

Dependent SiZer: goodness of fit tests for time series models

CHEOLWOO PARK

Statistical and Applied Mathematical Sciences Institute,
19 T. W. Alexander Drive, P.O. Box 14006,
Research Triangle Park, NC 27709-4006
cwpark@samsi.info

J. S. MARRON

Department of Statistics and Operations Research,
University of North Carolina,
Chapel Hill, NC 27599-3260
marron@email.unc.edu

AND VITALIANA RONDONOTTI

Monetary, Financial Institutions and Markets Statistics Division,
European Central Bank,
Frankfurt am Main, Germany
vitaliana.rondonotti@ecb.int

July 1, 2004

SUMMARY

In this paper, we extend SiZer (SIGNificant ZERo crossing of the derivatives) to dependent data for the purpose of goodness of fit tests for time series models. Dependent SiZer compares the observed data with a specific null model being tested by adjusting the statistical inference using an assumed autocovariance function. This new approach uses a SiZer type visualization to flag statistically significant differences between the data and a given null model. The power of this approach is demonstrated through some examples of time series of Internet traffic data. It is seen that such time series can have even more

burstiness than is predicted by the popular, long range dependent, Fractional Gaussian Noise model.

Key words: Autocovariance function; Dependent SiZer; Fractional Gaussian Noise; Internet traffic data; Goodness of fit test; SiZer; Time series.

1 Introduction

A goodness of fit test is a statistical technique to validate a null hypothesis without specifying an alternative hypothesis. It is useful to assess models being considered as possible candidates to explain the data well. In time series modeling, there have been several different goodness of fit tests, for example see Beran (1994), or Chen, Härdle, and Li (2003). In this paper, a new goodness of fit test method combined with visual insight is proposed by extending SiZer (SIGNificant ZERo crossing of the derivatives), originally proposed by Chaudhuri and Marron (1999). We illustrate the usefulness of our method in the challenging context of time series that arise in the statistical analysis of Internet traffic.

It is well known, see for example Paxson and Floyd (1995), and Leland, Taqqu, Willinger, and Wilson (1994), that many Internet time series are notoriously resistant to classical time series analysis methods. For example, they frequently tend to be long range dependent, so that conventional ARMA modeling can be quite inappropriate. In this type of context, Fractional Gaussian Noise (FGN), a natural, long range dependent, Gaussian process has been an appealing, and widely used model, see for example Taqqu (2003).

To allow efficient sharing of resources (i.e. simultaneous multiple transfers), data is divided into small units called packets, for transmission over the internet. In this paper, we study time series of packet counts (the numbers of packets arriving in consecutive 1 millisecond intervals). In other words, our raw data is a time series $Y = \{Y_i, i = 1, 2, \dots, n\}$ of binned packet arrival times, with a bin width of 1 millisecond.

Figure 1 about here.

Figure 1 displays a full packet count time series of Internet traffic coming into the University of North Carolina, Chapel Hill (UNC) from outside. They were measured at the main internet link of UNC on April 13, Saturday, from 1 p.m. to about 3 p.m., 2002 (Sat1300). The time series plot in Figure 1 shows a thick spike and two long thin spikes in the middle,

and a small spike at the right end. From this plot, however, it is not easy to see which features are meaningful.

SiZer analysis is a visualization method that enables statistical inference to discover meaningful structure within the data, while doing exploratory analysis using statistical smoothing methods. It brings clear and immediate insight into a central scientific issue in exploratory data analysis: Which features observed in a smooth of data are “really there”? This main question is critical in real data analysis, because discovery of a new feature can often lead to important new scientific insight, while on the other hand scientific effort can be wasted trying to find explanations for spurious features.

Figure 2 about here.

Figure 2 (a) shows the SiZer plot of the Sat1300 time series. The thin blue curves in the top panel display the family of smooths. These are kernel regression estimates, see e.g. Wand and Jones (1995), based on the data, (i, Y_i) for $i = 1, \dots, 7345555$. Some impression of the data is given by displaying a subsample as green dots. A random sample of 1000 is displayed to avoid overplotting. The Y coordinates of the green dots are integer valued, because they are counts. There are several blue curves corresponding to different levels of smoothing, that is bandwidths. Some of these curves are oversmoothed, and some are undersmoothed. The goal of SiZer is to understand which of the features in the smooths can be distinguished from the background white noise. In other words, to determine which features of the blue curves are important and which are just spurious sampling artifacts is a challenging problem.

The lower panel of Figure 2 (a) displays the SiZer map of the data. The horizontal locations in the SiZer map are the same as in the top panel, and the vertical locations correspond to the logarithm of bandwidths of the family of smooths shown as blue curves in the top panel. Each pixel shows a color that 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 bandwidth corresponding to that row. When the slope of the blue curve is significantly positive (negative), the pixel is colored blue (red, respectively). When the slope is not significant, which means the sampling noise is dominant, the pixel has the intermediate color of purple. There is less purple in this map than is typical, because the unusually large sample size of $n = 7345555$, results in a very powerful hypothesis test. There is a fourth SiZer color, that does not appear in Figure 2 (a), which is gray, used to show pixel locations where the data are too sparse for reasonable statistical inference. See Chaudhuri and Marron (1999)

for a more detailed introduction to SiZer.

In Figure 2 (a), the family of smooths indicates that the large spike in the middle dominates the other features. From the SiZer map, however, there are a lot of blue and red colors at several time and scale locations other than where the huge spike occurs. It clearly shows that the huge spike is significantly different from white noise, but so are all of the other wiggles (large and small) visible in the family of smooths. All these cases of burstiness are flagged as significant because of the very large sample size, and because SiZer is trying to distinguish the data from white noise, but the data have an intrinsic burstiness caused by natural dependence in the time series.

In addition to long range dependence, Internet traffic data have been known to exhibit the self-similar property, see e.g. Leland, Taqqu, Willinger, and Wilson (1994). Fractional Gaussian Noise (FGN) is a fundamental time series model which exhibits both long-range dependence and self-similarity, see Mandelbrot (1988) for example. To understand when such models are appropriate, it is useful to test whether a given time series is consistent with FGN or not. Figure 2 (b) will be discussed later.

Figure 3 about here.

Motivation for the development of dependent SiZer comes from comparing the analysis in Figure 2 (a) with a SiZer analysis of simulated FGN data (with parameters defined later), as shown in Figure 3 (a). While the two families of smooths exhibit very different behavior (a single very big burst for the Sat 1300 data, compared with more numerous and homogeneous bursts in the simulated data), the two SiZer maps look at least qualitatively similar, with very many red and blue significant regions. Thus conventional SiZer is not very effective at differentiating between these data sets (with apparently different burstiness properties). Much better differentiation will come from adjusting the SiZer inference, so that burstiness at the level inherent to FGN, as shown (for the same simulated data set) in Figure 3 (b), is the null hypothesis, and thus shows up as all purple. When such a method is applied to the Sat1300 time series, as in Figure 2 (b), it becomes clear that most of the burstiness apparent in Figure 2 (a) is of a magnitude expected under an FGN model, but the large central spike is a much different type of burst.

Figure 3 (b) and Figure 2 (b), are illustrations of dependent SiZer, which enables us to more specifically compare the data with a given model. This can be done by adjusting the statistical inference, using the autocovariance function of an assumed model, see section 2.2

for the details. Therefore, dependent SiZer provides not only a goodness of fit test of an assumed model but also visual insight as to how the data differ from an assumed model.

Figure 2 (b) shows the dependent SiZer plot (using adjusted statistical inference) of the Sat1300 time series. The coordinate axes and the colors are the same as those of the original SiZer. The only difference is that dependent SiZer compares the data against a null hypothesis of FGN with particular parameters, $H = 0.9$ and variance 42.416. The estimation of these parameters is described in Section 4.1. The SiZer map in the lower panel confirms that the spike in the middle can not be explained by FGN, as expected, because there is a significant increasing trend (flagged as blue) and a decreasing trend (flagged as red). However, in the other regions, there is no significant evidence that they are different from FGN because the SiZer map shows all purple. The input time series to the SiZer analysis in Figure 2 (b) is no longer the full packet count series of $n = 7345555$, but instead has been summarized by binning, to a series of length 400, for fast calculation of the dependent SiZer inference. The green dots in Figure 2 (b) look much different from those in part (a), because they show this summarized, i.e. binned, time series. No subsampling is done here because there are only 400 green dots.

Figure 4 about here.

To further explore this statistically significant spike, we zoom in for a closer look, by restricting attention to the data between the two vertical red lines in the top panel of Figure 2 (b). Figure 4 (a) shows the dependent SiZer within the spike (i.e. between the red lines), which is about 6 minutes long. The family of smooths shows a valley in the middle and the SiZer map indicates that this valley is not typical behavior of FGN with the same parameters above, by the red and blue regions. A further zoomed dependent SiZer, which corresponds to the two vertical red lines in the top panel of Figure 4 (a), is displayed in Figure 4 (b). This extracted subtrace is around 43 seconds long and its dependent SiZer plot surprisingly shows periodic behavior, with a period of about 3 seconds. This behavior is very inconsistent with a FGN model and the SiZer map in the lower panel confirms this by flagging the bumps in the top panel as significant. A likely explanation of this is an intense Internet Protocol (IP) port scan, where a hacker has attempted to find weak points in the system by sending query packets to all of the ports of all possible IP addresses in the UNC domain.

This example shows that dependent SiZer is a very useful tool, not only as a goodness of fit test that studies whether the given data are consistent with the model being tested, but

also to highlight unusual behavior of the data. This is done through the SiZer combination of visual insight and statistical inference.

Section 2 describes SiZer and dependent SiZer. Dependent SiZer is tested and compared with SiZer for three more simulated examples in Section 3. Section 4 describes parameter estimation for dependent SiZer and shows some more real data examples. Discussion of dependence fragility of dependent SiZer is described in Section 4 as well.

2 SiZer Methodology

Some details of conventional SiZer are reviewed in Section 2.1. The extension of these ideas to dependent SiZer is described in Section 2.2.

2.1 SiZer

SiZer is based on scale space idea from computer vision, see Lindeberg (1994) and ter Haar Romeny (2001) for example. Scale space is a family of kernel smooths indexed by the scale, which is a smoothing parameter h , that is the bandwidth. SiZer considers a very wide range of bandwidths to avoid the classical problem of bandwidth selection. The target of the SiZer analysis is shifted from the “true underlying curve” to “smoothed versions of the underlying curve”, with the idea that this contains all the information that is available in the data, for a given bandwidth, i. e. level of resolution of the data. SiZer extends the usefulness of a family of smooths plot by combining the SiZer map, which displays results of statistical inference. The SiZer map visually displays the significance of features over both location x and scale h . Multiple comparison tests based on confidence intervals for the derivatives of the underlying curve are involved in flagging significant features, using colors as described in Section 1.

Let us consider a regression problem with a fixed design setting. Given the data (t_i, Y_i) where $t_i = i/n$, for $i = 1, \dots, n$ and n is the sample size, a regression problem is described as

$$Y_i = m(t_i) + \sigma\epsilon_i, \tag{2.1}$$

where m is a regression function, $\sigma > 0$, and the ϵ_i 's are identically and independently distributed with $E(\epsilon_i) = 0$ and $Var(\epsilon_i) = 1$ for all i .

SiZer uses the local linear fitting method, see e.g. Fan and Gijbels (1996), to obtain a family of kernel estimates and derivatives of a regression function. Precisely, at a particular

point t_0 , they are obtained by minimizing

$$\sum_{i=1}^n \{Y_i - (\beta_0 + \beta_1(t_0 - t_i))\}^2 K_h(t_0 - t_i) \quad (2.2)$$

where $\beta = (\beta_0, \beta_1)'$ and $K_h(\cdot) = K(\cdot/h)/h$. K is called a kernel function which is taken here to be a Gaussian probability density function. By Taylor expansion, it is easy to show that $\beta_0 \approx m(t_0)$, and $\beta_1 \approx m'(t_0)$, thus the solution of (2.2) provides estimates of the regression function m and its first derivative m' for different bandwidths h . From this solution, we can construct a family of smooths and confidence intervals for the SiZer map.

The details underlying the SiZer statistical inference can be found in Chaudhuri and Marron (1999). The essential idea is that SiZer reports the results of a large number of simultaneous hypothesis tests. Of course it is very important to properly account for the simultaneous nature of these tests (to avoid the display of spurious structure in about 5% of the SiZer map). Because many of the tests are strongly correlated with each other, a Bonferroni approach is too crude. A more workable approximation is the “independent blocks” approach, described in Section 3 of Chaudhuri and Marron (1999), which is used to provide simultaneous inference here as well. Also, to achieve reasonable computational speed, fast binned implementation of the smoothers and the corresponding hypothesis tests, are used as discussed in Section 3.1 of Chaudhuri and Marron (1999). This is very useful in the context of internet traffic, where the data sets can be very large.

2.2 Dependent SiZer

The SiZer map of Figure 2 (a) shows a lot of significant trends. However, since the original SiZer assumes the data have independent errors, those features may be simple artifact of the natural dependence inherent to the time series. This motivates the need to properly incorporate dependence into the SiZer analysis. Rondonotti and Marron (2004) provided formulas for this, and proposed a family of data driven estimates of the level of dependence. In this paper, we use the same formulas for the SiZer inference, but use the FGN structure to provide simpler estimates of the level of dependence.

A time series setting can be viewed as a regression setting as in (2.1) with $t_i = i$. But a critical difference is that now the ϵ_i 's are no longer independent, and thus

$$Cov(\epsilon_i, \epsilon_j) = \gamma(|i - j|), \quad (2.3)$$

for all i and j where γ is an autocovariance function. In this case, the variance of the local linear estimator at i_0 is given by

$$V(\hat{\beta}) = (X'WX)^{-1}(X'\Sigma X)(X'WX)^{-1}$$

where $\hat{\beta}$ is an estimator of $\beta = (\beta_0, \beta_1)'$ in (2.2), $W = \text{diag}\{K_h(i - i_0)\}$, and the design matrix X at i_0 is

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

The kernel weighted covariance matrix of the data, Σ , has its elements σ_{ij} , which is defined by

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

SiZer for time series estimates γ , using the sample autocovariance function of the observed residuals from a “pilot smooth”. Therefore, it attempts to find meaningful features, while properly accounting for dependence.

The dependent SiZer proposed in this paper uses a given autocovariance function from an assumed model, instead of estimating it from the observed data. Our goal is to explore differences between the given data and the assumed model. In particular, a goodness of test is conducted against an assumed model by considering its true dependence structure.

For example, the autocovariance function of FGN is given by

$$\gamma(k) = \sigma^2((k + 1)^{2H} + (k - 1)^{2H} - 2k^{2H})/2, \quad k \geq 0, \quad (2.4)$$

where H is the Hurst parameter. When $H = 1/2$, $\gamma(k) = 0$ for $k \neq 0$ and ϵ_i becomes white noise. For $1/2 < H < 1$, as $k \rightarrow \infty$, γ has long range dependence. See Mandelbrot and Van Ness (1969) for example.

The details of the multiple comparison test procedure in dependent SiZer can be found in Rondonotti and Marron (2004).

3 Simulated examples

This section provides three simulated examples and comparisons between original SiZer and dependent SiZer. For this simulation study, n pairs (i, Y_i) are generated according to the no signal prescription

$$Y_i = \sigma \epsilon_i, \quad i = 1, \dots, n.$$

Because there is no signal, correct SiZer inference will show all purple in the SiZer map.

Figure 5 about here.

The first example is the independent case and the ϵ_i 's are drawn from a standard normal distribution with $\sigma = 1$. The sample size is chosen to be $n = 400$. Figure 5 displays the SiZer plot (shown in (a)) and the dependent SiZer plot (shown in (b)) of this time series, where the white noise covariance structure is specified. The observation that both SiZer maps show purple is consistent with the data being white noise. This example shows that dependent SiZer works equally to conventional SiZer as we expected when the errors are white noise (and that dependence model is used).

However, a very careful look at the SiZer map in the lower panel reveals some small unexpected colored regions. Such features have been observed in a number of other cases as well, including the following examples. Although such small features can be discounted as unimportant, it is still desirable to eliminate this. Hannig and Marron (2003) has proposed a new method using an advanced distribution theory to reduce these unexpected features in the SiZer map. In this paper, we use the original SiZer, and we leave extending this method to dependent SiZer for our future work.

Figure 6 about here.

The second example studies performance in the case of short range dependence. In this case, the ϵ_i 's are drawn from MA(1). Again, the sample size is chosen to be $n = 400$. The MA(1) dependency creates a number of significant features in the conventional SiZer map shown in Figure 6 (a) On the other hand, the dependent SiZer map of this example, which is shown in the lower panel of Figure 6 (b), shows purple as we expected because we test the data under the null hypothesis of MA(1). This clearly confirms that dependent SiZer can successfully conduct a goodness of fit test when the data are generated from MA(1).

The last example studies stronger dependence, there the ϵ_i 's are drawn from the FGN with $H = 0.9$ and $\sigma^2 = 20$. We generate this time series with sample size $n = 7,200,000$. Because of the heavy dependency, the SiZer map of Figure 3 (a) shows much more burstiness compared to that of MA(1) shown in the lower panel of Figure 6 (a). On the other hand, as discussed in Section 1, the dependent SiZer map in the lower panel of Figure 3 (b) displays mostly purple because the assumed model is FGN. Again, some unexpected structure appears at the second level, but it does not strongly suggest the data are different from FGN with the given parameters.

From these simulations, we observe that dependent SiZer works well in the case of white noise, and it also successfully conducts goodness of fit tests for various types of dependent error.

4 Real data

In this section, we present more results from the analysis of 20 two-hour packet count time series collected in 2002 from the high-speed access link connecting the campus of the University of North Carolina at Chapel Hill to its Internet Service Provider (ISP). The time series were measured every 1 millisecond at different times each day for a week-long period. The model being tested in this section is FGN.

4.1 Estimation of parameters

Since dependent SiZer assumes a given autocovariance function of the model being tested, this requires the estimation of parameters involved with an autocovariance function of the assumed model. In the case of a FGN series, the Hurst parameter H and the variance σ^2 need to be estimated to obtain the autocovariance function in (2.4).

For the estimation of H , we tried three different methods, which are the Aggregated Variance method, see Beran (1994), the Local Whittle method, see Robinson (1995), and the Wavelet method, see Abry and Veitch (1998). Interesting comparisons of all these three methods on all 2002 UNC packet counts time series can be found in Hernández-Campos et al. (2004). In this paper, estimated H values of time series were around or over 0.9, and

sometimes over 1.0. For numerical stability, we use $H = 0.9$ in the examples in this paper.

Figure 7 about here.

To determine the variance parameter, we draw a plot of the sample variances versus the sample means of packet counts of all 20 traces. Figure 7 clearly shows a linear relationship between the mean and the variance. Some data sets depart from this line, which may be an indicator of nonstationary behavior. Different modifications of SiZer, which study local nonstationarity, can be found in Park, Godtlielsen, Taqqu, Stoev and Marron (2004), and Stoev, Taqqu, Park, Michailidis and Marron (2004). From this plot, we apply a robust linear fit and estimate the variance σ^2 using the mean of packet counts as the input.

4.2 Real data examples

To save space, only dependent SiZer plots of a few data sets are presented in this subsection, but complete results and other interesting analyses can be found at http://www-dirt.cs.unc.edu/net_lrd/DepSiZer/SiZER_View.html.

Figure 8 about here.

Large spikes (e.g. as shown in Figure 2) and deep valleys are typical types of nonstationary behavior of packet counts time series, that may not be explainable by FGN. Another nonstationary behavior is a big increasing or decreasing trend in a time series. Figure 8 (a) shows the dependent SiZer plot of the packet counts time series measured on April 9, Tuesday, 2002, from 8 a.m. to 10 a.m. (Tue0800). The family of smooths shows a big increasing trend and the blue colors at coarse scales in the SiZer map confirm that this increasing trend is not typical behavior of FGN with $H = 0.9$ and $\sigma^2 = 39.9159$. This increasing (decreasing) trend was frequently found for morning (and also night) time blocks, because of the increasing (decreasing) number of network users.

Among the full collection of data sets considered here, there were several whose bursts could all be explained by a FGN model from the dependent SiZer point of view. For example, the packet counts time series measured on April 13, Saturday, 2002, from 7:30 p.m. to 9:30 p.m. (Sat1930), is studied in Figure 8 (b). The dependent SiZer of this time series and the consistent purple color in the lower panel indicates that this is not significantly different from FGN with $H = 0.9$ and $\sigma^2 = 38.5853$. Thus, FGN is not a bad choice for modeling much of

the behavior observed in the 2002 packet counts time series.

Since dependent SiZer depends on the estimation of parameters, it is important to investigate the sensitivity of dependent SiZer to changing parameters. To stress the importance of this issue, we tried other values for H and σ^2 . Since H is an index of dependency in the data, changing H values certainly affects dependent SiZer. On the other hand, the choice of σ^2 generally has somewhat less impact on the dependent SiZer inference.

Figure 9 about here.

However, one example of different conclusions by a slight change of σ^2 is the packet counts time series of April 12, Friday, 2002 from 3 a.m. to 5 a.m. (Fri0300). Figure 9 displays two dependent SiZer plots with different variance parameters. The Hurst parameter H is given by 0.9 for both plots. In Figure 9 (a), the variance is given by the estimated value from Figure 7, which is 29.539, and, in Figure 9 (b), the variance is given by 20. In spite of a clear decreasing trend in the family of smooths, the SiZer map of Figure 9 (a) does not flag this trend as significant. On the other hand, the SiZer map of Figure 9 (b) shows red at coarse scales in the SiZer map, which means this trend is significantly different from FGN with $H = 0.9$ and $\sigma^2 = 20$. This suggests that dependent SiZer needs precise estimation of parameters to give an accurate goodness of fit test.

4.3 Dependence fragility of dependent SiZer

Figure 10 about here.

In this section, we investigate the robustness of the SiZer goodness of fit test, against misspecified dependence model. We use the same data from Figure 3, which is a simulated FGN with $H = 0.9$ and $\sigma^2 = 20$, but model the dependence structure using a ARMA fit to the data. Figure 10 shows a goodness of fit test result against a null ARMA(5,3) dependence structure (which was estimated using standard methods). Even though FGN with $H = 0.9$ is viewed as long range dependent and ARMA (5,3) is classically short range dependent, the SiZer map in the lower panel does not show anything significant. This phenomenon appears related to the *fragility* ideas of Gong et al. (2001), who pointed out that the notion of “heavy tailed” distributions can be *fragile* in the sense that several very different models (including those that are classically considered to be both heavy and light tailed) can fit in the tail of the distribution of real data.

5 Discussion

As noted by a reviewer there are several interesting connections between SiZer and wavelets (good elementary introduction to wavelets can be found in Ogden (1997)). The first connection is that the scale space generated by a data set has strong parallels to the wavelet multi-resolution analysis. The different window widths, i.e. scales in SiZer (appearing in rows of the map), correspond strongly to the frequencies of a wavelet analysis. The columns in the SiZer map are analogs of the location indices in the wavelet decomposition. Deeper comparison of this type can be found in Lindeberg (1994).

Another connection comes from combining a wavelet decomposition with SiZer style statistical inference, as done by Park, Godtlielsen, Taqqu, Stoev and Marron (2004). There SiZer targets non-stationarity in the wavelet coefficients, which provides a particularly sensitive way of finding non-stationarity in time series.

References

- ABRY, P. AND VEITCH, D (1998). Wavelet analysis of long range dependent traffic. *IEEE Transactions on Information Theory* **44**, 2–15.
- BERAN, J. (1994). *Statistics for long-memory processes*. Chapman & Hall, London.
- CHAUDHURI, P. AND MARRON, J. S. (1999). SiZer for exploration of structures in curves. *Journal of the American Statistical Association* **94**, 807–823.
- CHEN, X., HÄRDLE, W. AND LI, M. (2003). An empirical likelihood goodness-of-fit test for time series. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **65**, 663–678.
- FAN, J., AND GIJBELS, I. (1996). *Local Polynomial modeling and Its Applications*. Chapman & Hall, London.
- 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.
- HANNIG, J. AND MARRON, J. S. (2003). Advanced Distribution Theory for SiZer. *Submitted*.

- HERNÁNDEZ CAMPOS, F., LE, L., MARRON, J. S., PARK, C., PARK, J., PIPIRAS, V., SMITH F. D., SMITH, R. L., TROVERO, M., AND ZHU, Z. (2004). Long range dependence analysis of Internet traffic. In preparation.
- MANDELBROT, B. B. (1988). *The Fractal Geometry of Nature*. W. H. Freeman and Co., New York, New York.
- MANDELBROT, B. AND VAN NESS, J. (1968). Fractional Brownian Motions, Fractional Noises and Applications. *SIAM Review*, **10(4)**, 422–437.
- LELAND, W., TAQQU, M., WILLINGER, W., AND WILSON, D. (1994). On the Self-Similar Nature of Ethernet Traffic. *IEEE/ACM Transactions on Networking* **2**, 1–15.
- LINDBERG, T. (1994). *Scale Space Theory in Computer Vision*. Kluwer, Boston.
- OGDEN, R. T. (1997). *Essential Wavelets for Statistical Applications and Data Analysis*. Birkhauser, New York.
- PARK, C., GODTLIEBSEN, F., TAQQU, M., STOEV, S., AND MARRON, J. S. (2004). Visualization and inference based on wavelet coefficients, SiZer and SiNos. *Submitted to Journal of Computational and Graphical Statistics*.
- PAXSON, V. AND FLOYD, S. (1995). Wide Area traffic: the failure of Poisson modeling. *IEEE/ACM Transactions on Networking* **3**, 226–244.
- ROBINSON, P. M. (1995). Gaussian semiparametric estimation of long range dependence. *The Annals of Statistics*. **23**, 1630–1661.
- RONDONOTTI, V. AND MARRON, J. S. (2004). SiZer for time series: a new approach to the analysis of trends. *Submitted to Journal of Time Series Analysis*.
- STOEV, S., TAQQU, M., PARK, C., MICHAILIDIS, G. AND MARRON, J. S. (2004). LASS: a tool for the local analysis of self-similarity. *Submitted to Computational Statistics and Data Analysis*. **23**, 1630–1661.
- TAQQU (2003). Fractional Brownian motion and long-range dependence, *in* P. Doukhan, G. Oppenheim & M. S. Taqqu, eds, *textitTheory and Applications of Long-range Dependence*, Birkhäuser, pp. 5–38.
- TER HAAR ROMENY, B. M. (2001). *Front-End Vision and Multiscale Image Analysis*. Kluwer Academic Publishers, Dordrecht, the Netherlands.
- WAND, M. P. AND JONES, M. C. (1995). *Kernel Smoothing*. Chapman & Hall, London.

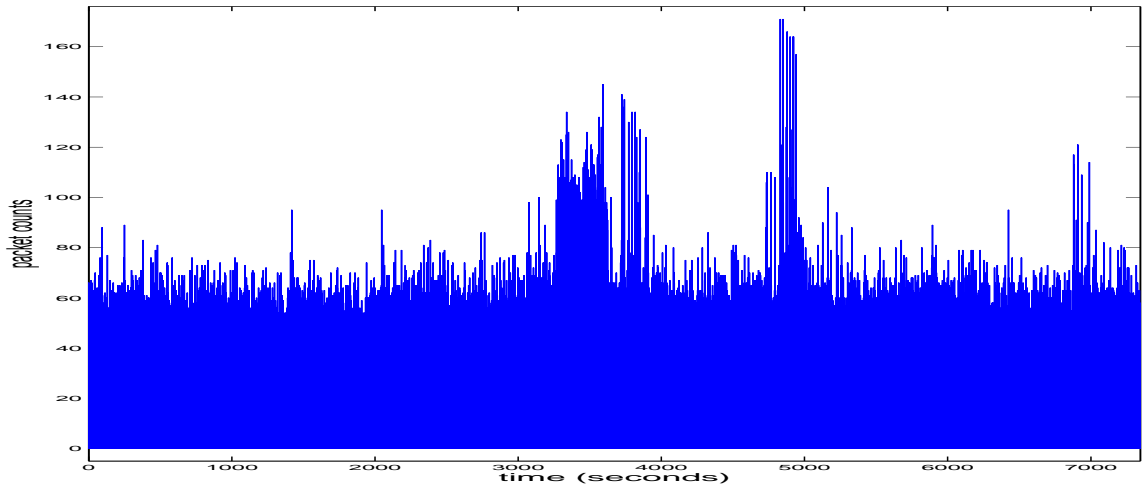


Figure 1: Time series plot of packet counts of the Sat1300 time series. The plot shows several spikes, but it is not clear which are meaningful features.

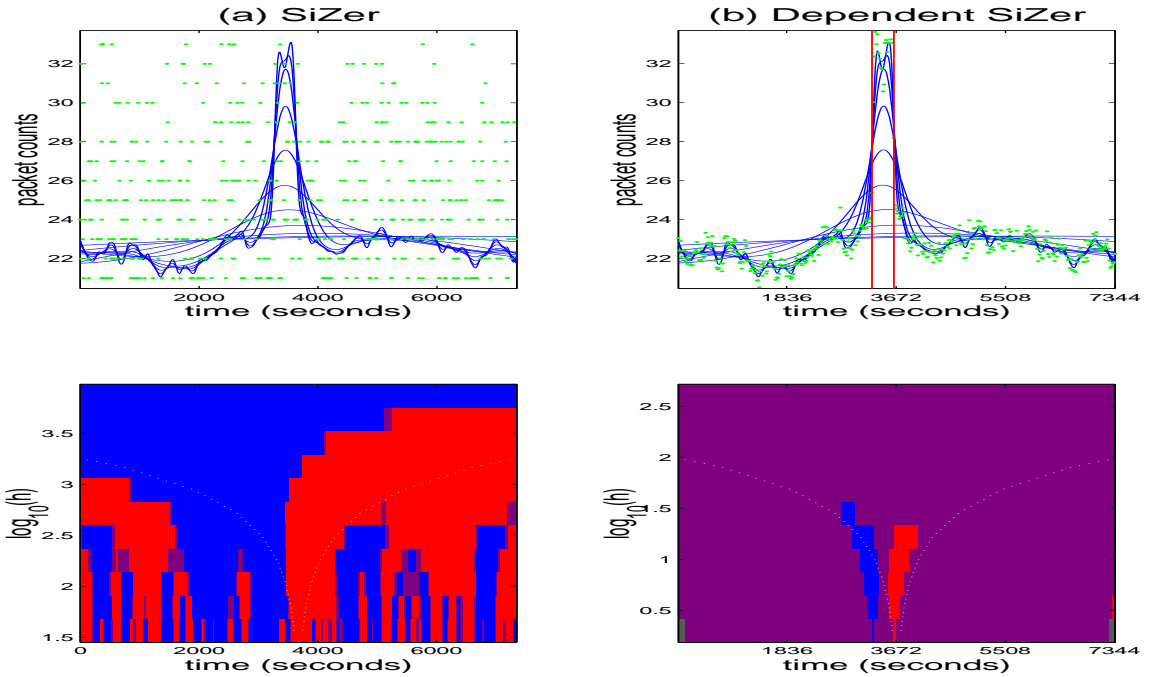


Figure 2: (a) Original SiZer plot of the Sat1300 time series. The top panel suggests that there is a huge spike in the middle. The lower panel shows many features distinguishable from the white noise including the huge spike in the middle. (b) Dependent SiZer plot of the Sat1300 time series. This series is tested against FGN. The spike in the middle is still significantly different from FGN with given parameters while other significant features in the original SiZer plot in (a) disappear, i.e. are consistent with FGN.

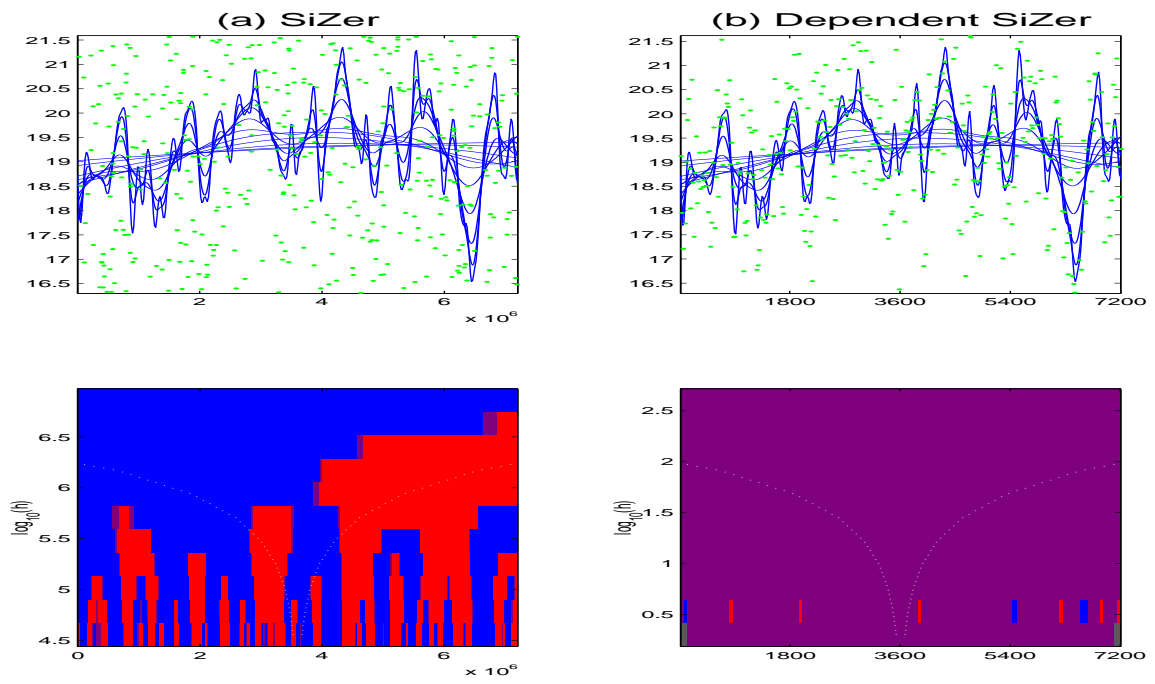


Figure 3: (a) Original SiZer of a simulated FGN. The SiZer map shows that many features at different scales are flagged as significant. (b) Dependent SiZer of the same simulated FGN. In the SiZer map, almost all pixels are purple, which is consistent with this model.

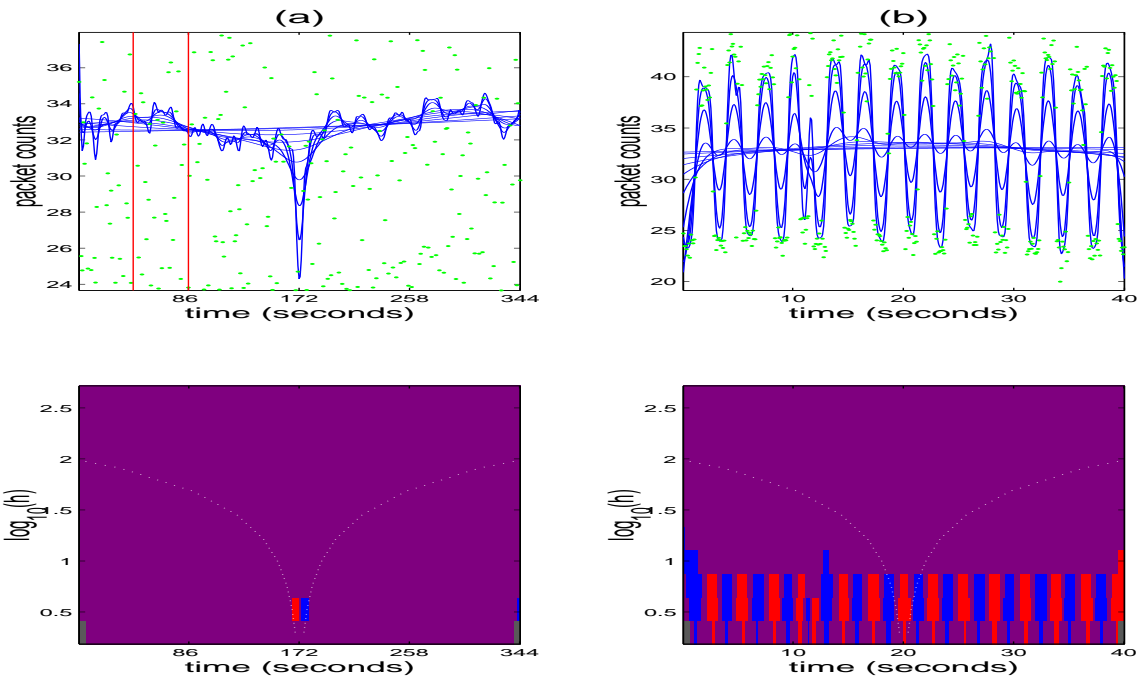


Figure 4: (a) A zoomed dependent SiZer (against FGN) plot of the Sat1300 data over a 6 minute range. The valley in the middle is significantly different from FGN. (b) A further zoomed dependent SiZer plot of the Sat1300 for 43 seconds. Surprisingly, there is a 3 second periodic behavior and which can not be explained by FGN.



Figure 5: (a) Original SiZer of simulated independent errors. Almost all pixels are purple in the lower panel, which is consistent with the given model. (b) Dependent SiZer of simulated independent errors. It shows that dependent SiZer gives the same result as that of conventional SiZer when the errors are white noise.

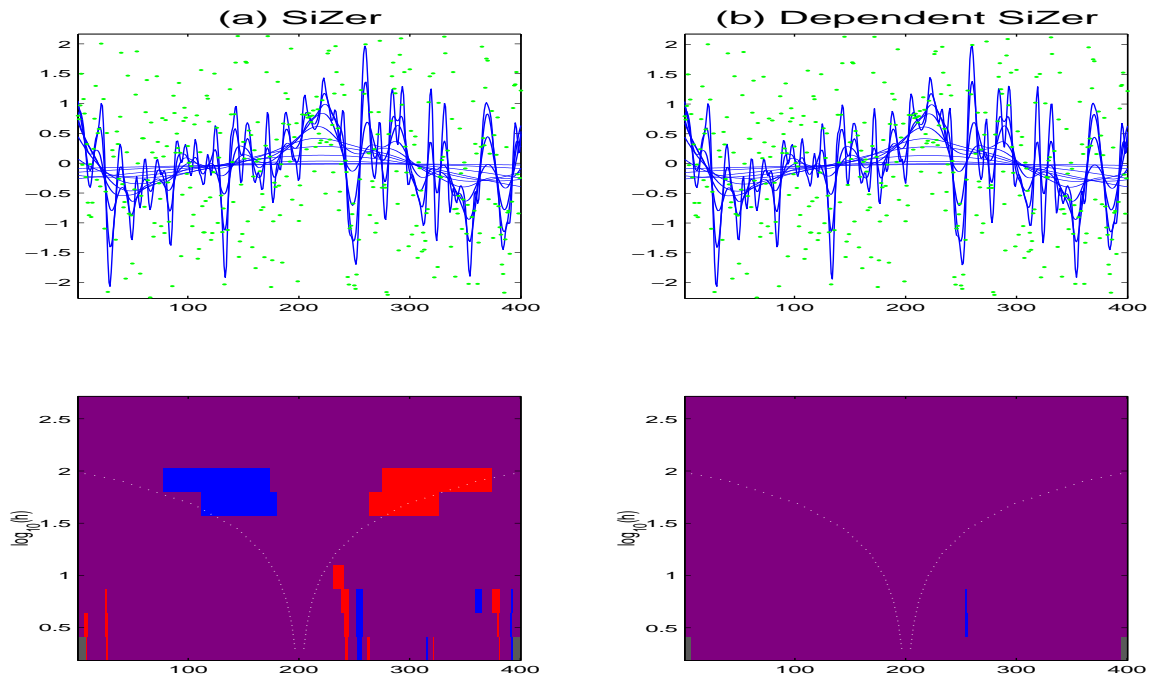


Figure 6: (a) Original SiZer of simulated MA(1) errors. The SiZer map in the lower panel flags some dependence artifacts as significant. (b) Dependent SiZer of simulated MA(1) errors. In the SiZer map, almost all pixels are purple, which is consistent with this model.

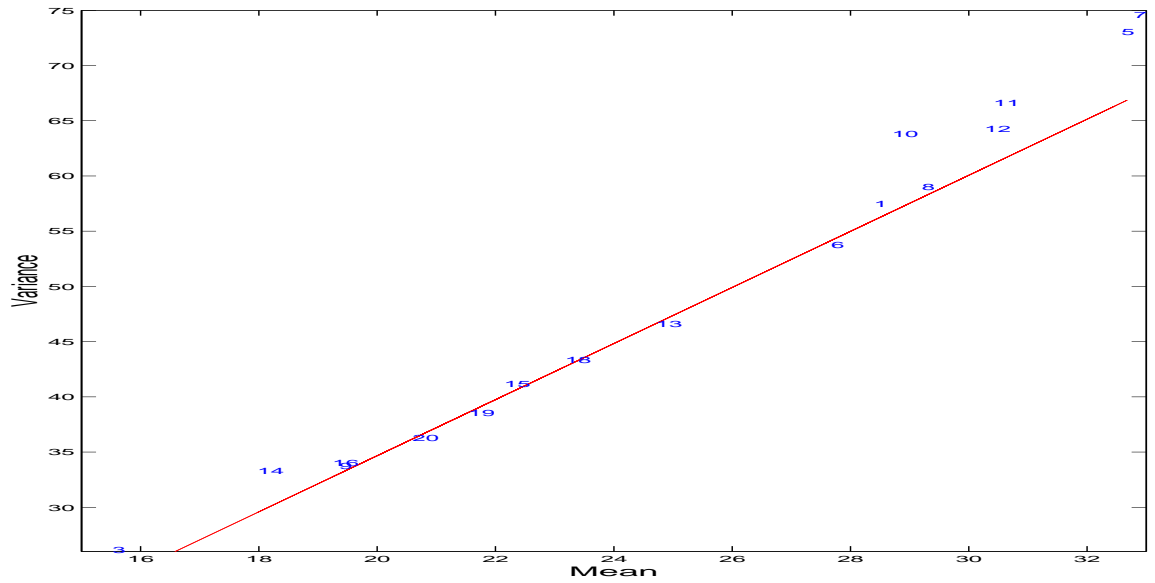


Figure 7: The variances of all 2002 UNC packet counts time series are plotted against their means. The numbers in the plot are identification numbers of time series. It shows a linear relationship between the mean and the variance.

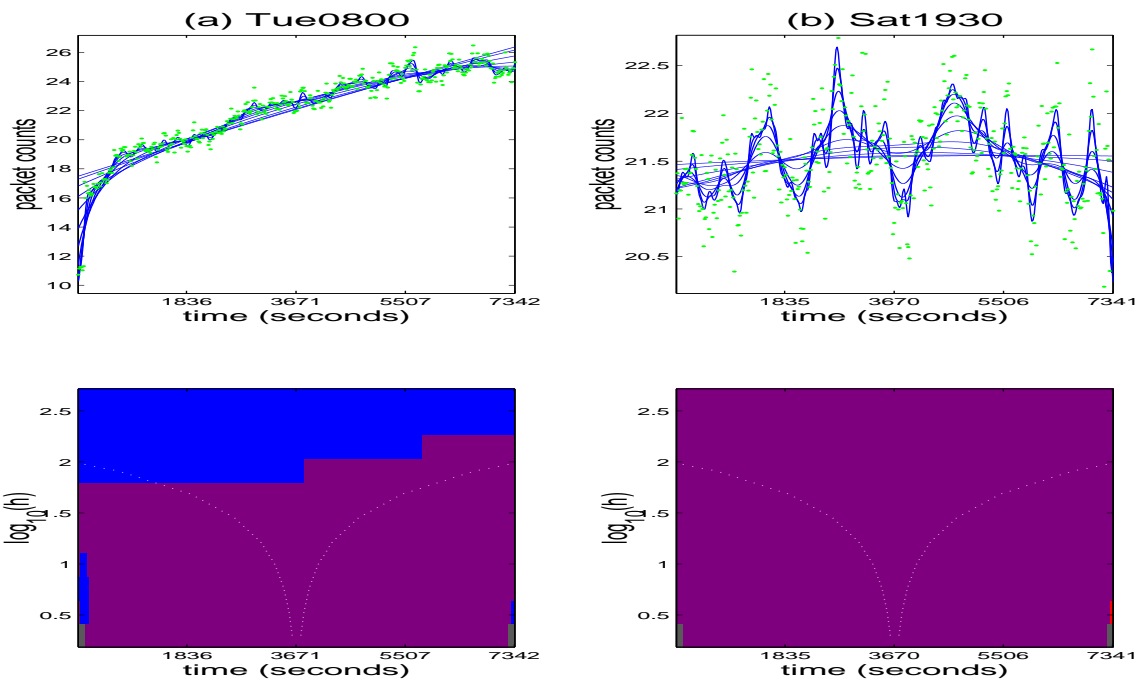


Figure 8: (a) Dependent SiZer plot of the Tue0800 time series. The time series is tested against FGN. A large increasing trend is present and this is significantly different from FGN. (b) Dependent SiZer plot of the Sat1930 time series. This time series has no significant features beyond what is expected from FGN.

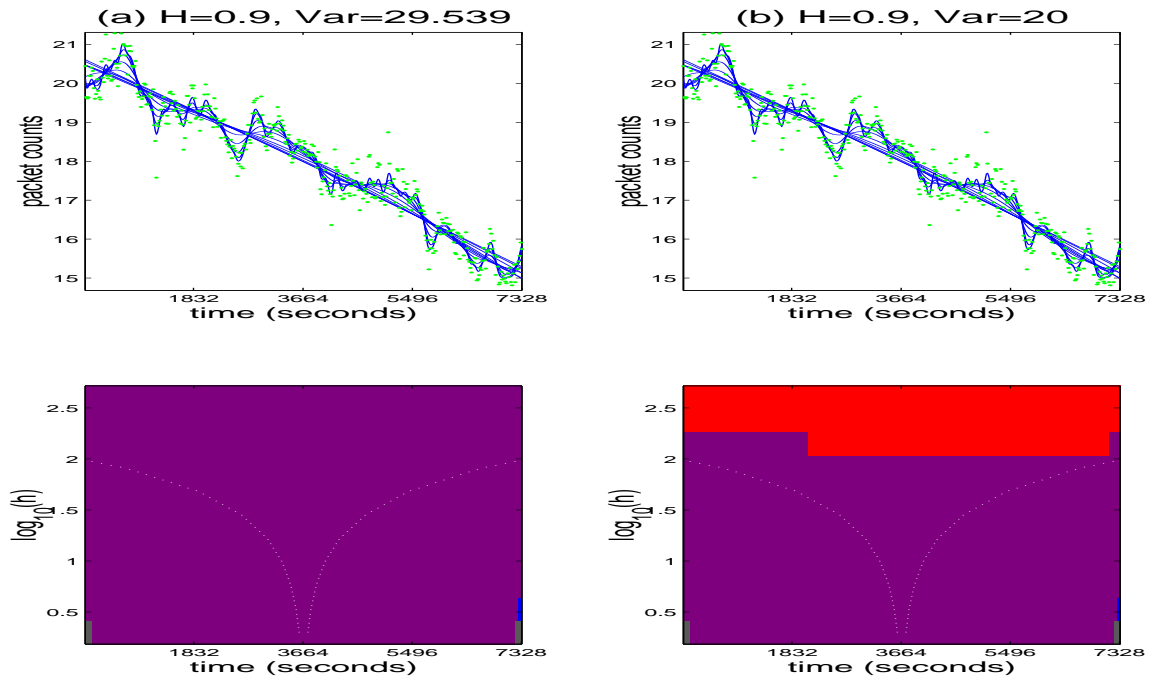


Figure 9: (a) Dependent SiZer of the Fri0300 time series tested against FGN with $H = 0.9$ and $\sigma^2 = 29.539$. The SiZer map does not show strong evidence of difference from FGN for this model. (b) Dependent SiZer of the Fri0300 time series tested against FGN with $H = 0.9$ and $\sigma^2 = 20$. The red colors at coarse scales of the SiZer map show there is a significant difference between the time series and FGN for these parameters. This example shows the importance of parameter estimation of an assumed model.

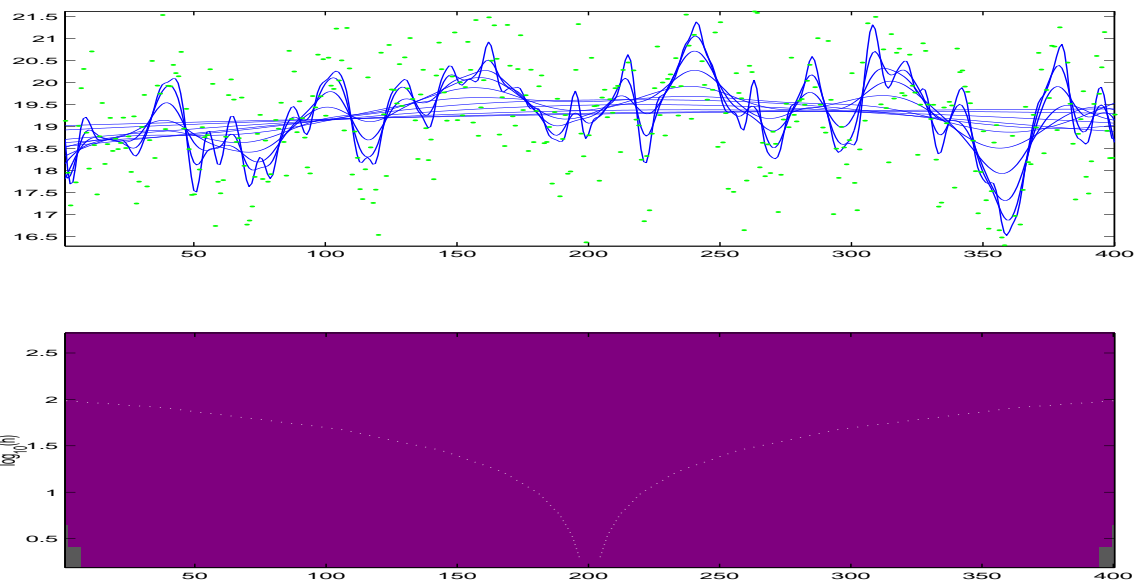


Figure 10: Dependent SiZer for simulated FGN data, with $H = 20$ and $\sigma^2 = 20$, against an ARMA(5,3) model. The SiZer map shows no significant features. This example demonstrates dependence fragility of dependent SiZer.