

# *SIZER* FOR TIME SERIES: A NEW APPROACH TO THE ANALYSIS OF TRENDS

V.Rondonotti and J.S.Marron \*

January 5, 2004

## Abstract

Smoothing methods and *SiZer* are a useful statistical tool for discovering statistically significant structure in data. Based on scale space ideas originally developed in the computer vision literature, *SiZer* (SIGNificant ZERO crossing of the derivatives) is a graphical device to assess which observed features are ‘really there’ and which are just spurious sampling artifacts.

In this paper, we develop *SiZer* like ideas in time series analysis to address the important issue of significance of trends. This is not a straightforward extension, since one data set does not contain the information needed to distinguish ‘trend’ from ‘dependence’. A new visualization is proposed, which shows the statistician the range of trade-offs that are available. Simulation and real data results illustrate the effectiveness of the method.

## 1 Introduction

Smoothing methods in statistics provide a useful tool for showing structure in data. Many monographs on smoothing are available in the statistical literature which in the last years include Green and Silverman [1994], Wand and Jones [1995], Simonoff [1996], Fan and Gijbels [1996], Hart [1997] and Bowman and Azzalini [1997].

---

\*V.Rondonotti is Economist-Statistician, Money and Banking Statistics Division, European Central Bank, Frankfurt am Main, Germany. Her work was supported in part by Università degli Studi di Roma ‘La Sapienza’, Dipartimento di Statistica, Probabilità e Statistica Applicata. J.S.Marron is Professor, Department of Statistics, University of North Carolina, Chapel Hill, NC 27599-3260. His research was supported by NSF Grant DMS-9971649. This work is from the PhD dissertation of the first author, under direction of the second.

When smoothing methods are used for exploratory data analysis, a question that often arises is: which features are ‘really there’ (i.e. represent important underlying structure) and which are just spurious sampling artifacts (i.e. reflect unimportant random variation)? Chaudhuri and Marron [1999] proposed answering this question with a graphical device called *SiZer* (Significant ZERo crossings of the derivatives), which is based on scale space ideas from computer vision.

Scale space is a family of kernel smooths indexed by the ‘scale’  $h$ , which is usefully viewed as the level of resolution of the data. The scale  $h$  is usually called the ‘smoothing parameter’, or the ‘bandwidth’ in the statistical literature. Marron and Chung [2001], for example, suggest using an overlay of these smooths for data analysis, and call it the ‘family plot’. This method is illustrated in Figure 1 (a), using the Chocolate data set, i.e. the monthly production of chocolate in Australia from the late 50’s to the early 90’s (kilotonnes). This data set comes with the software companion to Brockwell and Davis [1996]. The green dots show the Chocolate production after deseasonalising and linearly detrending the time series. The family of smooths suggests at different levels of resolution a dip around the year 78 preceded by two minor bumps and an increase in the last years. At the finest levels of resolution, many other features appear. Are all these features ‘really there’? *SiZer* answers this question by assessing statistical significance of such features.

In particular, *SiZer* extends the usefulness of the family plot by visually displaying the statistical significance of features over both location  $x$  and scale  $h$ . *SiZer* is based on confidence intervals for the derivative of the underlying function. The graphical device is a color map, reflecting statistical significance of the slope at  $(x, h)$  locations in scale space. At each  $(x, h)$  location, the curve is significantly increasing (decreasing) if the confidence interval is above (below) 0, so that map location is colored blue (red). If the confidence interval contains 0, the curve at the level of resolution  $h$  and at the point  $x$  does not have a statistically significant slope, so the intermediate color of purple is used. Finally, if there is not enough information in the data set (according to a rule that will be illustrated in the following sections), at this scale space  $(x, h)$  location, then no conclusion can be drawn, so the color gray is used to indicate that the data are too sparse.

The *SiZer* view of the Chocolate data is shown in Figure 1 (b). This data set was studied in Marron and Chaudhuri [1998a], who viewed the errors as independent and identically distributed (i.i.d.).

(put Figure 1 (a) and (b) approx. here)  
(caption: Exploratory analysis of the Chocolate data set through the *SiZer* approach: (a) The Family Plot; (b) The *SiZer* Plot)

At the coarsest levels of resolution (largest bandwidths), at the top of the *SiZer* map, the color purple appears everywhere, indicating that there is no statistically significant increase or decrease in the corresponding smooths. As it can be seen Figure 1 (a), these curves are close to the simple linear regression line, so the conclusion is that this line has no significant slope. Moving down in the plot, meaning decreasing the scale (i.e. the bandwidth gets smaller) the smooths are first significantly decreasing (red), then significantly increasing (blue) after a short time interval in which the smooths are neither significantly increasing nor decreasing (purple). This shows that the minimum near year 1978 is statistically significant. As the bandwidth gets smaller, i.e. further down the plot, some additional significant features appear but only for a limited number of bandwidth values. At the finest level of resolution (smallest bandwidth) there is not enough information to assess the significance of any structure in the data set (gray).

Overall, the *SiZer* map reveals that only the decrease followed at around year 78 by the increase in chocolate production are important features of the data.

As pointed out by Chaudhuri and Marron [1999], the statistical inference, which is the basis of *SiZer*, makes heavy use of the assumption of i.i.d. errors. This assumption is inappropriate in time series contexts, where dependence is omnipresent, and in fact is usually the focus of statistical analysis. For *SiZer* to fulfill its potential to flag significant trends in time series, its underlying confidence intervals must be adjusted to properly account for the correlation structure of the data. This adjustment is straightforward when the correlation structure is known. But in the more important and common case of unknown correlation, *SiZer* for time series is not a straightforward extension. This is because of the identifiability problem between ‘trend’ and ‘dependence artifacts’. Distinction between these can never be made on the basis of a single time series, as shown in Figure 2.

(put Figure 2 approx. here)  
(caption: An example for the identifiability problem between ‘trend’ and ‘dependence artifacts’: the Strikes data set)

Figure 2 (a) shows the Strikes data from Brockwell and Davis [1991]. One view of this series is that its underlying distribution is

nonstationary. In particular the mean function seems to decrease at the beginning, with an increase in the middle, and perhaps another decrease later. The variance also appears to be time varying, being large at both ends, and smaller in the middle. But this apparent ‘nonstationarity’ can not be proven on the basis of that single time series, as there are stationary stochastic processes whose realizations can look quite similar. An example is the simulated series shown in Figure 2 (b), which is a single realization of a stationary stochastic process. Note that the qualitative features, in terms of mean and variance, are very similar to those of Figure 2 (a). The remaining parts of Figure 2 reveal how this simulation was performed. The particular series of Figure 2 (b) was carefully selected from the much longer series shown in Figure 2 (d). It is the part between the vertical bars, chosen to match qualitatively the features of Figure 2 (a). The long series in Figure 2 (d) was generated as  $X_t = \mu_t + \sigma_t Z_t$  where the curves  $\mu_t$  and  $\sigma_t$  are generated in the same way as drivers of the intensity function of a Cox process (shown in Figure 2 (c)), and where the  $Z_t$  are independent standard normal realizations. Because  $\mu_t$  and  $\sigma_t$  are stationary stochastic processes, so is  $X_t$ .

Now consider this example from the viewpoint of ‘trend’. The time series in Figures 2 (a) and (b) can be viewed as containing a large amount of ‘trend’ in the mean, but that structure can be equally well explained as ‘serial dependence artifacts’, or in fact some of each. The challenge for applications of *SiZer* in the trend estimation context is that these cases cannot be distinguished. This paper proposes an approach to this dilemma via a visualization which displays the range of trade-offs to the statistician.

Background on the local linear fit and on the *SiZer* approach will be found in Section 2. In Section 3 we will discuss the development of *SiZer* for studying trends in time series. The performance of our method is studied via real and simulated data sets in Section 4.

## 2 Background

### 2.1 The local linear fit

One interesting approach to smoothing is local polynomial modelling (see e.g. Fan and Gijbels [1996] for a detailed overview of this method). Here are the main ideas in the time series context for trend estimation.

Trend estimation in time series can be viewed as a regression problem, for fixed and equispaced design, with correlated errors.

Given the data  $\{(i, Y(i)), i = 1, \dots, n\}$ , the regression setting is

$$Y_i = m(i) + \sigma\epsilon_i \quad i = 1, \dots, n \quad (1)$$

where  $m$  is assumed to be a smooth function and the error is assumed to be a zero mean weakly stationary process, i.e.

$$E(\epsilon_i) = 0 \quad \forall i, \quad (2)$$

$$V(\epsilon_i) = 1 \quad \forall i, \quad (3)$$

and

$$Cov(\epsilon_i, \epsilon_j) = \gamma(|i - j|) \quad \forall i, j. \quad (4)$$

Of interest is the estimation of the regression function  $m$  and its derivatives at  $i_0$ , i.e.  $m(i_0) = E(Y|i = i_0)$  and  $m'(i_0)$ ,  $m''(i_0)$ , ...,  $m^{(p)}(i_0)$ .

The flexibility of smoothing techniques comes from only assuming smoothness of the underlying curve. This means, intuitively, that data information about the value of the regression function  $m$  at  $i_0$  is given by the observations at  $i$  near  $i_0$  and therefore these observations can be used to construct an estimator of  $m(i_0)$ .

By assuming that the  $(p + 1)^{th}$  derivative of  $m(i)$  at the point  $i_0$  exists, the local polynomial approach approximates the regression function  $m(i)$  locally by a polynomial of order  $p$ . This polynomial is then fitted locally by solving a weighted least squares regression problem. In particular, in the local linear fit the function  $m(i)$  is approximated by Taylor expansion of order 1 for  $x$  in the neighborhood of  $x_0$ . The problem to be solved is then

$$\min_{\beta} \sum_{i=1}^n [Y_i - (\beta_0 + \beta_1(i - i_0))]^2 K_h(i - i_0) \quad (5)$$

where  $\beta = (\beta_0 \beta_1)'$ ,  $h$  is the bandwidth controlling the size of the local neighborhood and  $K_h(\cdot) = \frac{1}{h} K(\frac{\cdot}{h})$ , where  $K$  is a kernel function, often taken to be a symmetric probability density, assigning weights to each datum point.

By Taylor expansion  $\beta_0 = m(i_0)$  and  $\beta_1 = m'(i_0)$ , so the solution to this problem gives estimates of the regression function and its first derivative at  $i_0$ .

Because of the excellent interpretability properties of kernel type smoothers and because of their simplicity, Marron and Chung recommend the local linear fit to construct the family plot. Moreover,

for a better visualization of the family of smooths, it is suggested to use a ‘very wide range’ of  $h$  values in the log scale and for symmetry an odd number of curves should be chosen (for details see Marron and Chung [2001]).

## 2.2 The *SiZer* approach

To understand the development of *SiZer* for time series let us briefly illustrate the mathematical aspects of the original version of *SiZer* for regression function estimation in the specific case of fixed and equispaced design but where errors are i.i.d.. As previously stated, the *SiZer* approach is based on confidence intervals for the derivative of the underlying function. These are of the form

$$\hat{m}'_h(i) \pm q(h)\hat{sd}(\hat{m}'_h(i)) \quad (6)$$

where  $q(h)$  is an appropriate Gaussian quantile. There are two critical points of interpretation.

The first point is that from the scale space view, the center point of each confidence interval is automatically ‘correct’, i.e. the interval is unbiased. This requires replacing the regression curve  $m(i)$  as ‘target’, with its scale  $h$  smoothed version. This makes sense because it reflects the part of the regression curve that is available from the data at the level of resolution,  $h$ . I.e. for each value of the bandwidth, all the information available in the data set is considered in the inference process. For more details see Chaudhuri and Marron [1999].

The second point is that the confidence intervals are constructed in a simultaneous way. The approach is, for each level of resolution  $h$ , to approximate the full simultaneous confidence limit problem by  $m$  independent confidence interval problems. The quantity  $m$ , which reflects the number of ‘independent blocks’, is estimated through the quantity *Effective Sample Size* (*ESS*), i.e. for the bandwidth  $h$  and at  $i_0$

$$ESS(i_0, h) = \frac{\sum_{i=1}^n K_h(i_0 - i)}{K_h(0)}. \quad (7)$$

Note that when using the uniform kernel,  $ESS(i_0, h)$  is equal to the number of data points in the kernel window centered at  $i_0$  with the bandwidth  $h$ . The number of independent blocks is then approximated by the quotient:

$$m(h) = \frac{n}{ESS(i_0, h)}. \quad (8)$$

This results in the quantile

$$q(h) = \Phi^{-1}\left(\frac{1 + (1 - \alpha)^{\frac{1}{m(h)}}}{2}\right). \quad (9)$$

$ESS$  is also considered to decide when the normal approximation is inadequate, i.e. for  $ESS < 5$  no conclusion can be drawn. In particular, the blue, red, purple colors are used on the scale space set

$$\{(i_0, h) : ESS(i_0, h) \geq n_0\} \quad n_0 = 5. \quad (10)$$

Again detailed discussion of *SiZer* is available in Chaudhuri and Marron [1999].

### 3 *SiZer* for time series

For time series trend estimation, i.e. in the context of a regression problem with fixed and equispaced design with correlated errors, *SiZer* can be developed by modelling the error structure and by adjusting the confidence limits for the derivative of the underlying function according to the assumed error structure.

As suggested by Marron and Chung [2001], the family of smooths is constructed by considering a ‘very wide range’ of bandwidths  $h$  in the log scale and, in particular, the number of curves is here taken to be 11.

#### 3.1 The variance

For correlated data, the variance of the local polynomial estimator is given by

$$V(\hat{\beta}|X) = (X^T W X)^{-1} (X^T \Sigma X) (X^T W X)^{-1} \quad (11)$$

where the design matrix of the local linear fit at  $i_0$  is

$$X = \begin{pmatrix} 1 & (1 - i_0) \\ 1 & (2 - i_0) \\ \vdots & \vdots \\ 1 & (n - i_0) \end{pmatrix}$$

where  $W = \text{diag}\{K_h(i - i_0)\}$  and where, for the assumed correlation structure,  $\Sigma$  is the kernel weighted covariance matrix of the data where the generic element is given by

$$\sigma_{ij} = \gamma(|i - j|)K_h(i - i_0)K_h(j - i_0). \quad (12)$$

A sensible estimate of the variance (11) is based on estimating  $\gamma$ , by the sample autocovariance function of the observed residuals from a “pilot smooth”, using the pilot bandwidth  $h_p$ . One could take  $h_p = h$ , but this leads to a confounding of the different notions of “scale” and “dependence structure”. Hence, we treat  $h$  and  $h_p$  separately, which means that in the dependent case, another dimension needs to be added to the *SiZer* plot. We approach this via a series of *SiZer* plots, indexed by the pilot bandwidth  $h_p$ , which represent the different trade-offs available between trend and dependence. This is the key to our visualization, which is further developed in Section 4.

### 3.2 The quantile

For positively autocorrelated errors, the family of smooths typically varies substantially as the bandwidth is varied. The reason is that the strong sporadic patterns that characterize such errors behave similarly to high frequency regression components, which appear in the smooth for a wide range of bandwidths. On the other hand, for negatively autocorrelated errors, the tendency of data points to alternate above and below the regression function gives a family of smooths which changes less as a function of the bandwidth. This is one way of seeing that the amount of information, about the underlying smooth regression function, that is available in i.i.d. data, is not the same as the amount of information available in correlated data. Positively correlated data contain ‘less information’ about the regression function than i.i.d. data,



while negatively correlated data contain ‘more information’ about the regression function than i.i.d. data. Using statistical information ideas, a simple measure of ‘information in the data’, on the scale of sample size, is provided by the ratio

$$n^* = \frac{\sigma^2}{\text{Var}(\bar{Y})}, \quad (13)$$

where

$$\text{Var}(\bar{Y}) = \frac{\sigma^2}{n} + 2\frac{1}{n} \sum_{k=1}^{n-1} \left(1 - \frac{k}{n}\right) \gamma(k). \quad (14)$$

The ratio of  $n^*$  to  $n$  gives a version of the *ESS*, i.e. the *Effective Sample Size*, which properly reflects the type and the magnitude of the correlation structure to give correct simultaneous inference:

$$ESS^*(i, h) = \frac{n^* \sum_{i=1}^n K_h(i - i_0)}{n K_h(0)}. \quad (15)$$

For independent data,  $n^* = n$ , so  $ESS^* = ESS$  from conventional *SiZer*. But for correlated data,  $n^*$  is smaller or larger than  $n$  depending on the type of correlation, i.e. observations have less effect if they are positively correlated (since they are ‘less informative’ than i.i.d. observations) and more effect if they are negatively correlated (since they are ‘more informative’ than i.i.d. observations).

The computation of  $m$  in (8) and  $q$  in (9) remains unchanged.

When the data are positively autocorrelated, the number of independent problems is larger than in the i.i.d. case, so the resulting confidence intervals are longer. Longer intervals are more likely to contain the value 0, so fewer features are flagged as statistically significant. On the other hand, for negatively correlated data (a situation that is rare in real data but is worth considering), *SiZer* for time series can detect those features that could be hidden in the family of smooths by the alternating pattern of data points above and below the regression function.

Finally the quantity  $n^*$  has to be estimated (since  $\sigma^2$  and  $\gamma$  are unknown). The estimate given by the sample autocovariance function can’t be expected to give good results since it is inconsistent. A simple approach to this problem is to divide the sample into  $p = \sqrt{n}$  groups (so as to have a reasonably large number of groups with a reasonably large number of observations), and to estimate the variance of  $\bar{Y}$  by

$$\widehat{Var}(\bar{Y}) = \frac{1}{p} \frac{1}{p-1} \sum_{j=1}^p (\bar{Y}_j - \bar{\bar{Y}})^2 \quad (16)$$

where  $\bar{Y}_j$  is the mean in the group  $j$  and where

$$\bar{\bar{Y}} = \frac{1}{p} \sum_{j=1}^p \bar{Y}_j. \quad (17)$$

#### 4 Simulated and real data examples

In this section, we will illustrate some examples of *SiZer* for time series, chosen from many simulated and real data sets that we have explored. We will start with some examples from the simulation study to illustrate the development of the methodology and graphical presentation of the results.

The simulation study has been carried out by considering different combinations of trend, error structure and noise level. A series of different data sets have been simulated from the model

$$Y_i = m(i) + \sigma e_i \quad i = 1, \dots, n \quad (18)$$

for different choices of  $m$ ,  $\{e_i\}$  and  $\sigma$ , and for different values of  $n$ .

The trends that were considered are characterized by a different number of ‘peaks and valleys’, differently located and with different amplitudes. For each chosen trend we have considered i.i.d. errors and positively and negatively correlated errors. In particular, i.i.d. errors have been simulated from  $N(0, 1)$  while correlated errors have been simulated from an autoregressive process of order 2 with high correlation at lag one (for the examples of the simulation study shown in this paper, the value of the autocorrelation function at lag one for positively correlated errors is 0.97 while for negatively correlated errors the value of the autocorrelation function at lag one is -0.89). Moreover, in each case, three different noise levels were chosen to represent ‘low’, ‘medium’ and ‘high’ variability (in the examples shown below, these values were respectively 1, 20 and 50).

The real data sets shown here are the Deaths data set, i.e. the monthly number of accidental deaths in US from 1973 to 1978 (thousands) and the Chocolate data set, that is considered in the intro-

duction of this paper. As the Chocolate data set, the Deaths data set comes with the software companion to Brockwell and Davis [1996].

#### Example 4.1

In the simulation study, several interesting cases among the models considered had the trend:

$$m(i) = i + 10\sin\left(\frac{i}{40}2\pi\right) \quad i = 1, \dots, n. \quad (19)$$

Given the model in (18) and the trend in (19), let us consider a simulated time series, that we will indicate by sm1, where errors are positively correlated and a ‘medium’ level of variability is chosen. An important issue is how many *SiZer* plots, should be shown. There will be a different plot for each pilot bandwidth  $h_p$ , showing different trade-offs between trend and serial correlation. If  $h_p$  takes all 11 values of the bandwidth  $h$  then the complete series of *SiZer* plots will be 11 in number. The family plot and the complete series of 11 *SiZer* plots for the simulated time series sm1 is given in Figure 3.

(put Figure 3 approx. here)

(caption: Example 4.1: the family plot and the complete series of 11 *SiZer* plots, indexed by  $h_p$ , for the simulated time series sm1)

The family plot shows a family of smooths with a strong upward trend.

The first *SiZer* plot (titled  $hp(1)$ ) has very small  $h_p$ . Thus, the error structure used to construct this plot is estimated from the residuals of the closest curve of the family of smooths to the data points, which is essentially i.i.d.. From this plot we observe some significant structure at the finest levels of resolution, i.e. for the smallest values of the bandwidth. As the level of resolution decreases, less and less structure is significant until the curves are significantly increasing everywhere.

In the second *SiZer* plot (titled  $hp(2)$ ), which represents a situation of slightly correlated errors, less structure than in the previous plot, as expected, appears as significant. Moreover, at the finest level of resolution the data are too sparse, in terms of  $ESS^*$ , to draw any conclusion, i.e. the bottom line is now shaded gray, as discussed in Section 2.2.

As we move from the third to the last *SiZer* plot (i.e. from the plot titled  $hp(3)$  to the plot titled  $hp(11)$ ), an increasing amount of correlation appears in the error component, so that only the linear upward trend seems to be significant at every level of resolution, where conclusions can be drawn. Also, at the finest levels of resolution there is less perceived useful information in the data, which means more data sparsity, so more bottom lines of the *SiZer* plots are shaded gray.

Statistically significant structure in a curve can be hidden by a strong linear component. When this happens, as in this example, it can be useful to detrend the time series before using *SiZer*. Figure 4 shows the complete series of *SiZer* plots, for the same data set after linear detrending (indicated by sm1d).

(put Figure 4 approx. here)

(caption: Example 4.1: the family plot and the complete series of 11 *SiZer* plots, indexed by  $h_p$ , for the simulated and linearly detrended time series sm1d)

Now, many features appear in the family of smooths which is no longer closely following an increasing line. In particular, in the first *SiZer* plot, many peaks and valleys are significant for most of the levels of resolution. As the errors include increasing correlation, these features are significant but with a less precise location and for a smaller number of levels of resolution. From the seventh plot to the last, all of the structure in the data is explained by the error component and no significant structure is highlighted in the trend at any level of resolution.

This example clearly shows that the complete series of *SiZer* plots is too long. The simultaneous view of all 11 *SiZer* plots is hard to comprehend and the information contained in several such plots is often redundant. This motivates the choice of a subset of *SiZer* plots.

We develop a method for effective choice of a subset of these, which usually gives good representatives of the major different correlation structures. We found 4 plots usually conveyed the needed information. Our choice among the 11 plots is intended to reflect ‘a wide array of trade-offs between trend and dependence’. A simple numerical measure of this trade-off is the *IR* (Indicator of the Residual component), which takes values from 0 to 1:

$$IR(h_p) = \frac{\sum_{i=1}^n e_{h_p,i}^2}{\max_{h_p} \sum_{i=1}^n e_{h_p,i}^2}. \quad (20)$$

When the pilot bandwidth  $h_p$  is large, the ‘dependence component of the data’ appears strongly in the residuals, i.e. is viewed as ‘noise’, and  $IR$  is close to 1. On the other hand, when the pilot bandwidth  $h_p$  is small, the ‘dependence component of the data’ appears strongly in the pilot smooth, i.e. is viewed as ‘trend’, and the  $IR$  is close to 0. Intermediate values of  $IR$  reflect intermediate trade-offs.

A good reflection of the range of trade-offs came essentially from using the *SiZer* plots chosen (from the set of 11) for which  $IR$  is closest to 0%, 25%, 50% and 75%. This implies that the first plot to be chosen should always be the first of the 11 plots. However, since the error structure used to construct this plot is estimated from the residuals given by a smooth for which the degree of overfitting may be too high (mostly when data are positively correlated), a second option is considered. When this is the case, the second of the 11 plots is chosen as the closest to 0%, i.e. as ‘more representative of the independent case’. And if this plot corresponds to a value of  $IR$  which is also the closest to 25%, the 3 remaining plots will automatically be chosen by considering 25%, 50% and 75% of the difference between the values of  $IR$  for the 11th and the 2nd plots of the complete series.

When only 4 of the *SiZer* plots are displayed, it is useful to stay in touch with the trend-dependence trade-off that each represents. This is done by several graphical devices. First, all 11 values of  $IR$  that are considered are displayed as a bar graph, and the chosen 4 are highlighted. Second, for each chosen trade-off, we add plots showing the pilot smooth with bandwidth  $h_p$ , and the residuals from that smooth. The pilot smooths show which component of the data are viewed as trend, in that particular trade-off. The residuals give a visual impression of the component of the data used in the covariance estimate.

This is the visualization chosen for *SiZer* for time series. Figure 5 shows this graphical device for `smd1`.

(put Figure 5 approx. here)

(caption: Example 4.1: *SiZer* for time series for the simulated and linearly detrended time series `sm1d`)

The data are shown in the first plot above on the left (the continuous line shows  $m(i)$ , the deterministic part of the simulated time series), while the next graphic on the right is the family plot. Further right, is the bar diagram previously discussed.

The second and the third series of plots represent, respectively, the smooths and the residuals.

Note that the chosen *SiZer* plots give a representative sample of the complete series of *SiZer* plots shown in Figure 4. No relevant information is lost and the graphical representations of the smooths and the residuals associated to each *SiZer* give useful insights about each trend-dependence trade off. This demonstrates the power of the *SiZer* method for investigating trends in time series: one graphical presentation displays all of the relevant information about statistical significance.

#### Example 4.2

Figure 6 shows *SiZer* for time series for the Deaths data set, after deseasonalizing and linear detrending.

(put Figure 6 approx. here)

(caption: Example 4.2: *SiZer* for time series for the Deaths data set)

Assuming ‘i.i.d.’ errors (first *SiZer* plot on the left) only some features of the family of smooths are significant at the finest levels of resolution. The strongest feature is the minimum around  $i = 36$  for intermediate values of the bandwidth. But for smaller bandwidths, there is a significant increase near  $i = 20$ . For the largest value of the bandwidth the curve is neither significantly increasing nor decreasing.

For slightly correlated errors (second *SiZer* plot on the left), the major minimum again turns out to be significant but for a smaller range of bandwidths than in the previous plot (for the smallest values of the bandwidth no conclusion can be drawn and for the largest values no feature is significant). When the correlation increases, the number of values of the bandwidth for which no conclusion can be drawn increases while the minimum previously highlighted is still significant but only for a few values of the bandwidth and with a much less precise location. For a higher degree of correlation, no feature turns out to be significant, as we can see from the graphical presentation of the associated residuals and smooth. For such correlation structure every feature can be explained by the error component, thus not resulting in significant trend.

In this example, according to *SiZer* for time series, the only feature that appears to be significant at most of the levels of resolution, if we assume that errors are not strongly correlated, is the ‘valley’ around the third year of observation. For larger error correlation there is no significant trend.

### Example 4.3

Figure 7 shows *SiZer* for time series for the Chocolate data set. As in the previous example, this data set has been deseasonalized and linearly detrended. But here the strong upward trend, driven by increasing chocolate production due to increasing population size, seems to hide much more structure than in the previous case, as we can see once we detrend the data, especially when the errors are assumed to be ‘i.i.d.’. Nevertheless, *SiZer* for time series still gives useful insights about the trends under study, using the same 4 choices of *SiZer* plots from the complete series of 11.

(put Figure 7 approx. here)

(caption: Example 4.3: *SiZer* for time series for the Chocolate data set)

For ‘i.i.d.’ errors (first *SiZer* plot on the left) many ‘peaks and valleys’ are significant at the finest levels of resolution, i.e. for the smallest values of the bandwidth. As the level of resolution decreases, fewer and fewer features appear to be significant.

For slightly correlated errors, no conclusion can be drawn at the finest levels of resolution while for intermediate values of the bandwidth the major minimum around  $i = 250$  (which corresponds to year 78) is the only significant feature. For the highest value of the bandwidth the curve is neither significantly increasing nor decreasing.

As the correlation increases no significant structure is highlighted with the only exception being some intermediate values of the bandwidth where the curves are significantly increasing on the right end of the time domain.

For the strongly correlated error assumption, no significant structure appears at any level of resolution.

## 5 References

1. Bowman, A.W. and Azzalini, A. (1997), *Applied smoothing techniques for data analysis, the kernel approach with S-plus illustrations* Oxford Statistical Science Series 18, Oxford Science Publications, Oxford.
2. Brockwell, P.J. and Davis, R.A. (1991), *Time series: theory and methods*, 2nd Edition, Springer-Verlag, New York.

3. Brockwell, P.J. and Davis, R.A. (1996), *Introduction to time series and forecasting*, Springer-Verlag, New York.
4. Chaudhuri, P. and Marron, J.S. (1999), SiZer for exploration of structures in curves, *Journal of the American Statistical Association*, 94, 807-823.
5. Fan, J. and Gijbels, I. (1996), *Local polynomial modeling and its application*, Chapman and Hall, London.
6. Green, P.J. and Silverman, B.W. (1994), *Nonparametric regression and generalized linear models*, Chapman and Hall, London.
7. Hart, J.D. (1997), *Nonparametric smoothing and lack of fit tests*, Springer, New York.
8. Marron, J.S. and Chaudhuri, P. (1998a), Significance of features via SiZer, *Statistical Modelling, Proceedings of 13th International Workshop on Statistical Modelling*, Brian Marx and Herwig Friedl, Eds., 65-75.
9. Marron, J.S. and Chung, S.S. (2001), Presentation of smoothers: the family approach, *Computational Statistics*, to appear.
10. Simonoff, J.S. (1996), *Smoothing methods in statistics*, Springer-Verlag, New York.
11. Schmidt, G., Mattern, R. and Schüler, F. (1981), *Biomechanical investigation to determine physical and traumatological differentiation criteria for the maximum load capacity of head and vertebral column with and without helmet under effects of impact*, EEC Research program on Biomechanics of Impacts, Final Report Phase III, Project 65, Institut für Rechtsmedizin, Universität Heidelberg, West Germany.
12. Wand, M.P. and Jones, M.C. (1995), *Kernel smoothing*, Chapman and Hall, London.