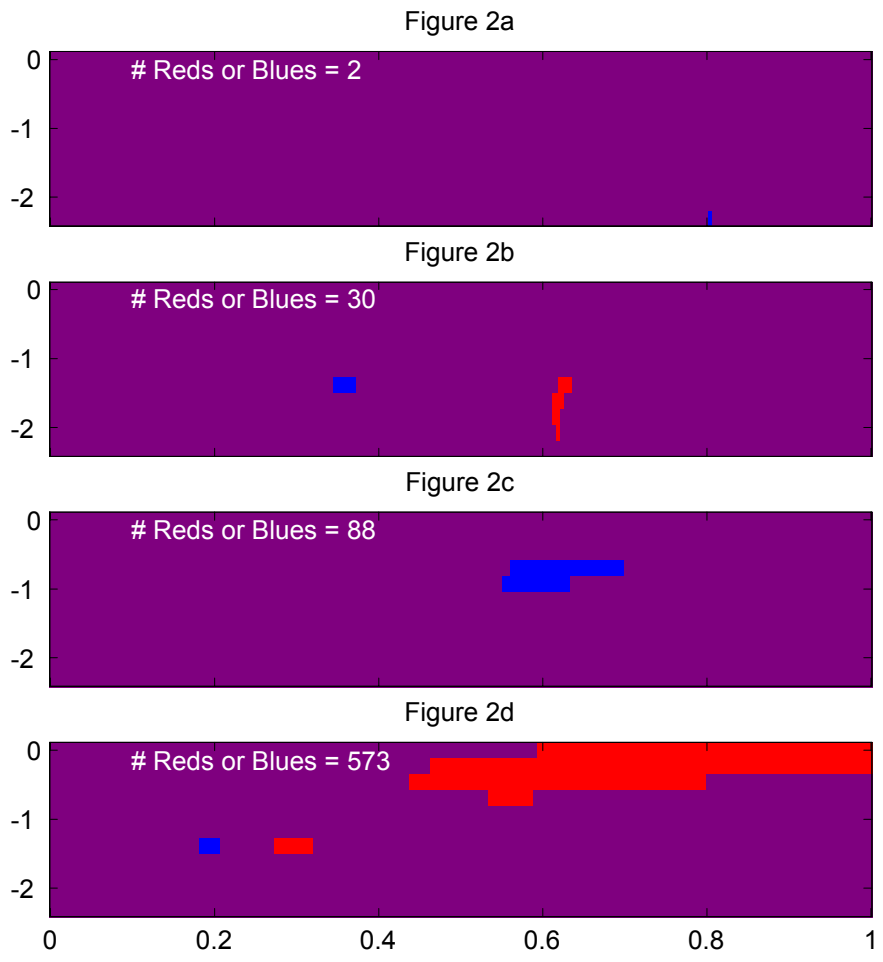


FIGURE 2: Conventional SiZer maps, based on simulated null distributions, for $n = 1600$ equally spaced regression data points. Figures 2a, b, c and d are for 0.5, 0.75, 0.85 and 0.95, respectively, quantiles of distribution.

The SiZer maps shown in Figure 2 illustrate the population of SiZer maps for this underlying distribution. They were drawn from a simulated sample of 1000 such SiZer maps. The population was ordered in terms of number of pixels that flag significant structure by being red or blue. Because these were drawn from the null distribution, it is desirable for the number of such pixels to be small. The first 405 of the 1000 ordered SiZer maps were completely purple (and are thus not shown to save space). Figure 2a shows the 500th of these (essentially the median of the population), where two pixels, at the finest scale, were flagged as significant. Figure 2b shows the 750th (the 3rd quartile), with substantially more significant pixels at medium fine scales. Figure 2c shows the 850th SiZer map, with quite a large blue region at medium coarse scales. Figure 2d is the 950th member of the ordered population, showing an even larger red

region at the coarsest scales, plus the suggestion of a small mode at medium scales. There appears to be a relationship between the number of spurious pixels, and the scales at which they appear, which is not surprising because at coarse scales adjacent pixels are strongly correlated.

This suggests a serious need for improvement in the size characteristics of the conventional SiZer. The ideal here is that Figures 2a-c should be completely purple, and that Figure 2d might or might not have some color. The goal of this paper is to improve this performance, by using a better approximation of the underlying distribution theory.

The results of our first attempt are shown in Figure 3. The format is the same as Figure 2, based on the same 1000 underlying data sets, but this time an improved version of the SiZer map is used. Again the maps were ordered, and the 500th, 750th, 850th and 950th of the 1000 maps are shown as Figures 3a, 3b, 3c and 3d respectively.




FIGURE 3: *SiZer* maps for simulated null distributions, for $n = 1600$ equally spaced regression data points, based on the new row-wise procedure. Figures 3a, b, c and d are for the 0.5, 0.75, 0.85 and 0.95, respectively, quantiles of the distribution.

The SiZer maps in Figure 3 flag far less spurious structure than was found, for the corresponding population quantiles, in Figure 2. In particular, in Figures 3a and 3b (representing the first 3 quartiles), there were no spurious results. Even for the 850th ordered SiZer map, in Figure 3c, the spurious structure is quite small. Hence the improved SiZer map studied in Figure 3 clearly has better size properties than the original SiZer shown in Figure 2. However, these results are still not completely satisfactory.

This size problem is driven by a number of factors, that are studied in Section 2, the most important of which is that the simultaneous inference is only row-wise in nature. This means that the SiZer inference in Figure 3 is only adjusted row by row. Hence it is not surprising that some spurious structure manages

to be flagged here, since each of the maps in Figure 3 includes 11 such rows (so just by chance, the test flags significant structure more often than 5% of the time).

This problem is addressed in Section 2 by using Bonferroni type adjustments. The corresponding globally adjusted version of Figure 3 has been plotted, but is not shown here (to save space) because each of the panels is completely purple, indicating that the size problem has been solved.

The distribution theory that drives the improvements in the statistical performance of SiZer shown in Figures 2 and 3 is developed in Section 2 with the main recommendations summarized in Section 2.5. Detailed analysis of the impact of these improvements is done in Section 3.

An important issue is that the improved size properties, further investigated in Section 3.1, come at a sometimes substantial cost in terms of power. Power issues are studied for simulated data in Section 3.2 and for real data sets in Section 3.3.

In our personal opinion, the substantial loss of power by the global methods makes the row-wise improved SiZer more useful for data analysis than the global versions. The reason is that away from the null distribution (i.e. when the underlying target curve actually has some interesting structure), the spurious features of the type illustrated in Figures 2 and 3 tend to come up far less frequently than suggested by the size analysis. We view this as an acceptable price to pay for most exploratory data analyses. However, we anticipate that others will disagree, and furthermore recognize situations where statistical rigor is imperative, and thus our software allows a choice between row-wise and global implementations.

In addition to this important row-wise vs. global issue, there are also a number of other issues, such as the impact of smoothing boundary effects, that are also dealt with in Section 2.

2 Improved Distributions

To aid in the development of the distributional properties of SiZer, some basics of kernel smoothing are first reviewed.

Convenient notation for density estimation is X_1, \dots, X_n for a random sample from a probability density $f(x)$. The kernel density estimate of f is

$$\hat{f}_h(x) = n^{-1} \sum_{i=1}^n K_h(x - X_i), \quad (1)$$

where K_h , is a “kernel function”, indexed by a “window-width” h . The estimator $\hat{f}_h(x)$ is simply interpreted as “putting probability mass $1/n$ in a region near each data point”, where the window width controls the critical amount of spread of this mass. The window width h is important enough to appear as a subscript in \hat{f}_h . In all examples in this paper, K_h is taken to be the Gaussian density function, with standard deviation h , because of its very natural scale

space interpretations. It is also important to point out that the scale-space ideas naturally lead to making inference about the smoothed density $\int f(t)K_h(x-t)dt$ rather than about the density $f(x)$. See Chaudhuri and Marron (1999, 2000) for discussion on these subjects.

Our notation for regression data is $(X_1, Y_1), \dots, (X_n, Y_n)$. Such data arise in several ways, and admit several mathematical models. The term “equally spaced design regression” is used to mean that the X_i are deterministic, and equally spaced (in order), and that $Y_i = m(X_i) + \varepsilon_i$, where m is the regression function, and where $\varepsilon_1, \dots, \varepsilon_n$ are independent and identically distributed. While the term “random design” means that $(X_1, Y_1), \dots, (X_n, Y_n)$ are a random sample from a bivariate distribution, with $E(Y_i|X_i) = m(X_i)$, so that again m is the regression function. For random design regression, it can also be useful to think of “residuals”, defined as $\varepsilon_i = Y_i - m(X_i)$. For both settings a common estimator is the local linear smoother, defined at each location x as

$$\hat{m}(x) = a_0, \text{ where } (a_0, a_1) = \arg \min_{a_0, a_1} \sum_{i=1}^n \{Y_i - [a_1(X_i - x) + a_0]\}^2 K_h(x - X_i). \quad (2)$$

This estimator is simply interpreted as providing a local linear fit, in a window centered at x determined by K_h , which is then “moved along” over the range of x values. Again there are many competing estimators, but the local linear smoother is the focus of this paper, for the same reasons as the kernel density estimator. As above, the kernel window function K_h is the Gaussian density function, with standard deviation h .

Because SiZer requires evaluation of a number of smooths (indexed by the window width h), the fast binned implementation discussed in Fan and Marron (1994) is important, especially for larger data sets.

The distributional underpinnings of SiZer are reviewed in Section 2.1. The improved distribution theory of this paper is based on the Gaussian Power distribution, for reasons given in Section 2.2. This distribution is shown to be empirically sensible for the SiZer inference, using simulations in Section 2.3. Useful probability theory for approximation of the Gaussian Power distribution is developed in Section 2.4. Application of the best of these approximations improves the statistical inference in SiZer, as detailed in Section 2.5.

2.1 SiZer distribution theory

Like other hypothesis tests, part of the performance of SiZer is driven by the distribution of SiZer under the null hypothesis of “no signal”. It is desired to set the size of the test, i.e. the probability of “false positives”, to be a small pre-set value α . There are two natural approaches to addressing the multiple comparison problems. The first, called “row-wise” simultaneous inference, seeks to have at most $\alpha 100\%$ of the rows containing “false positives”. The second, called “global” simultaneous inference, aims at having at most $\alpha 100\%$ of the SiZer maps containing false positives.

The “row-wise” problem is analyzed. This means, that a row of a SiZer map is fixed and its distributional properties are considered. After suggesting a solution to the “row-wise” problem in Section 2.5, the rows are combined to obtain the solution of the “global” problem.

Each row of the SiZer map contains colored pixel values, which report the results of a family of hypothesis tests. The distribution theory for each row is that of a sequence of test statistics (modeled as random variables) at each grid point in the domain of the smoother, i.e. at each pixel location in the SiZer map. Let T_1, \dots, T_g , where g is the number of grid points, denote these test statistics. The pixels are equidistant and we can assume without loss of generality that the i -th pixel is in the location $i\tilde{\Delta}$ for some $\tilde{\Delta} > 0$. At the i th pixel in this given row of the SiZer, the color blue (significantly increasing) is used when $T_i > C$, and the color red when $T_i < -C$. The overall size of the row-wise simultaneous SiZer inference will be α when C is chosen such that, under the null distribution of the target curve being constant,

$$P[\{T_i > C \text{ or } T_i < -C\} \text{ for some } i] = \alpha. \quad (3)$$

In the rest of this section we will show that the sequence T_1, \dots, T_g could be approximated by a stationary Gaussian process with mean zero, variance one and correlation

$$\text{corr}(T_i, T_{i+j}) = e^{-j^2/(4b^2)} \left[1 - \frac{j^2}{2b^2} \right], \quad (4)$$

where $b = h/\tilde{\Delta}$ is the bandwidth associated with the SiZer map row expressed on the scale of the pixels. In this approximation, boundary issues that introduce non-stationarity are ignored.

For simplicity it makes sense to think first of equally spaced regression, but our formulas also apply to other settings, including random design regression and density estimation, because these setting have some very strong connections. For some interesting mathematics that makes these connections precise, see Nussbaum (1996), Brown and Low (1996), Brown et. al. (2002) and Grama and Nussbaum (1998, 2002).

When dealing with regression data, SiZer uses the local linear smoother defined by (2). To color the pixels SiZer checks whether the estimate of the first derivative

$$\begin{aligned} a_1 &= -c^{-1} \left[\sum_{i=1}^n K_h(x - X_i) \right] \left[\sum_{i=1}^n (x - X_i) K_h(x - X_i) Y_i \right] \\ &+ c^{-1} \left[\sum_{i=1}^n (x - X_i) K_h(x - X_i) \right] \left[\sum_{i=1}^n K_h(x - X_i) Y_i \right], \quad (5) \\ c &= \left[\sum_{i=1}^n K_h(x - X_i) \right] \left[\sum_{i=1}^n (x - X_i)^2 K_h(x - X_i) \right] \\ &- \left[\sum_{i=1}^n (x - X_i) K_h(x - X_i) \right]^2. \end{aligned}$$

is significantly different from 0. In the particular case of fixed design regression the design points X_i satisfy $X_i = i\Delta$, where $\Delta > 0$ is fixed. If x is away from the boundary, it follows from symmetry of the kernel that

$$\sum_{l=1}^n (x - X_l) K_h(x - X_l) \approx 0. \quad (6)$$

This means that the second term in (5) disappears.

Denote $k = \tilde{\Delta}/\Delta$. The number k is “the number of data points per SiZer column”. For simplicity of notation we can assume that k is a positive integer. This is supported by the fact that SiZer colors the pixel gray if the data are too sparse.

Thus T_j is proportional to the estimate of the first derivative a_1 calculated for $x = j\tilde{\Delta} = jk\Delta$. In particular

$$T_j \approx \sum_{l=1}^n W_{jk-l} Y_l. \quad (7)$$

The exact form of the W_{jk-l} can be read from the first term of (5). For our purpose it suffices to realize that W_{jk-l} is proportional to $-(jk-l) K_{h/\Delta}(jk-l)$. Thus the weights W_i are proportional to the derivative of the Gaussian kernel with standard deviation h/Δ .

If the null hypothesis of no signal is true, Y_i are independent, identically distributed random variables. If additionally the Y ’s have their first two moments, the linear approximation (7) greatly simplifies the distribution theory, because the Central Limit Theorem gives an approximate Gaussian distribution, with mean 0 (under the SiZer null hypothesis) and variance 1, by appropriate scaling. The full joint distribution of T_1, \dots, T_g also depends on the correlation between them. This correlation is approximated by

$$\text{corr}(T_i, T_{i+j}) = \frac{\sum_l W_{l-jk} W_l}{\sum_l W_l^2} \approx e^{-(jk\Delta)^2/(4h^2)} \left[1 - \frac{j^2}{2b^2} \right], \quad (8)$$

where the last step follows from replacing the sums by integral approximations. The equation (4) now follows by observing that $k\Delta/h = \tilde{\Delta}/h = 1/b$.

We only remark that the approximations in equations (6, 7, 8) become limits, if we consider our design to be doubly infinite replacing the sums by $\sum_{l=-\infty}^{\infty}$ (to remove any edge effects), keep h , g and $\tilde{\Delta}$ fixed, and let $\Delta \rightarrow 0$.

In practice we do not know the standard deviation of the noise ε_i . This is needed to scale T_1, \dots, T_g to have variance 1. For this reason it must be estimated from the data introducing additional dependence as well as other issues. However, this is not a problem in theory, as consistent estimators of this standard deviation are available and therefore (7, 8) will still be valid asymptotically (as $\Delta \rightarrow 0$, with h , g and $\tilde{\Delta}$ kept fixed). This is confirmed by our simulation reported in section 3.1, where the estimation of the standard deviation from the data seems to create problems only for small sample sizes at fine scales.

A similar argument can be constructed in the other cases considered here. It is encouraging to see that Figure 7 gives an evidence that the equivalence between the different models considered works for the purposes of SiZer distribution theory.

2.2 Motivation of the Gaussian Power distribution

The row-wise simultaneous inference used in SiZer depends on finding approximate solutions, in C , to the equation (3). Chaudhuri and Marron (1999) used a “number of independent blocks” approach to give a first approximate solution. In this paper, much more precise approximations are developed. These come from

$$\begin{aligned} P[\{T_i > C \text{ or } T_i < -C\} \text{ for some } i] &= P[|T_i| > C \text{ for some } i] \\ &= 1 - P[|T_i| < C \text{ for all } i] \\ &= 1 - P\left[\max_i |T_i| < C\right]. \end{aligned}$$

If T_1, \dots, T_g were independent, then the needed distribution is simply a power of the distribution of the absolute value of a Gaussian random variable, since

$$P[|T_i| < C \text{ for all } i] = \prod_{i=1}^g P[|T_i| < C] = P[|Z| < C]^g, \quad (9)$$

where Z is a standard Gaussian random variable.

The calculation in (9) requires independence of T_1, \dots, T_g . In the dependent, but still stationary Gaussian, case Berman (1964) has found very mild conditions under which the same result holds in the limit as $g \rightarrow \infty$, suggesting that this distribution continue to be considered. However, the rate of convergence is extremely slow, so more careful approximation is needed. Such improved approximations, including careful tuning of the power, are considered in Section 2.4.

The Gaussian power distribution will be used for approximating the distributions of both the row-wise and global maxima. Therefore we first study the appropriateness of the Gaussian power distribution in Section 2.3. Notice that in Section 2.3 we devote our attention solely to the maximum of the whole SiZer map. This is done because there is less theoretical backing for using the Gaussian power distribution for maxima of a spatial random field rather than for a maxima of a stationary random sequence. The parallel analysis for the row-wise maxima has not been done, because we expect to see entirely similar lessons.

2.3 Empirical verification of the Gaussian Power distribution

Here we use the graphical device of the Quantile-Quantile (Q-Q) plot to study how well the Gaussian Power distribution fits the simulated data that was stud-

ied in Figures 2 and 3. See Fisher (1983) for an overview of Q-Q plots and a number of related graphical devices.

The setting is again fixed design regression, for sample size $n = 1600$, based on a 0 regression function, with standard Gaussian noise. For each of 1000 realizations, we compute the maximum over all of the pixels in the SiZer map, of the test statistics used to do inference (i.e. decide on the SiZer color). The distribution of these 1000 maxima is studied in Figure 4, where it is compared to a theoretical Gaussian Power distribution.

The Q-Q plot is essentially a scatterplot, of the data quantiles (just the ordered data values) on the vertical axis vs. the corresponding theoretical quantiles, from the Gaussian Power distribution, on the horizontal axis. Connecting the dots give the red curve. If the theoretical distribution were correct, and there was no sampling variation, the red curve would lie exactly on the 45 degree line, shown in green. Sampling variation leads to some departure from the green line. An important question is whether the amount of variation is explainable by the sampling process, or if it represents a serious departure of the data distribution from the theoretical distribution. This issue is addressed by the family of blue curves, which are 100 simulated Q-Q plots, from data having the theoretical distribution. Note that in this case, the red curve lies nearly completely inside the blue envelope, suggesting that the theoretical fit is good. This shows that the Gaussian Power distribution provides a good fit to the simulation results, and suggests continuing with the ideas for Section 2.2.




FIGURE 4 *QQ Plot showing that a power of Gaussian provides a good fit to the maxima of the 1000 simulated SiZer maps under the null hypothesis. This plot was generated using the same simulated data set as in figures 2 and 3.*

The theoretical distribution considered in Figure 4 is a member of a parametric family. In particular, the Gaussian Power distribution is parametrized by a scale parameter σ (the standard deviation of the underlying Gaussian distribution), and a shape parameter α (the power of the Gaussian c.d.f., i.e. the number of independent Gaussians to maximize). These parameters are estimated in Figure 4 by quantile matching. In particular, they are solutions of the equations that make the Gaussian power distribution correct at the .5 and .75 quantile (these were chosen to give good visual impression).

The estimated value of $\sigma = 1$ is very good, because the underlying Gaussian distribution here has standard deviation 1. The estimated value of $\alpha = 238$ is very interesting, because it is substantially less than 4,400 the total number of pixels in the SiZer map. This shows that the Berman approximation to the null distribution is too crude, and motivates the more refined distributions that are developed in Section 2.4. However, the estimated value of the shape parameter α appears to be unstable, being greatly affected by small changes in the value of σ and the quantiles we decide to match. For example, if we set $\sigma = 1$ we get $\alpha = 552$. Moreover if we then decide to approximately match quantiles .8 and .95 we get $\alpha = 689$. In all of these cases the Q-Q plot shows a reasonable fit.

This phenomenon is related to the “distributional fragility” ideas of Gong et al (2001).

Similar Q-Q plots have been constructed for other simulation settings (detailed in Section 3.1). The results were generally similar (i.e. the Gaussian Power distribution gave a good fit) for the density estimation settings, and for the larger sample sizes. For the smaller sample sizes, in the regression settings, there were no values of the parameters σ and α that left the red curve within the blue envelope. The values that gave the best visual fit, resulted in estimates of σ that were far from 1, and unreasonable values of α . This occasional poor performance seems to be due to t distribution issues, discussed further in Section 3.1.

Another natural distribution for modelling data of this type is the Gumbel distribution, also called the double exponential or the extreme value distribution, see for example Leadbetter, et al (1983). These plots are not shown here, to save space, and because the conclusions are essentially the same as for the Gaussian Power distribution. The fact that both distributions fit well is not surprising because the Gaussian power distribution converges to the Gumbel as $\alpha \rightarrow \infty$.

2.4 Approximation of the Gaussian Power distribution

Let us consider a stationary, mean zero, variance one, Gaussian series T_1, \dots, T_g , with a j step correlation denoted ρ_j . We will be interested in the distribution of $\max(T_1, \dots, T_g)$. As mentioned in section 2.2, Berman (1964) has proved that if $\log(j)\rho_j \rightarrow 0$ the distribution function of $\max(T_1, \dots, T_g)$ behaves asymptotically as the g -th power of the distribution function of a standard Gaussian random variable, i.e.,

$$|P[\max(T_1, \dots, T_g) \leq z] - \Phi(z)^g| \rightarrow 0 \text{ as } g \rightarrow \infty. \quad (10)$$

Unfortunately this approximation is usually of little practical significance as the speed of convergence is very slow. There are at least two alternate approximations in the literature based upon more detailed asymptotics, with the aim of improving the small sample properties of (10).

The first approach, discovered by Rootzen (1983), shows that if the time series is m dependent, if $g(1 - \Phi(x_g)) \rightarrow \kappa$, and if $\max(\rho_1, \dots, \rho_m) > 0$ then

$$P[\max(T_1, \dots, T_g) \leq x_g] - \Phi(x_g)^g \sim e^{-\kappa} R_g \text{ as } g \rightarrow \infty,$$

where R_g is positive quantity depending only on the ρ_j 's, g and κ . The formula for R_g is very complicated, so we will not reproduce it here. An interested reader can consult section 4.6 of Leadbetter, Lindgren and Rootzen (1983) for details.

The second approach is discussed in Hsing, Husler and Reiss (1996). Their approach is based on the observation that for dependent data it is often better to approximate $P[\max(T_1, \dots, T_g) \leq x]$ by $\Phi(x)^{\theta g}$ where $\theta < 1$. Their main idea

is to find θ using asymptotic considerations. In order to get $\theta < 1$ they need the correlation ρ_j to increase with g for each fixed j .

In particular they embed the series in a triangular array $T_{j,\gamma}$, where rows are indexed by γ . For each fixed γ , the random variables $T_{j,\gamma}, j = 1, 2, \dots$ are mean zero, variance one, Gaussian series with the j step correlations $\rho_{j,\gamma}$ satisfying

$$\log(\gamma) (1 - \rho_{j,\gamma}) \rightarrow \delta_j \text{ as } \gamma \rightarrow \infty, \text{ for all } j,$$

where δ_j is a finite positive number. Define

$$\theta = P \left[E/2 + \sqrt{\delta_k} H_k \leq \delta_k \text{ for all } k \geq 1 \right], \quad (11)$$

where E is a standard exponential random variable and H_k is a mean zero Gaussian process independent of E and satisfying $EH_i H_j = \frac{\delta_i + \delta_j - \delta_{|i-j|}}{2\sqrt{\delta_i \delta_j}}$. The authors then claim that under certain technical conditions on $\rho_{j,\gamma}$ the distribution function $P[\max(T_{1,\gamma}, \dots, T_{\gamma,\gamma}) \leq x]$ could be approximated by $\Phi(x)^{\theta\gamma}$. The parameter θ has been called the ‘‘cluster index’’.

Recently Wilhelm (2002) performed an extensive simulation study comparing the three possible approaches for a wide class of stationary Gaussian processes. The simulation study proved inconclusive as neither of the methods clearly dominated the other two. In fact none of the approaches seemed to give reliable answers in the case of highly dependent stationary series. In our simulation study, described in section 3.1, the implementation of the Rootzen method was seen to have even worse performance than the conventional SiZer approach, based on the independent block calculation, whose size characteristics are illustrated in Figure 2. Thus in this paper, we only work with the second method of simultaneous inference, based on the paper of Hsing, Husler and Reiss (1996), which improves the size of SiZer dramatically, as seen in Figure 3. For this reason we will concentrate on the latter method in the remainder of this section.

In the particular case of SiZer, as noted in Section 2.1, it is reasonable to assume that under the null hypothesis T_1, \dots, T_g are Gaussian, with mean 0 and variance 1 and j step correlation $\rho_j = e^{-j^2/(4b^2)} [1 - j^2/(2b^2)]$, where b is the bandwidth expressed at the scale of the SiZer map as defined in Section 2.1. The next step is to embed the series in an appropriate triangular array.

Consider a triangular array $T_{1,\gamma}$ determined by

$$\rho_{j,\gamma} = \left(1 - \frac{\xi}{\log(\gamma)} \right)^{j^2} \left(1 - 2j^2 \log \left(1 - \frac{\xi}{\log(\gamma)} \right) \right). \quad (12)$$

Notice that if we set $\gamma = g$ and $\xi = \log(g) [1 - \exp(-1/(2b)^2)]$ we get $\rho_j = \rho_{j,g}$ and therefore our series can be considered as a part of this triangular array defined by (12). Straightforward calculation shows that $\delta_j = \lim_{\gamma \rightarrow \infty} \log(\gamma) (1 - \rho_{j,\gamma}) = 3j^2\xi$.

Finally, to determine the value of θ we need to calculate the probability in (11). This could be a rather difficult task in general, however in this case we are

helped by the fact that $EH_iH_j = \frac{\delta_i + \delta_j - \delta_{|i-j|}}{2\sqrt{\delta_i\delta_j}} = \frac{i^2 + j^2 - |i-j|^2}{2ij} = 1$ and therefore $Z = H_1 = H_2 = \dots$, where Z is a standard Gaussian random variable. Thus the problem in (11) transforms to

$$\theta = P \left[E/2 + k\sqrt{3\xi}Z \leq 3\xi k^2 \text{ for all } k \geq 1 \right]. \quad (13)$$

Since E is a non-negative random variable, the inequality in (13) implies $Z < \sqrt{3\xi}$. Moreover under this condition $E/2 + k\sqrt{3\xi}Z - 3\xi k^2$ is decreasing as a function of k and therefore

$$\begin{aligned} \theta &= P \left[E/2 + \sqrt{3\xi}Z \leq 3\xi \right] \\ &= E \left(P \left[E \leq 2(3\xi - \sqrt{3\xi}Z) \mid Z \right] \right) \\ &= \int_{-\infty}^{\sqrt{3\xi}} (1 - e^{-2(3\xi - \sqrt{3\xi}Z)}) \frac{e^{-\frac{z^2}{2}}}{\sqrt{2\pi}} dz \\ &= 2\Phi \left(\sqrt{3\xi} \right) - 1. \end{aligned}$$

From here we finally get the cluster index for our SiZer row to be

$$\theta(b) = 2\Phi \left(\sqrt{3 \log(g) [1 - \exp(-1/(2b^2))]} \right) - 1 \quad (14)$$

and consequently we will approximate $[\max(T_1, \dots, T_g) \leq x]$ by $\Phi(x)^{\theta g}$. The verification of the technical conditions of Hsing, Husler and Reiss (1996) is rather tedious and we will not present it here. It can be found in the simpler case of $\rho_j = e^{-j^2/(4b^2)}$ in Wilhelm (2002).

Chaudhuri and Marron (2002) have shown that in a number of real data situations, interesting structure can be found in the data using a curvature based version of SiZer, which sometimes is not flagged as statistically significant by the slope version. Hence, we derive an analogous formula that can be used for this curvature version. Using a similar approximation as for the slope version of SiZer we conclude that under the null hypothesis the curvature SiZer version test statistics $\bar{T}_1, \dots, \bar{T}_g$ are approximately Gaussian, with mean 0 and variance 1 and j step correlation $\bar{\rho}_j = e^{-j^2/(4b^2)} (1 - j^2/b^2 + j^4/(12b^4))$. This leads to $\bar{\delta}_j = \lim_{\gamma \rightarrow \infty} \log(\gamma) (1 - \bar{\rho}_{j,\gamma}) = 5j^2\xi$ and the cluster index of

$$\bar{\theta}(b) = 2\Phi \left(\sqrt{5 \log(g) [1 - \exp(-1/(2b^2))]} \right) - 1. \quad (15)$$

Detailed discussion, with examples, are of some interest. However, they are not included here (except for Figure 11), because the general ideas are the same as for the slope version of SiZer, so it does not seem to be worth the space.

2.5 Proposed Improvements

As mentioned at the beginning of Section 2, there are two natural goals when considering the size of SiZer. The first, called ‘‘row-wise’’ simultaneous inference,

seeks to have at most $\alpha 100\%$ of the rows containing “false positives”, i.e., pixels flagged as statistically significant when no noise is present in the data. The second, called “global” simultaneous inference, aims at having at most $\alpha 100\%$ of the SiZer maps containing false positives.

The row-wise adjustment follows directly from the mathematical considerations of Section 2.4. Define

$$C_R = \Phi^{-1} \left(\left(1 - \frac{\alpha}{2} \right)^{1/(\theta(b)g)} \right),$$

where $\theta(b)$ was defined in (14) and b (defined just after (4)) is the window-width used to smooth the data, scaled so it is compatible with the scale of the T 's. Then color the i th pixel in the g th row blue if the corresponding $T_i > C_R$ and red if $T_i < -C_R$. Notice that under the null hypothesis the distribution of $\max(T_1, \dots, T_g)$ is the same as the distribution of $-\min(T_1, \dots, T_g)$. It follows that if the data contains no signal the probability there is a spurious color on the g th row is

$$\begin{aligned} P [T_i < -C_R \text{ or } T_i > C_R \text{ for some } i = 1, \dots, g] &\leq P [\min(T_1, \dots, T_g) < -C_R] + \\ &\quad P [\max(T_1, \dots, T_g) > C_R] \\ &= 2(1 - P [\max(T_1, \dots, T_g) < C_R]) \\ &\approx 2 \left(1 - \Phi(C_R)^{\theta g} \right) = \alpha. \end{aligned}$$

Thus no more than about $\alpha 100\%$ of the rows will have spurious colors as desired.

We consider two possible global modifications of our row-wise procedure. The first is a simple Bonferroni type adjustment, i.e., if r is the number of rows, we allocate size α/r to each row. In other words we set

$$C_B = \Phi^{-1} \left(\left(1 - \frac{\alpha}{2r} \right)^{1/(\theta(b)g)} \right),$$

and color the i th pixel in the g th row blue if the corresponding $T_i > C_B$ and red if $T_i < -C_B$. Notice that for both the Bonferroni global as well as the row-wise simultaneous adjustment the constants C_R and C_B are different for each row.

The second global adjustment is derived as follows. If all the rows were independent then the distribution of the maxima across all rows and columns would be just a product across all rows, i.e.

$$P \left[\max_{\text{all rows}} \max_{i=1, \dots, g} T_i < x \right] \approx \Phi(x)^{(\theta(b_1) + \dots + \theta(b_r))g},$$

where b_1, \dots, b_r are the window-widths used to smooth the different rows of the SiZer map. Proceeding as though dependence between rows is not too important leads to coloring the i th pixel in the g th row blue if the corresponding $T_i > C_I$ and red if $T_i < -C_I$, where

$$C_I = \Phi^{-1} \left(\left(1 - \frac{\alpha}{2} \right)^{1/((\theta(b_1) + \dots + \theta(b_r))g)} \right).$$

Observe that the constant C_I is the same for all the rows. Of course, since the SiZer rows were calculated from the same data they are likely to be correlated, especially at the coarse scales. On the other hand a careful analysis reveals that the rows are almost independent at the finest scale. This makes this “independent row” approximation less conservative than the Bonferroni method as the fine scales account for most of the sum $\theta(b_1) + \dots + \theta(b_r)$. It is worth noting that a similar effect can be achieved using a Bonferroni type adjustment by allocating a large fraction of the size α to the finer scales.

This explains the main difference between the two global methods. The first (Bonferroni) method is more sensitive at the coarse scales and less sensitive at the fine scales while the second (independent rows) method is less sensitive for the coarse scales and more sensitive for the fine scales. Both of these methods can be considered rather crude and a careful examination of the distributional properties of the random field of the SiZer test considered for all rows and all columns can lead to a slightly less conservative answer. However, the analysis of the random field seems very challenging and our simulations suggest that the independent rows method might be in fact a very good approximation already, for large sample sizes. This is demonstrated in Figure 6 of Section 3.1. In the particular case studied in Figure 4 of section 2.3 the theoretical value of the shape parameter α suggested by the independent rows global method is 882 which is slightly higher than the highest of the various values of the shape parameter estimated there. This affirms our believe that though slightly conservative the independent rows global row-wise method is already rather good.

3 Analysis of Improvements

In this section we investigate the properties of these improvements of SiZer. First the size properties are studied via a simulation study in Section 3.1. The amount of power that is sacrificed to get the size correct, is studied via simulation in Section 3.2, and through some real data examples in Section 3.3.

3.1 Size Simulations

To compare the size performance of the conventional SiZer with our new row-wise SiZer, and with the two new global versions of SiZer, we did an array of simulations against several variations of “the null hypothesis”. We tried:

Settings Each of the settings of:

1. (KDE) kernel density estimation, for the Uniform(0,1) density,
2. (FDR) fixed design regression, for an equally spaced design, with standard Gaussian noise, but no signal,
3. (RDR) random design regression, where the X_i are chosen from the Uniform(0,1) density, and the Y_i are independent standard Gaussian.

Sample sizes For each of the above settings, the following sample sizes were tested:

1. $n = 100$,
2. $n = 400$,
3. $n = 1600$.

For each of the 9 combinations above, 1000 pseudo data sets were drawn, and the various SiZer maps were calculated, and the numbers of red and blue pixels (ideally none, since there are no signals in any of these examples) was counted.

One way to summarize these numbers is row-wise in the SiZer maps. In particular, for each setting, each sample, and each row, report the percentage of realizations of the data where there were some red or blue pixels in that row. Figure 5 shows these summaries. Notice that if no red or blue pixels are present in a particular row, the $\max_{i=1,\dots,g} |T_i|$ for T_i 's corresponding to this row was less than the preset value of the cut-off C . In particular, if we set C using the row-wise approximation of Section 2 and the simulated proportion of red and blue pixels is equal to the nominal value $\alpha = .05$ then we have some evidence that our approximation is working. As seen in Figure 5 this is often the case.

Instead of showing long tables of numbers, the main ideas are made more accessible by displaying the results with a parallel coordinate plot, see Inselberg (1985). Figure 5a summarizes performance for the Kernel Density Estimation setting, Figure 5b does the same for the Fixed Design Regression setting, and Figure 5c is for the Random Design Regression setting. The coordinates (points on the horizontal axes) represent rows of the SiZer map, and thus are quantified via $\log_{10} h$ (only shown on the bottom panel, to avoid overplotting with the Figure titles), just as on the vertical axes of the SiZer maps. The vertical axes are the percentage of rows (across the 1000 replications) that showed some significant structure (i.e. red or blue pixels). Each curve represents one setting (indicated by color as shown) and one sample size (indicated by line type as shown). The curves are piecewise linear, with nodes at each row of the map (i.e. each window width h). The heights at the nodes contain the useful information, and the connecting line segments simply make it easier to understand the relationships.

Ideally, all of these values should be close to $\alpha = 0.05$ for the row-wise procedures such as the conventional SiZer, and our new row-wise SiZer. Hence, this level is shown by a horizontal black line.




FIGURE 5 Row-wise summaries of the percent of significant pixels for SiZer under the null hypothesis, allowing comparison of the different simultaneity adjustments and sample sizes. Clearly shows relationship between sample sizes. Figures 5a, b and c are for the settings of KDE, FDR and LDR, respectively.

Note that in almost every case the conventional SiZer flags significant structure far too often. This again verifies the main idea in this paper: it is well worth finding less crude approaches to this multiple comparison problem.

Similarly, in a large majority of the cases, the new Row-Wise SiZer is quite close to the desired $\alpha = 0.05$.

As expected, the global methods are almost always quite far below the desired level, because they aim at a global size of $\alpha = 0.05$, which requires them to be deliberately conservative when studied in this row-wise sense.

A perhaps surprising feature in the KDE setting, studied in Figure 5a, is the 0 values everywhere for the second and third coarsest scales. This is due to the crude type of boundary adjustment used. Boundary adjustment is essential for estimating the Uniform(0,1) density with kernel estimates, because these methods tend to “round off the corners” at both edges. If the summaries of Figure 5 are computed with no adjustment, far too many percentages are %100, since every realization of most rows has some significant pixels flagged at the edges. To avoid this boundary problem, the simple “circular design” device was used (where the data are treated as periodic, and shifted copies of the data are added at each end). While this crude adjustment is reasonably effective at most scales, there are a few where it introduces artifacts such as the zeros shown in Figure 5a. Such boundary effects are not a serious issue for the regression settings, because the local linear smoother that is used in both does an automatic first order boundary adjustment.

Another departure from the expected size occurs for the regression settings, shown in Figures 5b and 5c. These are substantial increases in the percentage of realizations flagged as significant at finer scales. At these scales, there can be few points in the kernel window, so that the underlying null distributions are better approximated by a t distribution, than by the Gaussian. This idea is verified by the fact that it is generally worst for $n = 100$, better for $n = 400$, and the problem is nearly nonexistent for $n = 1600$. Exceptions include the FDR in Figure 5b, where the dotted curves for $n = 100$ disappear for fine scales (because there are never enough data points in the kernel windows, i.e. the SiZer color is always gray), and the RDR in Figure 5C, where the dotted curves for $n = 100$ actually go down for finer scales, because there are typically just a few locations where the data are rich enough to do any inference (thus most of the pixels are colored gray), and in those remaining locations the SiZer color is often completely purple.

A simple approach to this problem is to replace the Gaussian distribution with the t distribution. This was attempted, but the results were too conservative to be useful. The reason seems to be the complicated interaction of the t distribution with the correlation structure.

The comparison in Figure 5 is for the row-wise size of the statistical inference. But also of keen interest is the global size, for the multiple comparison problem over the entire map, not just within individual rows. Global size, for the same simulation settings, is studied in Figure 6.

Figure 6 is a parallel coordinate display of the percent of realizations (out of 1000) for which there were some significant pixels in the SiZer map. Again color is used to indicate SiZer type, with the same color scheme. The coordinates now are taken to be the sample size n , different from SiZer map row as in Figure 5), to highlight the perhaps surprising impact of n on the results. Line type is now used to show the setting.




FIGURE 6: *The global size summaries showing the percent of significant pixels in the full SiZer maps, under the null hypothesis, grouped by settings.*

In this sense, the size problems of the conventional SiZer map are even worse than in the row-wise sense indicated in Figure 5 (note the larger vertical axis). The new Row-Wise SiZer is also always far above the nominal level of $\alpha = 0.05$, which not surprisingly shows that there is substantial difference between row-wise and global statistical inference. This is consistent with the global methods appearing as generally too conservative in Figure 5.

Performance of the global SiZer approaches, is quite dependent on the setting. For Kernel Density Estimation, both methods are generally conservative. This is caused by the boundary effect and adjustment discussed above, and by data sparseness issues at the finest scales. In particular, the 0's at the second and third coarsest scales mean that the “effective number of rows” r , is essentially too large in our calculations. For Fixed Design Regression and Random Design Regression, the percentages are often too large. The exception is $n = 1600$, where the Global Independent SiZer has excellent size performance (for both FDR and RDR), and the Bonferroni is a little conservative (both as expected). For $n = 400$, the percentage of maps flagged as significant increases substantially, because of the t effect described above (most of which occurs at the finest scales where there are relatively few points in each kernel window, so the number of degrees of freedom can be as low as 4). As noted above, many of the curves are lower for $n = 100$, because of data sparsity effects.

Figure 7 is a reorganization of the parallel coordinates plot in Figure 5, which highlights an important lesson about how the settings compare, that is obscured there because the settings are in different panels. This time the panels show the sample sizes n , with $n = 100, 400, 1600$ in Figures 7a, b and c respectively. As in Figures 5 and 6, color represents SiZer type, using the same scheme. The line type is consistent with Figure 6, representing the setting. Again the coordinates represent rows of the SiZer map, and are indexed by $\log_{10} h$.




FIGURE 7 *Row-wise summaries of the percent of significant pixels for SiZer under the null hypothesis, allowing comparison of the different simultaneity adjustments and settings. This organization shows that settings are very similar. Figures 7a, b and c are for the sample sizes of $n = 100$, 400 and 1600, respectively.*

The main lesson of Figure 7 is that curves of the same color tend to be very close to each other, i.e. the settings are very similar. While there are important differences in the simultaneity type (expressed by colors), and the sample sizes (different panels), the settings are similar. This validates the approach of using the common mathematical structure, as developed in Section 2.1.

Another useful feature of the view shown in Figure 7, is that it provides another way of seeing that the Row-wise method is best in this sense, and that the best results are for the larger sample sizes. In particular, it is very clear that

at $n = 1600$, the percentages virtually achieve their goal of $\alpha = 0.05$, uniformly over both rows and settings (except for density estimation at large scales).

A similar simulation study has been carried out to investigate the size properties of the curvature version of SiZer. The results were similar to those summarized in Figures 5 - 7 for the slope version of SiZer and are not explicitly reported to save space. The main differences between the results were that both the boundary effect in the kernel density estimation and the t effect for the small sample sizes of regression were even more severe in the curvature version than in the slope version.

3.2 Power Simulations

The previous section showed that our global versions of SiZer were quite good at achieving the desired overall size for the statistical inference. In this and the next section, by analyzing some simulated and real data sets, it is seen that this entails substantial cost in terms of power.

The first example is the same as shown in Figure 1, the Donoho Johnson blocks regression function, with high noise, as shown in Figure 1a. Figure 8 allows direct comparison between the conventional SiZer shown in Figure 8a, the new row-wise SiZer shown in Figure 8b, the Bonferroni Global SiZer shown in Figure 8c and the Global Independent Row SiZer shown in Figure 8d.




FIGURE 8: *Full range of SiZer analyses of the Donoho - Johnstone Blocks regression, with high noise. Figures 8a, b, c and d show conventional, row-wise, Bonferroni and independent rows global SiZer versions, respectively.*

As shown in Figure 1b, the conventional SiZer flags all 11 jumps as statistically significant, but it also indicates a spurious jump near $x = 0.58$. As expected, the new row-wise SiZer (Figure 8b) flags fewer pixels as significant, but still finds all 11 jumps. The spurious jump near $x = 0.58$ is still present, but smaller. For the two global methods the spurious feature disappears, but also the jump near $x = 0.15$ disappears as well. This reflects the loss of power from insisting on global simultaneous inference.

If one were to use only the global analysis, the upward jump near $x = 0.78$, would be flagged as statistically significant by a very small blue region. Thinking from the viewpoint of conventional SiZer, it might be tempting to ignore this. However an important lesson is that any significant pixel (regardless of how small it is) at all that is found by a global method, should be regarded

as important underlying structure.

Figure 9 shows a simulated density estimation example, with the same four panels as in Figure 8. The underlying density is the Trimodal Gaussian Mixture Density from Marron and Wand (1992), and the sample size is $n = 10,000$. Both the conventional and new row-wise SiZer show three statistically significant modes. However, the conventional SiZer also flags a spurious fine scale feature near $x = 1.4$, which correctly disappears for the new row-wise version. The global SiZers show some loss of significant structure, in particular the small blue region just left of $x = 0$, again reflecting some loss of power.

Figure 9 also shows a feature predicted by the analysis of Section 2.5, that the independent rows SiZer is more sensitive at the finer scales than the Bonferroni version.




FIGURE 9: Full range of SiZer analyses of the Trimodal mixture of Gaussians. Figures 9a, b, c and d show conventional, row-wise, Bonferroni and independent rows global SiZer versions, respectively.

Similar plots have been constructed for all of the Marron-Wand Gaussian mixture densities, for the sample sizes $n = 100, 1000, 10,000$. Overall, the different versions of SiZer tended to flag very similar structure as being statistically significant. There was generally substantial erosion of the red and blue regions for the methods with better size properties (to a similar extent to that shown in Figure 9). Sometimes this erosion was enough that significant features actually disappeared, as in Figures 9c and d, but most often they did not. Spurious features, such as the very small red region, near $x = 1.4$ in Figure 9a, were fairly rare, perhaps because at most locations, these densities are not close to flat (as at the null distributions studied in Section 3.1), but instead have substantial slope.

3.3 Real Data Examples

Another approach to studying the trade-off between size and power that is made by these different versions of SiZer is through the analysis of real data. Figure 10 shows the density estimation example of the 1975 British Family Incomes data, that was carefully analyzed by Schmitz and Marron (1992), again using the same 4 panels of Figures 8 and 9. The conventional SiZer analysis shows two significant modes, which has been independently confirmed by a parametric analysis as discussed in Schmitz and Marron (1992). The red region between modes nearly disappears for the new row-wise SiZer shown in Figure 10b, although again, greater credence needs to be placed in this more precise version. Unfortunately this red region completely disappears in both global SiZer maps. This loss of power is particularly unfortunate, since the bimodality is the important feature of this data set.




FIGURE 10: *Full range of SiZer analyses of the British Family Incomes data. Figures 10a, b, c and d show conventional, row-wise, Bonferroni and independent rows global SiZer versions, respectively.*

While the global slope versions were unable to find the important bimodal characteristics of the British Family Incomes data in Figure 10, it is interesting to note that the global curvature versions of SiZer do flag this feature of the data as statistically significant, as shown in Figure 11. The conventional curvature version of SiZer was proposed by Chaudhuri and Marron (2002). Here we similarly improved the simultaneity, using the ideas from Section 2. The main difference is that formula (15) replaces (14).

To clearly distinguish it from the slope version of SiZer, the curvature version uses a different color scheme. Pixels with significant concavity (second derivative strongly negative) are indicated by cyan (light blue). Those with significant convexity are colored orange. Locations in scale space where there is no significant curvature are colored green. Again gray is used in regions

where the data are too sparse.

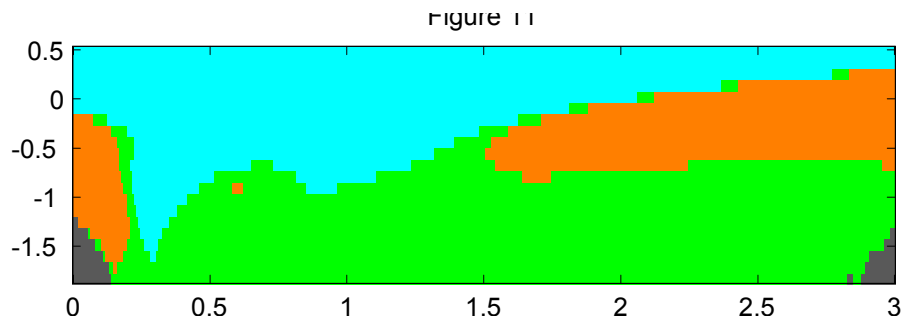


FIGURE 11: *Independent rows global curvature SiZer analyses of the British Family Incomes data. This finds the bimodality that is known to be an important feature of this data set.*

The bimodality of this data set is shown to be strongly significant, by the very small orange region near $x = 0.6$. While the region is very small, again it is important to keep in mind that when using global versions of SiZer, any significance at all should be regarded as strong evidence.

Figure 12 shows an example from flow cytometry, where the presence and percentage of fluorescence marked antibodies on cells are measured. The medical goal is the determination of quantities such as the percentage of lymphocytes among cells. The data come from the laboratory of Drs. S. Mentzer and J. Rawn, Brigham and Women’s Hospital, Boston, Massachusetts, and we are grateful to M. P. Wand for putting us in contact with them. In a single experiment, many cells are run through a laser, and the intensity of fluorescence of each cell is measured, and the data are stored as 256 bin counts, where bins are called “channels”. These bin counts are traditionally viewed on the square root scale. An important question is how many “bumps” there are in this square root histogram. Here we treat this as a regression problem.

Figure 12 shows again the same 4 panels, comparing the different simultaneity methods. Figure 12a, conventional SiZer, shows two clear modes, and a small fine scale feature near $x = 20$. This small feature is already seen to be spurious by the new row-wise SiZer map in Figure 12b. This time the effect of the global is representative of many of the examples we have seen: the significant red and blue regions are substantially eroded, but indicate essentially the same lessons.

Figure 12c confirms one more point that was predicted in Section 2.5, the Global Bonferroni method is more powerful than the Global Independent Rows approach for larger scales. This appears as a larger blue region on the left side.




FIGURE 12: *Full range of SiZer analyses of a flow cytometry data set. Figures 11a, b, c and d show conventional, row-wise, Bonferroni and independent rows global SiZer versions, respectively.*

Based on this experience, and a number of other examples studied during this research, we recommend that the default version of SiZer be the new row-wise approach. This choice is made to give reasonable power, but it needs to be kept in mind that the statistical inference is not completely valid in the classical sense, which is often acceptable in exploratory data analysis situations. When statistical rigor is essential (e.g. before making a large investment of research effort in understanding “phenomena found”) it is recommended that the global versions be used. Both global versions are useful, because they have different sensitivities. The Bonferroni Global method is somewhat more powerful at finding structure at coarser scales, while the Independent Row Global version has more power at finer scales.

4 Future Work

While the methods developed in this paper are intended to enhance the applicability of the SiZer method, there are a number of remaining open problems, including:

1. Replace the independent rows global approximation, of Section 2.5, by an approximation that takes the full random field distribution of the SiZer inference into account.
2. More careful boundary adjustment, as discussed in Section 3.1.
3. Improved incorporation of the t distribution, for regression settings, with careful accounting of the correlation structure, as discussed in Section 3.1.

References

- [1] Berman, S. (1964) Limit theorems for the maximum term in stationary sequences, *Annals of Mathematical Statistics*, 35, 502-516.
- [2] Brown, L. D. and Low, M. (1996). Asymptotic equivalence of nonparametric regression and white noise, *Annals of Statistics*, 24, 2384-2398.
- [3] Brown, L. D., Cai, T. T. and Low, M. (2002). Asymptotic equivalence theory for nonparametric regression with random design. *Annals of Statistics*, 30, 688-707.
- [4] Chaudhuri, P. and Marron, J. S. (1999) SiZer for exploration of structure in curves, *Journal of the American Statistical Association*, 94, 807-823.
- [5] Chaudhuri, P. and Marron, J. S. (2000) Scale space view of curve estimation, *Annals of Statistics*, 28, 408-428.
- [6] Chaudhuri, P. and Marron, J. S. (2002) Curvature vs. Slope Inference for Features in Nonparametric Curve Estimates, unpublished manuscript.
- [7] Donoho, D. L. and Johnstone, I. M. (1994). Ideal spatial adaptation by wavelet shrinkage, *Biometrika*, 81, 425-455.
- [8] Fan, J. and Gijbels, I. (1996) *Local polynomial modeling and its applications*, Chapman and Hall, London.
- [9] Fan, J. and Marron, J. S. (1994) Fast implementations of nonparametric curve estimators, *Journal of Computational and Graphical Statistics*, 3, 35-56.
- [10] Fisher, N. I. (1983) Graphical Methods in Nonparametric Statistics: A Review and Annotated Bibliography, *International Statistical Review*, 51, 25-58.

- [11] Gong, W., Liu, Y., Misra, V. and Towsley, D. (2001) On the tails of web file size distributions, *Proceedings of 39-th Allerton Conference on Communication, Control, and Computing*. Oct. 2001. Internet available at: <http://www-net.cs.umass.edu/networks/publications.html>.
- [12] Grama, I. and Nussbaum, M. (1998) Asymptotic equivalence for nonparametric generalized linear models, *Probability Theory and Related Fields*, 111, 167-214.
- [13] Grama, I. and Nussbaum, M. (2002) Asymptotic equivalence for nonparametric regression, *Mathematical Methods of Statistics*, 1-36.
- [14] Hsing, T., Husler, J. and Riess, R. D. (1996) The extremes of a triangular array of normal random variables, *Annals of Applied Probability*, 6, 671-686.
- [15] Inselberg, A. (1985) The plane with parallel coordinates, *The Visual Computer*, 1, 69-91.
- [16] Leadbetter, M. R., Lindgren, G. and Rootzen, H. (1983) *Extremes and Related Properties of Random Sequences and Processes*, Springer Verlag, Berlin.
- [17] Lindeberg, T. (1994) *Scale-Space Theory in Computer Vision*, Kluwer, Dordrecht.
- [18] Marron, J. S. and Wand, M. P. (1992) Exact mean integrated squared error, *Annals of Statistics*, 20, 712-736.
- [19] Marron, J. S. (1996) A personal view of smoothing and statistics, in *Statistical Theory and Computational Aspects of Smoothing*, eds. W. Härdle and M. Schimek, 1-9 (with discussion, and rejoinder 103-112).
- [20] Marron, J. S., Adak, S., Johnstone, I. M. Neumann, M. and Patil, P. (1998), Exact risk analysis of wavelet regression, *Journal of Computational and Graphical Statistics*, 7, 278-309.
- [21] Jones, M. C., Marron, J. S. and Sheather, S. J. (1996a) A brief survey of bandwidth selection for density estimation, *Journal of the American Statistical Association*, 91, 401-407.
- [22] Jones, M. C., Marron, J. S. and Sheather, S. J. (1996b) Progress in data-based bandwidth selection for kernel density estimation, *Computational Statistics*, 11, 337-381.
- [23] Nussbaum, M. (1996) Asymptotic equivalence of density estimation and Gaussian white noise, *Annals of Statistics*, 24. 2399-2430.
- [24] Rootzen, H. (1983) The rate of convergence of extremes of stationary normal sequences, *Advances in Applied Probability*, 15, 54-80.

- [25] Scott, D. W. (1992) *Multivariate density estimation, theory, practice and visualization*, John Wiley: New York.
- [26] Schmitz, H. P. and Marron, J. S. (1992) Simultaneous estimation of several size distributions of income, *Econometric Theory*, 8, 476-488.
- [27] Silverman, B. W. (1986) *Density estimation for statistics and data analysis*, Chapman and Hall: London.
- [28] ter Haar Romeny, B. M. (2001) *Front-End Vision and Multiscale Image Analysis*, Kluwer Academic Publishers, Dordrecht, the Netherlands.
- [29] Wand, M. P. and Jones, M. C. (1995) *Kernel Smoothing*, Chapman and Hall: London.
- [30] Wilhelm, J. R. (2002) *A simulation study on competing distributions for the maxima of stationary normal processes*, M. S. Thesis, Department of Statistics, Colorado State University.