Thesis Proposal
The Short Time Fourier Transform and Local Signals

Shuhei Okumura

June 2, 2010

1 Introduction/Summary

In this thesis, I am going to investigate the properties of the short time Fourier transform (STFT) with overlapping windows and its uses in detecting local (in time) signals. The STFT applies the Fourier transform within a sliding window and is useful for analyzing local features of time series data. Applied to a white noise time series, the STFT results in a multivariate complex-valued stationary time series with cross-covariance functions derived here. This series can be well approximated by autoregressive models with a small order. Using such STFT properties, I aim to develop methods to detect local signals.

2 Previous Work

A large number of the Fourier transform applications can be found with partitioned, non-overlapping windows which can be seen as partitioned discrete Fourier transform (DFT). Some authors call this the STFT (Shumway and Stoffer 2006). Percival and Walden (1993) mention a variety of methods to smooth neighboring DFT coefficients, for example.

However, I have not found any research paper that deals with the STFT with an overlapping window that moves one data point at a time.

The wavelet transform, which does not segment data, is also widely used for studying local signals, but is currently beyond the scope of this thesis.

3 DFT and STFT

Throughout this document, suppose we have a univariate (possibly complex-valued) time series $\{X_t\}_{t=0}^{M-1}$ of length M.

The discrete Fourier transform (DFT) is given by

$$A_k = \frac{1}{\sqrt{M}} \sum_{j=0}^{M-1} X_j \omega_M^{-jk} , \qquad k = 0, 1, \ldots, M-1$$

where

$$\omega_M = \exp(i2\pi/M).$$

When the time series is real-valued, we focus on the range $k = 0, \ldots, \lfloor M/2 \rfloor$ because the ignored range is just the complex conjugate of the focused range. We can recover the original time series by applying the inverse DFT to the DFT coefficients, where the inverse DFT is given by

$$X_j = \frac{1}{\sqrt{M}} \sum_{k=0}^{M-1} A_k \omega_M^{jk} , \qquad j = 0, 1, \ldots, M-1$$

The short time Fourier transform (STFT) is computed by applying the DFT to a continuous subset of the data with $N \leq M$ data points. We call this range the window. Then we move the starting position of the window over the appropriate index range;

$$A_k^t = \frac{1}{\sqrt{N}} \sum_{j=0}^{N-1} X_{j+t-N+1} \omega_N^{-jk} , \qquad k = 0, 1, \ldots, N-1; \quad t = N-1, \ldots, M-1.$$

Now we have a complex-valued array $\mathcal{A}$ of size N-by-$(M - N + 1)$. $\{A^t\}$ can be seen as an N-dimensional complex-valued time series of length $(M - N + 1)$. We think of the time index t in the horizontal direction and k in the vertical direction of the array $\mathcal{A}$.

3.1 STFT on a White Noise Time Series

When the original time series X_t is a white noise time series, its STFT will form a multivariate time series of interest. A real-valued time series X_t is called white noise, if

$$\mathrm{E}[X_t] = 0 \ \forall t, \qquad \mathrm{Var}[X_t] = \sigma^2 \text{ (constant) } \forall t, \qquad \text{and} \qquad \mathrm{Cov}[X_i, X_j] = 0 \ \forall i \neq j.$$

Since I have not found any universally accepted definition of complex-valued white noise, I just define it here. We assume $\mathrm{E}[\mathrm{Re}(X_t)] = \mathrm{E}[\mathrm{Im}(X_t)] = 0$ and denote $\mathrm{Var}[\mathrm{Re}(X_t)] = \sigma_{RR}$, $\mathrm{Var}[\mathrm{Im}(X_t)] = \sigma_{II}$, $\mathrm{E}[\mathrm{Re}(X_t)\mathrm{Im}(X_t)] = \sigma_{RI}$, $\mathrm{E}[(\mathrm{Re}(X_t))^2(\mathrm{Im}(X_t))^2] = \mu_{R^2I^2}$, $\mathrm{E}[(\mathrm{Re}(X_t))^4] = \mu_{R^4}$, $\mathrm{E}[(\mathrm{Im}(X_t))^4] = \mu_{I^4}$ $\forall t$ and $\mathrm{Cov}[X_i, X_j] = 0 \ \forall i \neq j$, which means $\mathrm{Cov}[\mathrm{Re}(X_i),\mathrm{Re}(X_j)] = \mathrm{Cov}[\mathrm{Im}(X_i),\mathrm{Im}(X_j)] = \mathrm{Cov}[\mathrm{Re}(X_i),\mathrm{Im}(X_j)] = \mathrm{Cov}[\mathrm{Im}(X_i),\mathrm{Re}(X_j)] = 0$.

There are four ways of looking at one k (one row) of the STFT.
(1) complex-valued $\{A_k^t\}_t$,
(2) the real part $\{\mathrm{Re}(A_k^t)\}_t$,
(3) the imaginary part $\{\mathrm{Im}(A_k^t)\}_t$, and
(4) the squared modulus $\{|A_k^t|^2\}_t$.

Note that each of these time series is stationary. We have found the autocovariance function, cross-covariance function (between different k's), and the spectrum for each of the four time series. We note (1) is a univariate complex-valued moving average process; (2) and (3) are the widely-used univariate real-valued moving average processes (Brockwell and Davis 1991) with order $N - 1$; and (4) is the sum of the squared real part and the squared imaginary part, and therefore is univariate real-valued. It does not have a universal name. We will refer to it as the squared modulus time series.

Another time series is the phase of A_k^t, or the angle between the positive real-axis of the complex plane and the complex value, $\mathrm{angle}(A_k^t) = \mathrm{Arg}(A_k^t) = \mathrm{atan2}(\mathrm{Im}(A_k^t),\mathrm{Re}(A_k^t))$. This is a nonlinear time series. Known properties are: 1) the marginal distribution is uniform on $[-\pi, \pi]$ when X_t is Gaussian; and 2) the cross-covariance (between two different k's) functions

are zero at lags greater than or equal to the window size N. The bivariate probability density function of (angle(A_k^t), angle(A_ℓ^{t+h})) is hard to find analytically, because it involves transforming 4 random variables with non-independence structure. Later we will look at one example of simulated data and its STFT angle time series in Figure 2, where we clearly see nonlinearity. In such a case, the cross-covariance functions are no longer appropriate measures for the dependence of nonlinear time series (Fan and Yao 2003). One problem is that the time series angle(A_k^t) "jumps" near the boundaries. For example, if $A_k^t = -1 + 0.01i$ and $A_k^{t+1} = -1 - 0.01i$, then angle(A_k^t) = 3.131593 and angle(A_k^{t+1}) = $-3.131593 \neq \text{Arg}(1 + 0.01i) + \pi = 3.151592 = -3.131593 + 2\pi$. This phenomenon makes it hard to utilize this angle time series for detecting local signals, while other series (1)-(4) can do the task, as shown later.

3.2 Theoretical Properties

Let $CCF_{X,Y}(h)$ denote the cross-covariance function of two time series X_t and Y_t at lag h. For each of the four types of time series, $CCF_{X,Y}(h) = 0$ for $|h| \geq N$. Also, $E[A_k^t] = 0$ and $E[|A_k^t|^2] = \sigma_{RR} + \sigma_{II}$ $\forall t$ and k. The autocovariance function of a time series is the CCF of itself, so we only compute the CCF. For $0 \leq h < N$, where $*$ denotes the complex conjugate,

$$\begin{aligned}
\text{with letting } c_k(m) &= \cos\left(\frac{2\pi km}{N}\right) \text{ and } s_k(m) = \sin\left(\frac{2\pi km}{N}\right) \\
CCF_{A_k,A_\ell}(h) &= E[(A_k^{t+h} - E[A_k^{t+h}])(A_\ell^t - E[A_\ell^t])^*] = \frac{\sigma_{RR} + \sigma_{II}}{N} \sum_{m=0}^{N-h-1} \omega_N^{-km+\ell(m+h)} \\
CCF_{Re(A_k),Re(A_\ell)}(h) &= \frac{1}{N}\left[\sigma_{RR} \sum_{m=0}^{N-h-1} c_k(m)c_\ell(m+h) + \sigma_{II} \sum_{m=0}^{N-h-1} s_k(m)s_\ell(m+h)\right. \\
&\quad \left. + \sigma_{RI} \sum_{m=0}^{N-h-1} c_k(m)s_\ell(m+h) + \sigma_{RI} \sum_{m=0}^{N-h-1} s_k(m)c_\ell(m+h)\right] \\
CCF_{Im(A_k),Im(A_\ell)}(h) &= \frac{1}{N}\left[\sigma_{RR} \sum_{m=0}^{N-h-1} s_k(m)s_\ell(m+h) + \sigma_{II} \sum_{m=0}^{N-h-1} c_k(m)c_\ell(m+h)\right. \\
&\quad \left. - \sigma_{RI} \sum_{m=0}^{N-h-1} c_k(m)s_\ell(m+h) - \sigma_{RI} \sum_{m=0}^{N-h-1} s_k(m)c_\ell(m+h)\right] \\
CCF_{Re(A_k),Im(A_\ell)}(h) &= \frac{1}{N}\left[-\sigma_{RR} \sum_{m=0}^{N-h-1} c_k(m)s_\ell(m+h) + \sigma_{II} \sum_{m=0}^{N-h-1} s_k(m)c_\ell(m+h)\right. \\
&\quad \left. + \sigma_{RI} \sum_{m=0}^{N-h-1} c_k(m)c_\ell(m+h) - \sigma_{RI} \sum_{m=0}^{N-h-1} s_k(m)s_\ell(m+h)\right] \\
CCF_{Im(A_k),Re(A_\ell)}(h) &= \frac{1}{N}\left[-\sigma_{RR} \sum_{m=0}^{N-h-1} s_k(m)c_\ell(m+h) + \sigma_{II} \sum_{m=0}^{N-h-1} c_k(m)s_\ell(m+h)\right. \\
&\quad \left. + \sigma_{RI} \sum_{m=0}^{N-h-1} c_k(m)c_\ell(m+h) - \sigma_{RI} \sum_{m=0}^{N-h-1} s_k(m)s_\ell(m+h)\right]
\end{aligned}$$

$$
\begin{aligned}
CCF_{|A_k|^2,|A_\ell|^2}(h) &= -(\sigma_{RR}+\sigma_{II})^2 + \frac{1}{N^2}\left[(N-h)(\mu_{R^4}+\mu_{I^4}+2\mu_{R^2I^2})\right. \\
&\quad +(N^2-N+h)((\sigma_{RR})^2+(\sigma_{II})^2+2\sigma_{RR}\sigma_{II}) \\
&\quad +4(\sigma_{RI})^2 \sum_{p=0}^{N-h-1}\sum_{q=0}^{N-h-1} s_k(p-q)s_\ell(q-p) \\
&\quad +4\mu_{R^2I^2} \sum_{p=0}^{N-h-1}\sum_{q=0}^{N-h-1} s_k(p-q)s_\ell(p-q) \\
&\quad \left.+2((\sigma_{RR})^2+(\sigma_{II})^2+2(\sigma_{RI})^2) \sum_{p=0,p\neq q}^{N-h-1}\sum_{q=0}^{N-h-1} c_k(p-q)c_\ell(p-q)\right]
\end{aligned}
$$

These hold for both $k \leq \ell$ and $k \geq \ell$ and for any Gaussian or non-Gaussian input, as long as the moments up to the fourth are the same. For a real-valued time series, we can let $\sigma_{II} = \sigma_{RI} = \mu_{I^4} = \mu_{R^2I^2} = 0$. We can also compute the spectral density of a univariate time series X_t with $f(\nu) = \sum_{h=-\infty}^{\infty} CCF_{x,x}(h)e^{-i2\pi\nu h}$, and the cross spectrum between two series X_t and Y_t, $f_{xy}(\nu) = \sum_{h=-\infty}^{\infty} CCF_{x,y}(h)e^{-i2\pi\nu h}$ for $-0.5 \leq \nu \leq 0.5$.

3.3 Approximation by Autoregressive Models

When we know the theoretical autocovariance function of a stationary time series, we can also compute the partial autocorrelation function (PACF) with the Durbin-Levinson algorithm. It is known that an autoregressive process with order p (AR(p)) has PACF exactly equal to zero at lags greater than p. It can be shown that each of the three real-valued time series $\{\mathrm{Re}(A_k^t)\}_t$, $\{\mathrm{Im}(A_k^t)\}_t$, and $\{|A_k^t|^2\}_t$ has PACF very close to zero after a few lags for any k and for any window size N. This indicates that we can approximate these time series with AR(p) with relatively small p, possibly p=1.

As mentioned earlier, the two time series $\{\mathrm{Re}(A_k^t)\}_t$ and $\{\mathrm{Im}(A_k^t)\}_t$ are moving average processes with order $N-1$. It can be shown that each time series is equivalent to a seasonal autoregressive moving average model SARMA(1, 0)$\times$(0, 1). The conditions under which such model equivalence holds are under investigation as well.

3.4 An Example

Here we look at the STFT of a generated Gaussian white noise time series. Figure 1 shows a complex-valued Gaussian white noise time series of length 50 (with $E[\mathrm{Re}(X_t)] = E[\mathrm{Im}(X_t)] = 0$, $Var[\mathrm{Re}(X_t)] = Var[\mathrm{Im}(X_t)] = 1$, and $E[\mathrm{Re}(X_t)\mathrm{Im}(X_t)] = 0.5$), and the squared modulus, the real and imaginary parts of the resulting STFT with window size $N = 10$ (and thus $k = 0, \ldots, 9$). As expected, it is hard to visually see any pattern in the STFT, except that neighboring values are often similar.

Figure 2 shows two time series angle(A_2^t) and angle(A_3^t) and their one- and two-step functions. The two time series are bounded on $[-\pi, \pi]$ and show "jumps" as we discussed earlier. The four scatter plots are 1) angle(A_2^{t-1}) against angle(A_2^t), 2) angle(A_2^{t-2}) against angle(A_2^t), 3) angle(A_3^{t-1}) against angle(A_2^t), and 4) angle(A_3^{t-1}) against angle(A_2^t). They clearly show nonlinearity and indicate that the cross-covariance functions are not appropriate measures for the dependence of these nonlinear time series.

4 STFT on a Simple Local Signal

In this section, we examine the STFT of data that contains a sinusoidal function and discuss the detection of such functions. The traditional DFT may not detect the presence of local signals while the STFT can carry out the task better.

4.1 On A Simple Global Signal Without Noise

Suppose we have a complex-valued time series Y_t of length M with real-valued amplitudes A and B, the number of cycles K (not necessarily an integer) and real-valued phases ϕ_A and ϕ_B, where

$$Y_t = A\cos\left(\frac{2\pi Kt}{M} + \phi_A\right) + iB\cos\left(\frac{2\pi Kt}{M} + \phi_B\right) \text{ for } t = 0, 1, \ldots, M-1.$$

Let us call such signal with the same sinusoidal form throughout the time a *global signal*, (global in time) as opposed to a *local signal* (local in time) that will be seen in the next section. Letting k^* = the number of cycles within the STFT window (suppose this is an integer and $k^* \leq N/2$), we have explicit forms of the resulting STFT, denoted by G_k^t instead of A_k^t, for $t = N-1, \ldots, M-1$,

$$\begin{aligned}
G_{k^*}^t &= \frac{A\sqrt{N}}{2}\exp\left(i\left(\phi_A + \frac{2\pi k^*}{N}(t-N+1)\right)\right) + \frac{iB\sqrt{N}}{2}\exp\left(i\left(\phi_B + \frac{2\pi k^*}{N}(t-N+1)\right)\right) \\
&= \left[\frac{A\sqrt{N}}{2}\cos\left(\phi_A + \frac{2\pi k^*}{N}(t-N+1)\right) + \frac{-B\sqrt{N}}{2}\sin\left(\phi_B + \frac{2\pi k^*}{N}(t-N+1)\right)\right] \\
&\quad + i\left[\frac{A\sqrt{N}}{2}\sin\left(\phi_A + \frac{2\pi k^*}{N}(t-N+1)\right) + \frac{B\sqrt{N}}{2}\cos\left(\phi_B + \frac{2\pi k^*}{N}(t-N+1)\right)\right] \\
|G_{k^*}^t|^2 &= \frac{N(A^2+B^2)}{4} + \frac{ABN}{2}\sin(\phi_A - \phi_B) \\
G_{k^{**}}^t &= \left[\frac{A\sqrt{N}}{2}\cos\left(\phi_A + \frac{2\pi k^*}{N}(t-N+1)\right) + \frac{B\sqrt{N}}{2}\sin\left(\phi_B + \frac{2\pi k^*}{N}(t-N+1)\right)\right] \\
&\quad + i\left[\frac{-A\sqrt{N}}{2}\sin\left(\phi_A + \frac{2\pi k^*}{N}(t-N+1)\right) + \frac{B\sqrt{N}}{2}\cos\left(\phi_B + \frac{2\pi k^*}{N}(t-N+1)\right)\right] \\
|G_{k^{**}}^t|^2 &= \frac{N(A^2+B^2)}{4} + \frac{-ABN}{2}\sin(\phi_A - \phi_B)
\end{aligned}$$

G_k^t equals to $\mathbf{0}$ for all t at any k other than k^* and $k^{**} = N - k^*$. If k^* is not an integer, we will have a phenomenon called "leakage" from difference between the signal's frequency and the sampling frequency, which results in non-zero STFT at other k's than k^* and k^{**} (Cristi 2004).

Sliding the window is the same as applying the DFT to the same signal with the phases ϕ_A and ϕ_B changing. Perhaps the two squared modulus time series staying constant as a function of time is intuitive because the squared modulus of the DFT measures the amplitude of the signal and ignores the phases. In contrast, the real-part and imaginary part time series at each of k^* and k^{**} form sinusoidal signal outputs as the STFT window moves along.

4.2 On A Simple Local Signal Without Noise

Unlike global signals, local signals appear only at parts of the data between the starting point S and the ending point E ($0 < S < E < M-1$). We assume the signal value to be zero where the signal is not present. Such a signal can be obtained by applying an appropriate indicator function to the above global signal Y_t for $t = 0, 1, \ldots, M-1$;

$$X_t = I_{(S \le t \le E)} \cdot Y_t = I_{(S \le t \le E)} \cdot \left[A \cos\left(\frac{2\pi K t}{M} + \phi_A \right) + iB \cos\left(\frac{2\pi K t}{M} + \phi_B \right) \right]$$

This representation of a simple local function X_t has seven parameters; amplitudes A and B, the number of cycles K, phases ϕ_A and ϕ_B, starting point S, and ending point E.

When applied to the zero-valued region at the beginning, the resulting STFT is zero at all the k's for $t = N-1, \ldots, S-1$, and also at the ending; for $t = E+N, \ldots, M-1$.

When the STFT window is only on the local signal, for $S^* \le t \le E$, where $S^* = S+N-1$, the STFT A_k^t is exactly equal to the STFT G_k^t applied to the global signal Y_t; sinusoidal signal outputs at $k = k^*$ and k^{**}, and zero at any other k.

Now, the STFT gets more complicated when the window covers both the zero-valued region at the beginning and the local signal; for $S \le t < S^*$. The STFT results in non-zero at other k's as well. This phenomenon is called "ringing" (Percival and Walden 1993). Ringing occurs when the DFT is applied to a region with discontinuity, which in this particular case is the change from the zero constant to the periodic function. For any k and $1 \le d \le N-1$,

$$A_k^{S^*-d} = G_k^{S^*-d} - \frac{1}{\sqrt{N}} \sum_{j=0}^{d-1} Y_{S-d+j} \omega_N^{-jk}$$

Similarly, when the window covers both the end of the local signal and the following zero-valued region (for $E < t \le E+N-1$),

$$A_k^{E+d} = G_k^{E+d} - \frac{1}{\sqrt{N}} \sum_{j=0}^{d-1} Y_{E+d-j} \omega_N^{-(N-1-j)k}$$

When N/k^* is an integer, simpler expressions exist at k^* (and k^{**}) and various time points. For $j = 0, 1, \ldots, k^*$,

$$\begin{aligned} A_{k^*}^{S^*-(N/k^*)j} &= \left(1 - \frac{j}{k^*}\right) A_{k^*}^{S^*} = \left(1 - \frac{j}{k^*}\right) G_{k^*}^{S^*} \\ A_{k^*}^{E+(N/k^*)j} &= \left(1 - \frac{j}{k^*}\right) A_{k^*}^{E} = \left(1 - \frac{j}{k^*}\right) G_{k^*}^{E} \end{aligned}$$

4.3 An Example

We provide a simple example to illustrate how the STFT works on a local signal. The time series $\{X_t\}_{t=0}^{49}$ in Figure 3 consists of zeros at the beginning and at the end, and a cosine function of length 20 with periodicity 4 and amplitude 5 in the middle from $t = 15$ to $t = 34$. This local signal can be described with the representation in Section 4.1 with $M = 50$, $A = 5$, $B = 0$, $K = 10$, $\phi_A = \phi_B = 0$, $S = 15$, and $E = 34$.

We use a window size $N = 10$ that results in $k = 0, \ldots, 9$. The k that matches the signal's frequency is $k^* = 2$ (and thus $k^{**} = 8$). We examine the STFT in three paragraphs below.

When the STFT window is on the zero valued region at the beginning and at the end (for $t = 9, \ldots, 14, 44, \ldots, 49$), the complex-valued STFT is zero at all the k's.

When the STFT window is only on the cosine function (for $t = 24, \ldots, 34$), the STFT behaves exactly the same way as it does for a global signal: the squared modulus STFT is constant at $k = 2$ and 8 zero at all other k's over the region $|A^t|^2 = (0, 0, 62.5, 0, 0, 0, 0, 0, 62.5, 0)$, and the real and imaginary STFT produce sinusoidal signal outputs at $k = 2$ and 8, and equal to zero at all other k's.

When the STFT window is on both a zero-valued region and the local cosine function (for $t = 15, \ldots, 23, 35, \ldots, 43$), we have $0 \leq |A_2^t| \leq 62.5$, and the more the window covers the cosine function, the higher the squared modulus is. As given by the simple expressions, $A_2^{19} = 3.952847 - 0i = A_2^{24}/2 = (7.905694 - 0i)/2$ and $A_2^{39} = 3.952847 - 0i = A_2^{34}/2 = (7.905694 - 0i)/2$. In general, the expression for the resulting STFT in these regions does not simplify because of the ringing phenomenon.

4.4 The Linearity of DFT and STFT

Here we mention an important property of the DFT and STFT, that is, they are a linear function which means that the sum of the DFT (or STFT) of data X_t and the DFT (or STFT) of data Y_t equals to the DFT (or STFT) of $X_t + Y_t$. This holds for the complex-valued time series $\{A_k^t\}_t$ (and therefore also for $\{\mathrm{Re}(A_k^t)\}_t$ and $\{\mathrm{Im}(A_k^t)\}_t$), but not for the squared modulus time series $\{|A_k^t|^2\}_t$. By this property, we would know how the STFT behaves when a cosine function exists in the middle of the white noise time series, because the resulting STFT would be the sum of the first STFT and the STFT of the white noise, as in the previous example, The subsequent examples are to illustrate how the STFT changes as the noise level increases.

4.5 With White Noise in the Background

Here we look at some examples of how STFT results when a local signal exists with a Gaussian white noise background with noise levels, 0.5, 2, 4 and 10 standard deviations . The data displayed on the top panel of Figures 4 through 7 were generated by adding the simple cosine function from the earlier example to IID Gaussian white noise with zero mean and each standard deviation for the real and imaginary parts. The correlation between the real and imaginary parts is 0.5 for each case. The STFT of the data are shown in the other plots. When the noise level is low, we can clearly see the local signal and the resulting STFT is not very different from Figure 3 of the simple cosine example. As the noise level increases (relative to the amplitude of the cosine function), the pattern becomes less obvious.

5 STFT on a General Local Signal

Here we consider more general cases. By the (inverse) Fourier representation, each of the global signal $\{g_t\}_{t=0}^{M-1}$ with length M and the local signal $\{\ell_t\}_{t=S}^{E}$ with length $P(= S - E + 1)$ can be described as a linear sum of periodic functions with different frequencies. (The above example consisted of only two non-zero Fourier coefficients.)

$$\text{Global Signal: } g_t = \frac{1}{\sqrt{M}} \sum_{a=0}^{M-1} G_a \omega_M^{at} \qquad \text{for } t = 0, 1, \ldots, M-1.$$

$$\text{Local Signal: } \ell_t = \begin{cases} \frac{1}{\sqrt{P}} \sum_{b=0}^{P-1} L_b \omega_P^{bt} & \text{for } t = S, \ldots, E \\ 0 & \text{otherwise} \end{cases}$$

This representation allows the local signal ℓ_t to be an arbitrary function. For $N \leq P$ and $S \leq t \leq E - N + 1$ where the STFT window is completely inside the local signal,

$$A_k^t = \frac{1}{\sqrt{MN}} \sum_{\substack{a=1 \\ aN/M \notin \mathbb{Z}}}^{M-1} G_a \omega_M^{a(t-N+1)} \sum_{j=0}^{N-1} \omega_N^{j(\frac{N}{M}a-k)} + \frac{1}{\sqrt{NP}} \sum_{\substack{b=1 \\ bN/P \notin \mathbb{Z}}}^{P-1} L_b \omega_P^{b(t-N+1)} \sum_{j=0}^{N-1} \omega_N^{j(\frac{N}{P}b-k)}$$

As before, if M and N have no common positive factor other than 1, we will have leakage.

Now,
1) when the window is crossing S ($S \leq t \leq S + N - 1$),
2) when the window is crossing E ($E \leq t \leq E + N - 1$), or
3) $N < P$ and the local signal is completely inside the window ($S + N - 1 \leq t \leq E$),

$$A_k^t = \frac{1}{\sqrt{MN}} \sum_{\substack{a=1 \\ aN/M \notin \mathbb{N}}}^{M-1} G_a \omega_M^{a(t-N+1)} \sum_{j=0}^{N-1} \omega_N^{j(\frac{N}{M}a-k)} + \frac{1}{\sqrt{NP}} \sum_{b=0}^{P-1} L_b \omega_P^{bt} \sum_{j=\max\{0,S-t+N-1\}}^{\min\{N-1,E-t+N-1\}} \omega_N^{j(\frac{N}{P}b-k)}$$

Cases 1) and 2) will result in the phenomenon called "ringing" as we saw in the simple example without noise.

6 Proposed Work

6.1 Generalizations to More Complicated Signals

I would like to show how the STFT behaves when the local signal is more complicated. One obvious extensions of the above example would be a case where the window size does not match the frequency of the local signal, thus resulting in "leakage." We can also consider a case where multiple local signals exist with different starting and ending points. So far the uses of DFT and STFT have had something to do with sinusoidal functions. We can further consider a case where the local signal is not a sinusoidal function. Lastly, we can complicate the background noise, where the background noise is not a white noise time series but some stationary time series such as AR(1). Its STFT would be stationary and we can find ACF, CCF, and spectrum, since the STFT is a linear filter, whereas the (squared or non-squared) modulus STFT is not.

6.2 Detection

Now that we know how the STFT behaves, we can construct ways to detect such local signals. One way is to simply look for a cluster where $|A_k^t|^2$ observations are considerably high. Another is to regress $\{|A_k^t|^2\}_t$ with AR(p) with the appropriate coefficient and look for

a cluster of unusual residuals where a local cosine function is present and the window covers it partially. As can be seen in Figure 3, the real and imaginary parts of the STFT of a local signal are characteristic. Regressing them with AR(p) results in unusually volatile residuals when the noise level is small.

One time domain approach is to find the MLE of a local signal with eight parameters: 1-2) amplitudes A and B, 3-4) phases ϕ_A and ϕ_B, 5) starting point S, 5) ending point E, and 6-8) background noise covariance matrix, although this does not utilize STFT.

I aim to develop methods to quantify the likelihood of the local signal existence such as p-values with the frequentist and Bayesian approaches.

6.3 Application to Biomedical Engineering

Although Magnetoencephalography (MEG) data is most likely to be too noisy, it is worth trying, because of its power, sampling frequency (can be set at more than 1 kHz), and wide-ranging applications. One advantage of using MEG is that it is relatively easy to obtain hundreds, or even thousands, of trials of one simple experiment. Another advantage is that an MEG has 306 channels with neighboring channels measuring the same signal source, so we will have a multi-dimensional original time series with which to work. Often brain signals are assumed sinusoidal, so the use of the STFT may be sensible.

7 References

Brockwell, P. J., and Davis, R. A. (1991) *Time Series: Theory and Methods* (Second Edition). New York: Springer-Verlag.

Chonavel, T. (2002) *Statistical Signal Processing.* New York: Springer-Verlag.

Cristi, R. (2004) *Modern Digital Signal Processing.* Pacific Grove, California: Brooks/Cole-Thomson Learning.

Fan, J., and Yao, Q. (2003) *Nonlinear Time Series: Nonparametric and Parametric Methods.* New York: Springer-Verlag.

Gray, R. M., and Davidsson, L. D. (2005) *An Introduction to Statistical Signal Processing.* Cambridge: Cambridge University Press.

Higgins, J. R., and Stens, R. L. (1996) *Sampling Theory in Fourier and Signal Analysis.* Oxford: Oxford University Press.

Hippenstiel, R. D. (2001) *Detection Theory: Applications and Digital Signal Processing.* Boca Raton, Florida: CRC Press.

Niedzwiecki, M. (2000) *Identification of Time-Varying Processes.* New York: John Wiley & Sons.

O Ruanaidh, J. J. K., and Fitzgerald, W. J. (1996) *Numerical Bayesian Methods Applied to Signal Processing.* New York: Springer-Verlag.

Percival, D. B., and Walden, A. T. (1993) *Spectral Analysis for Physical Applications.* Cambridge: Cambridge University Press.

Shumway, R. H., and Stoffer, D.S. (2006) *Time Series Analysis and Its Applications: With R Examples.* New York: Springer-Verlag.

8 Figures

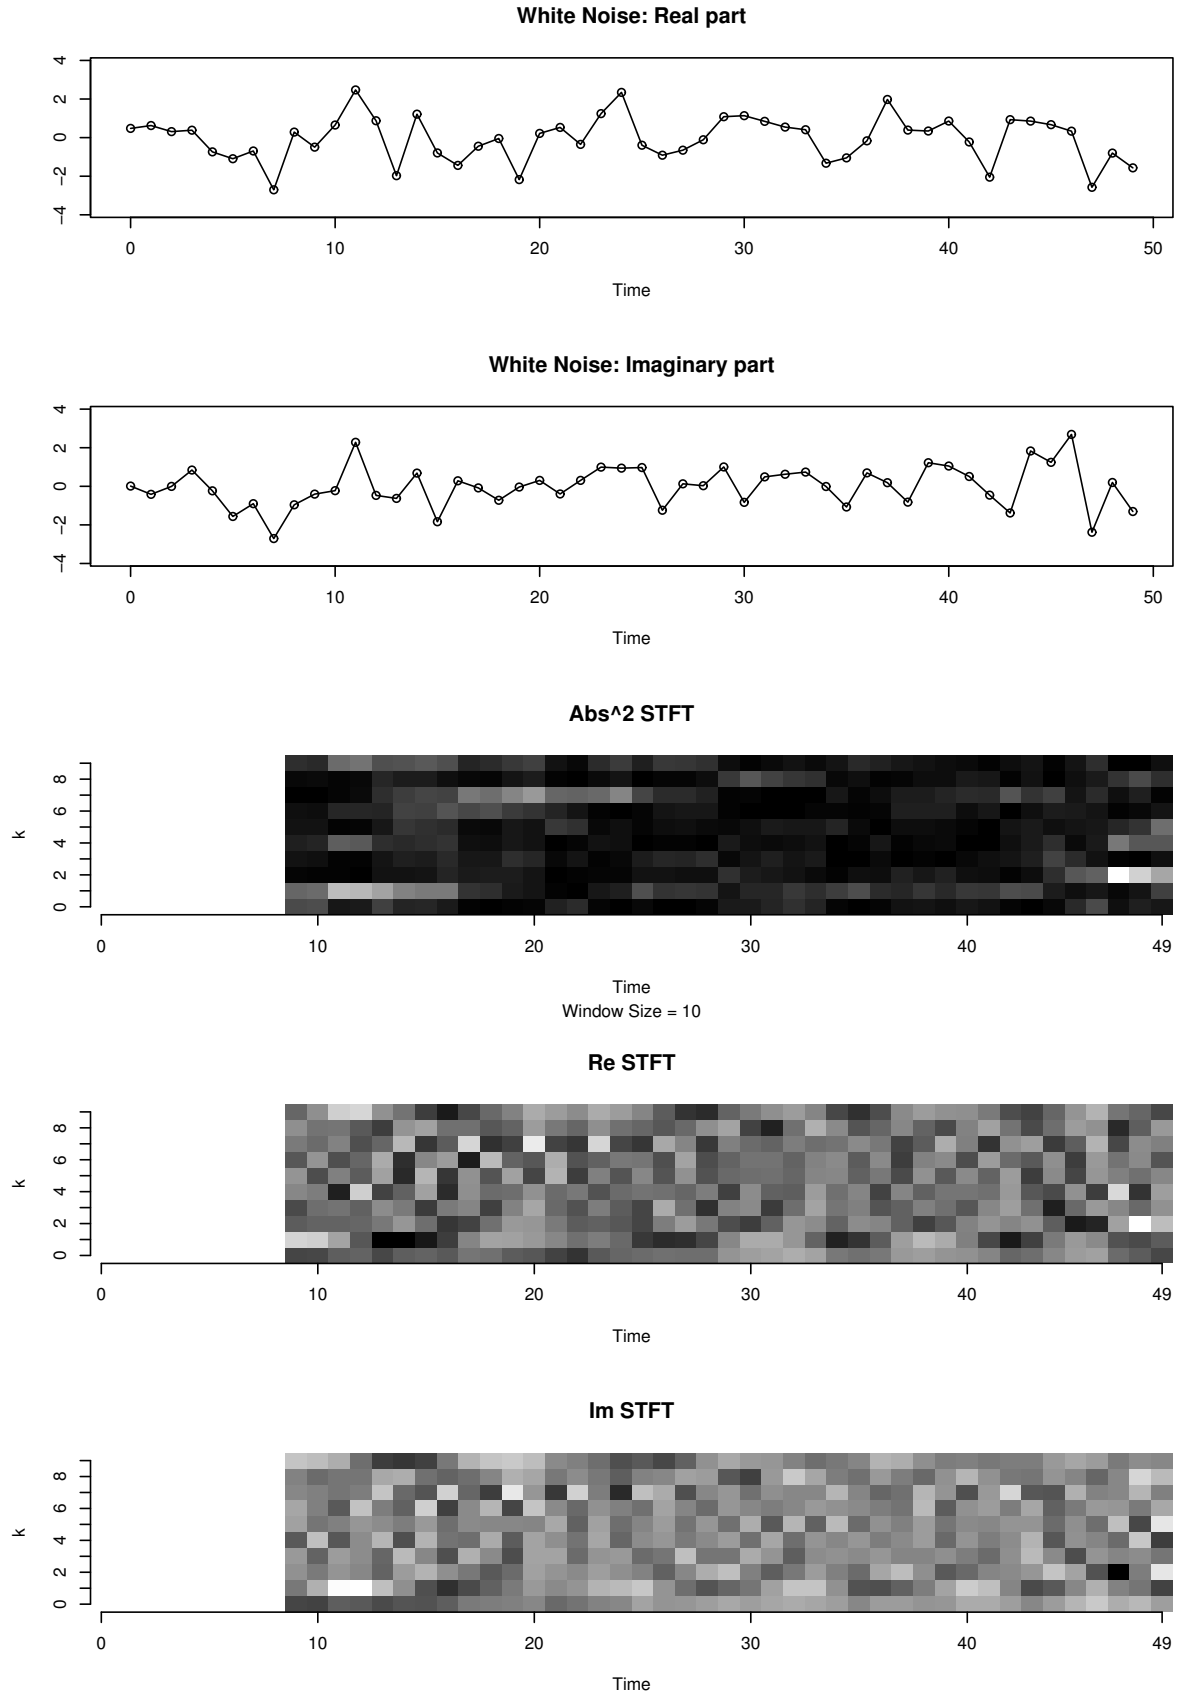

Figure 1: STFT of a Gaussian white noise time series. No visually obvious pattern exists, except neighboring points are often similar.

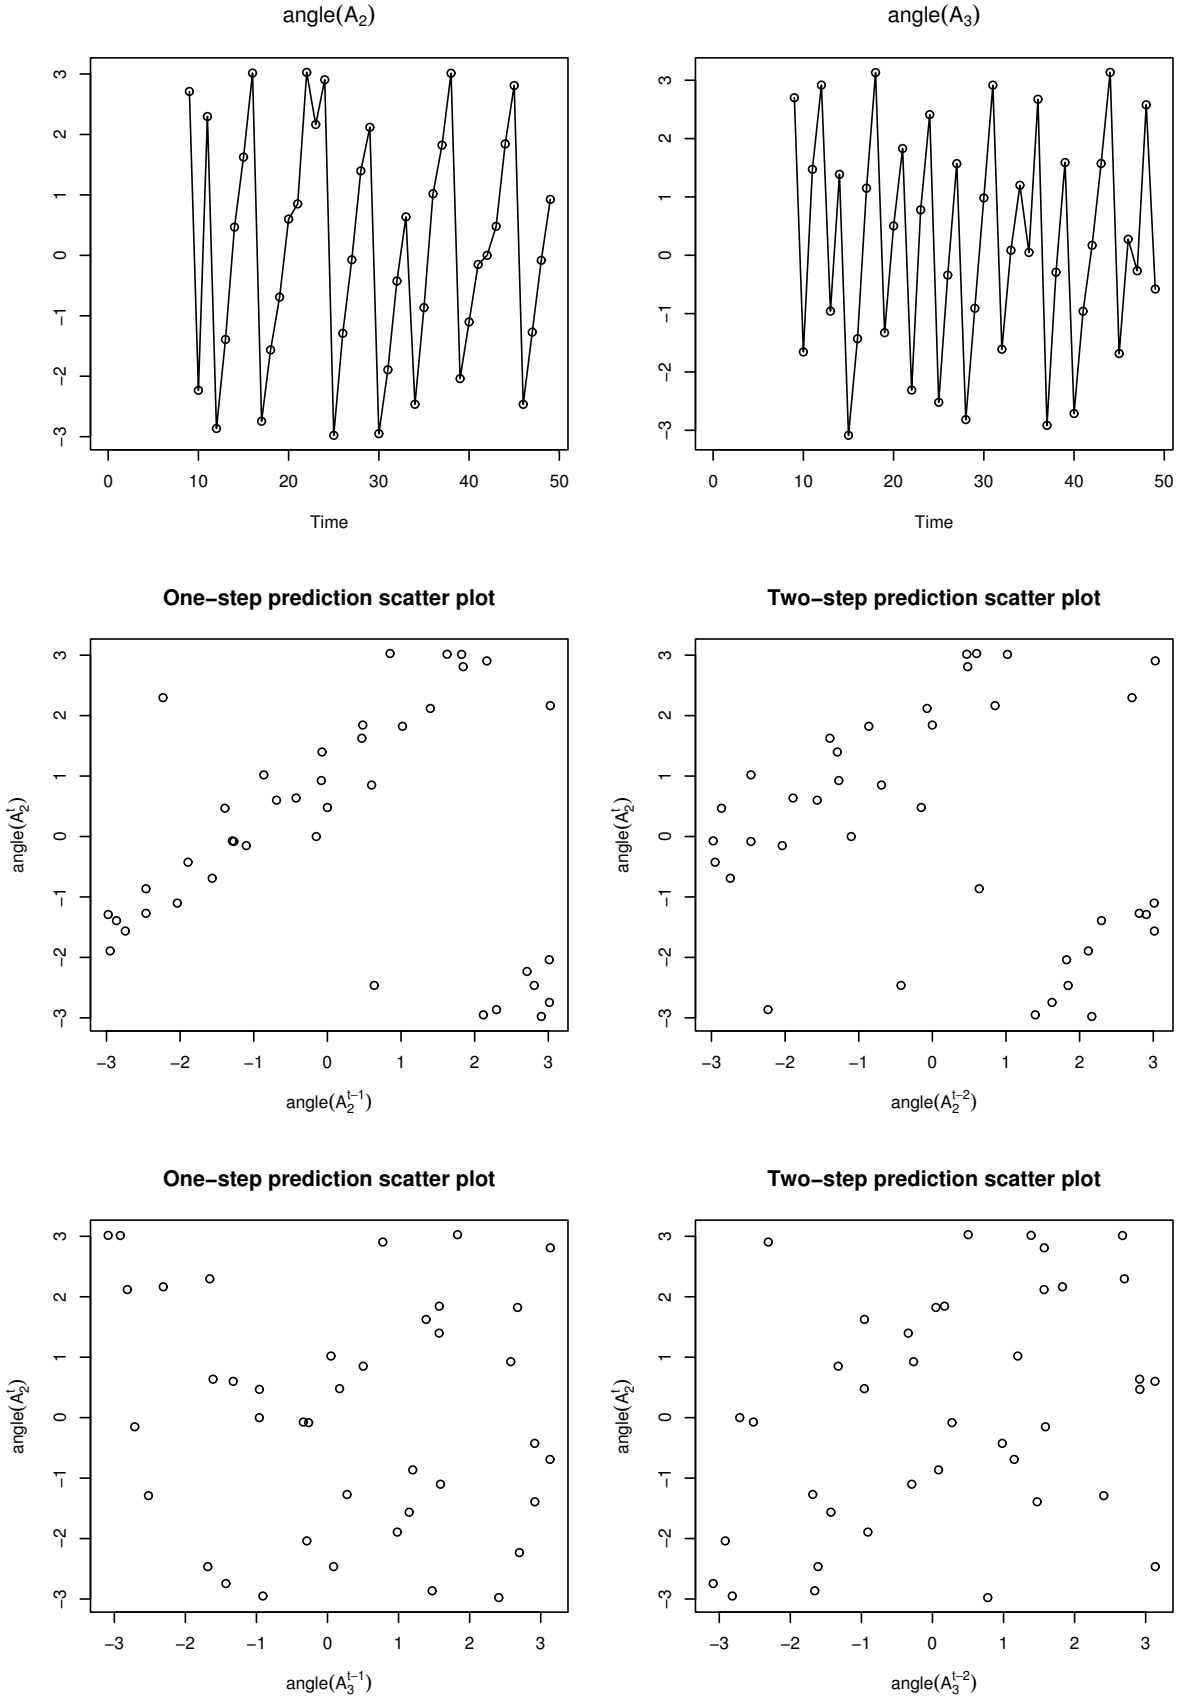

Figure 2: Two time series of angle(A_2^t) and angle(A_3^t) from the example in Figure 1. The scatter plots of one- and two-step functions indicate that the cross-covariance functions are not appropriate measures for the dependence of these nonlinear time series.

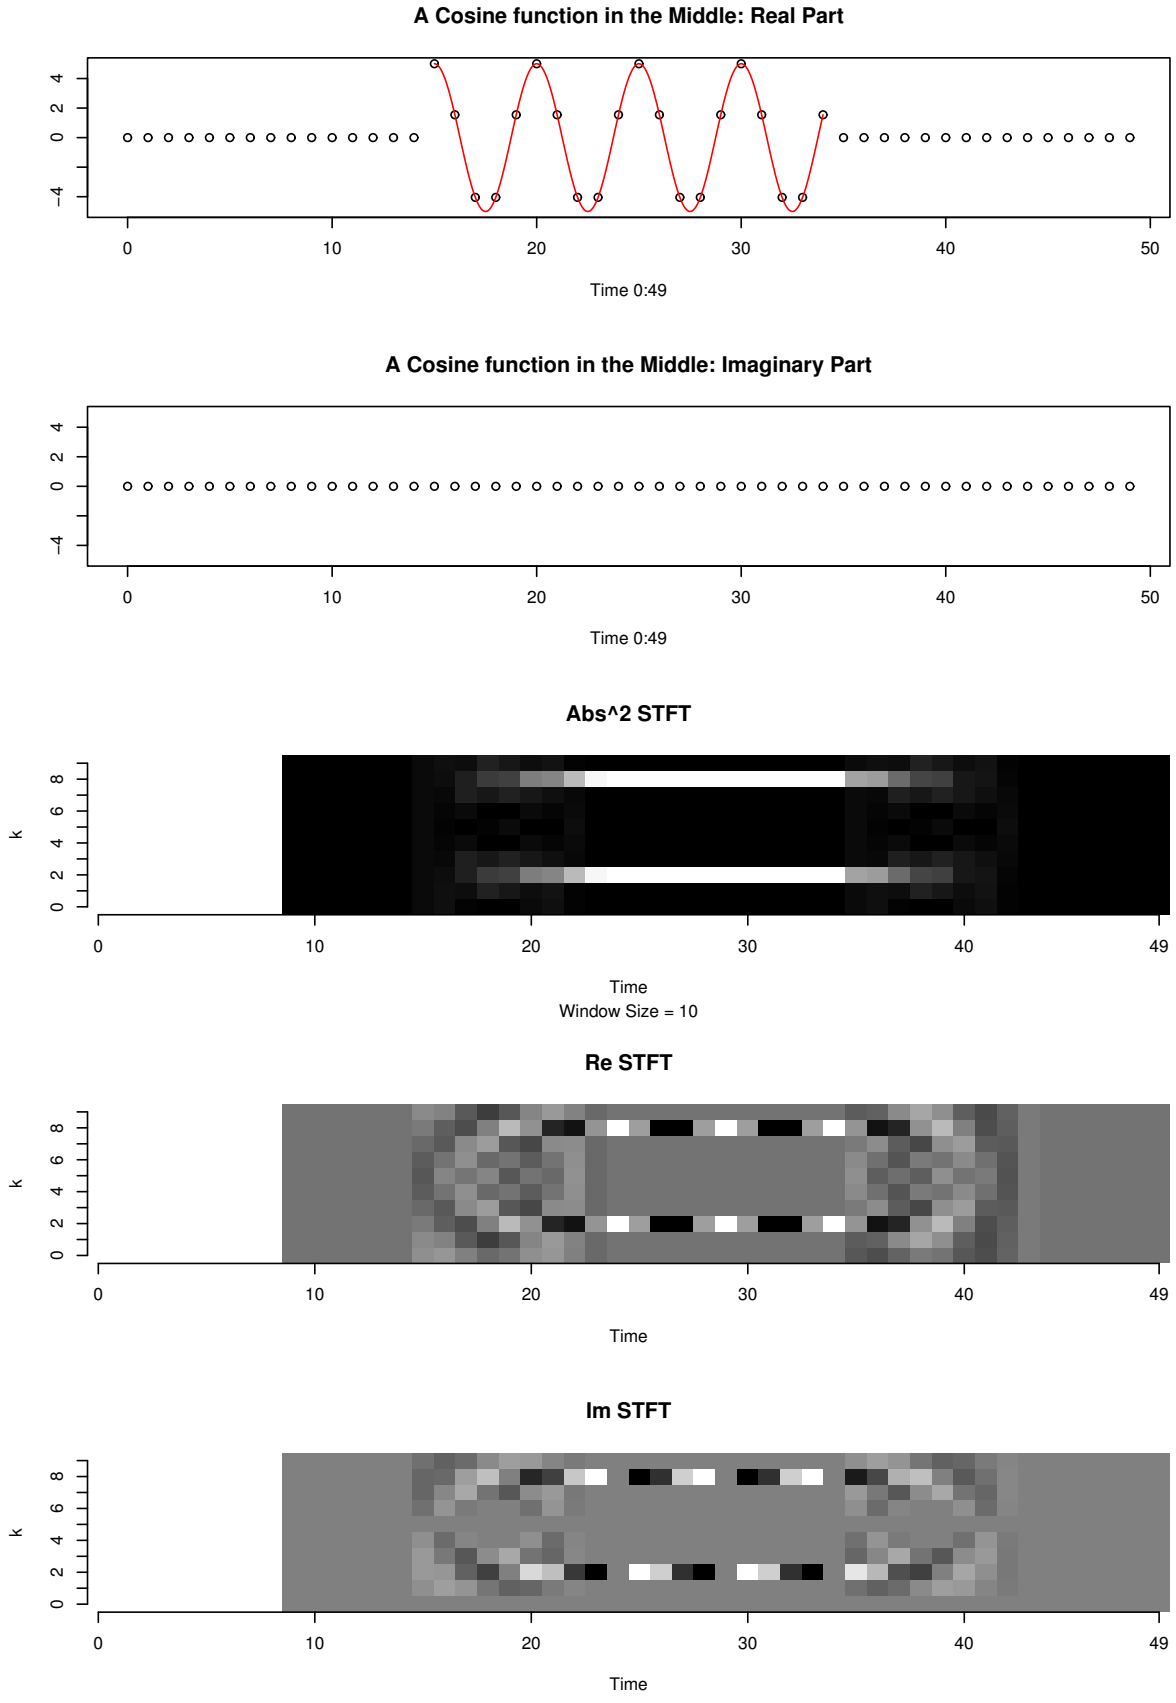

Figure 3: A simple example: there is a consine function in the middle and the squared modulus, real and imaginary parts of the resulting complex-valued STFT are shown.

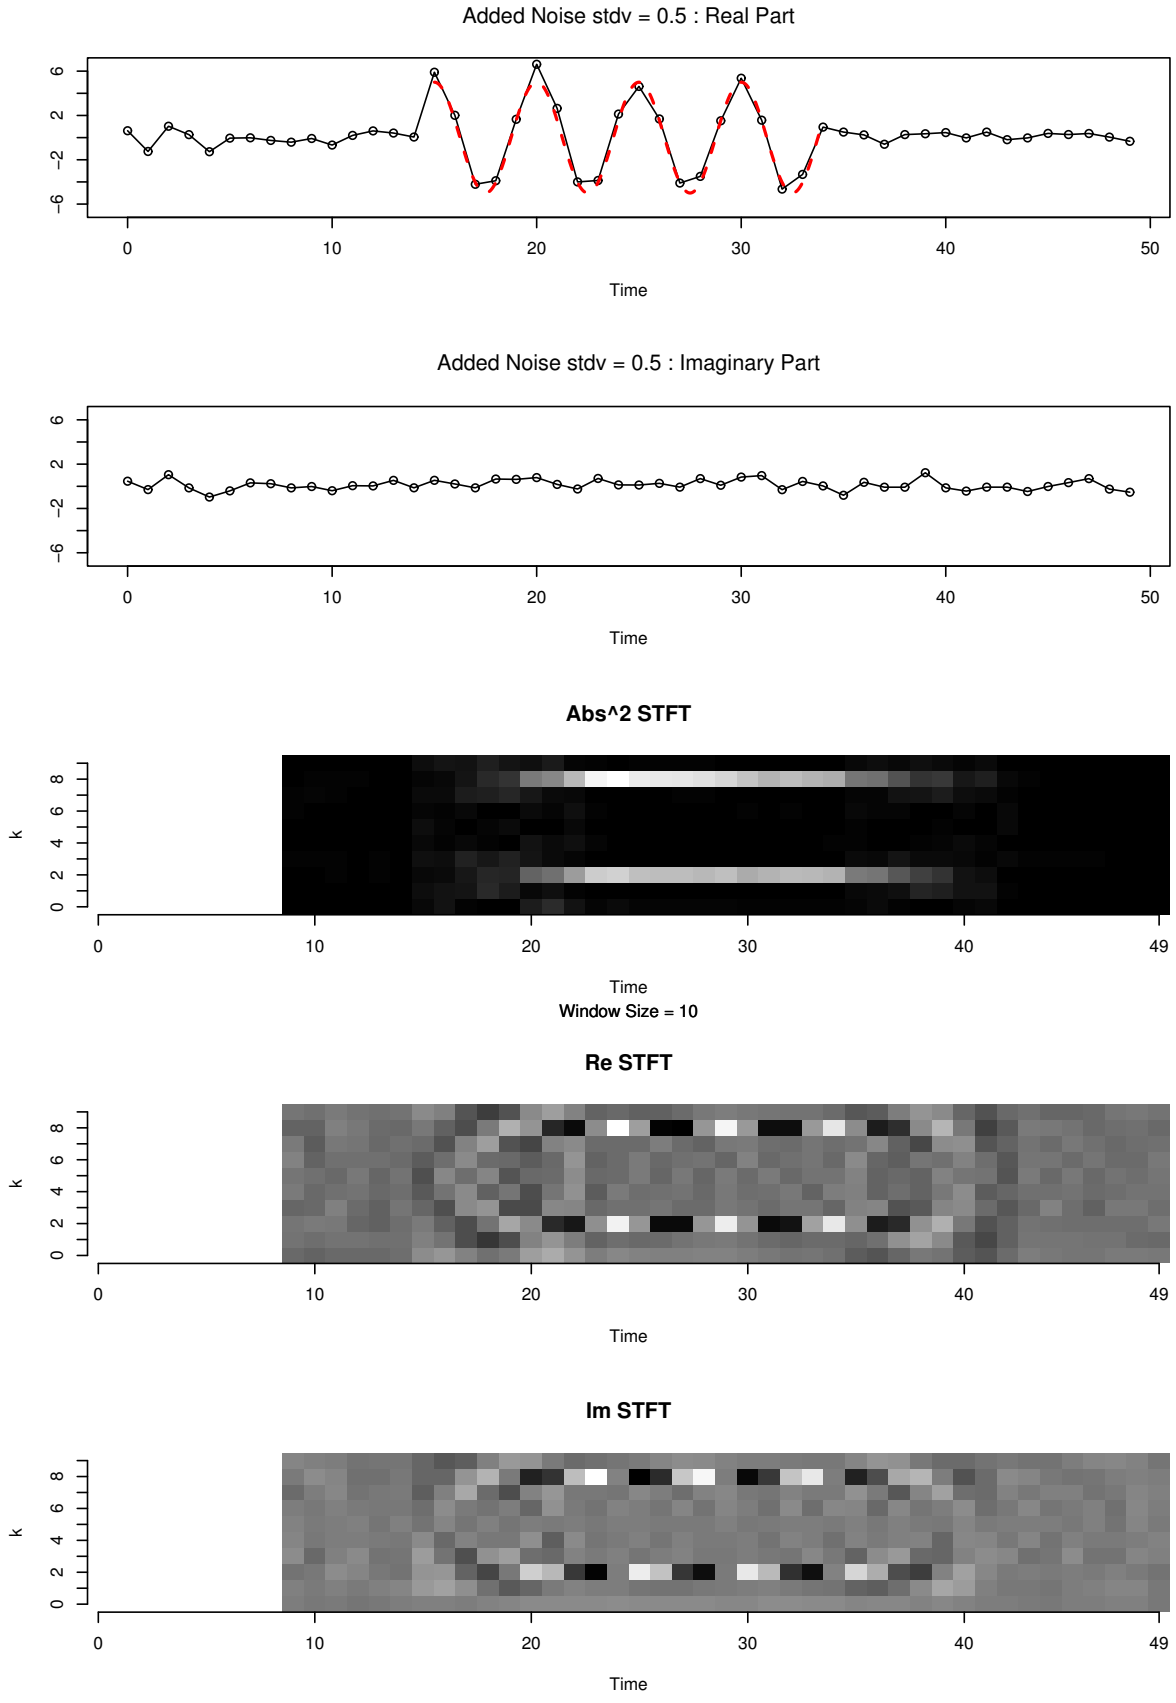

Figure 4: A local signal with a white noise series in background. As the noise level increases, the time series plot becomes noisier and the STFT less similar to Figure 3.

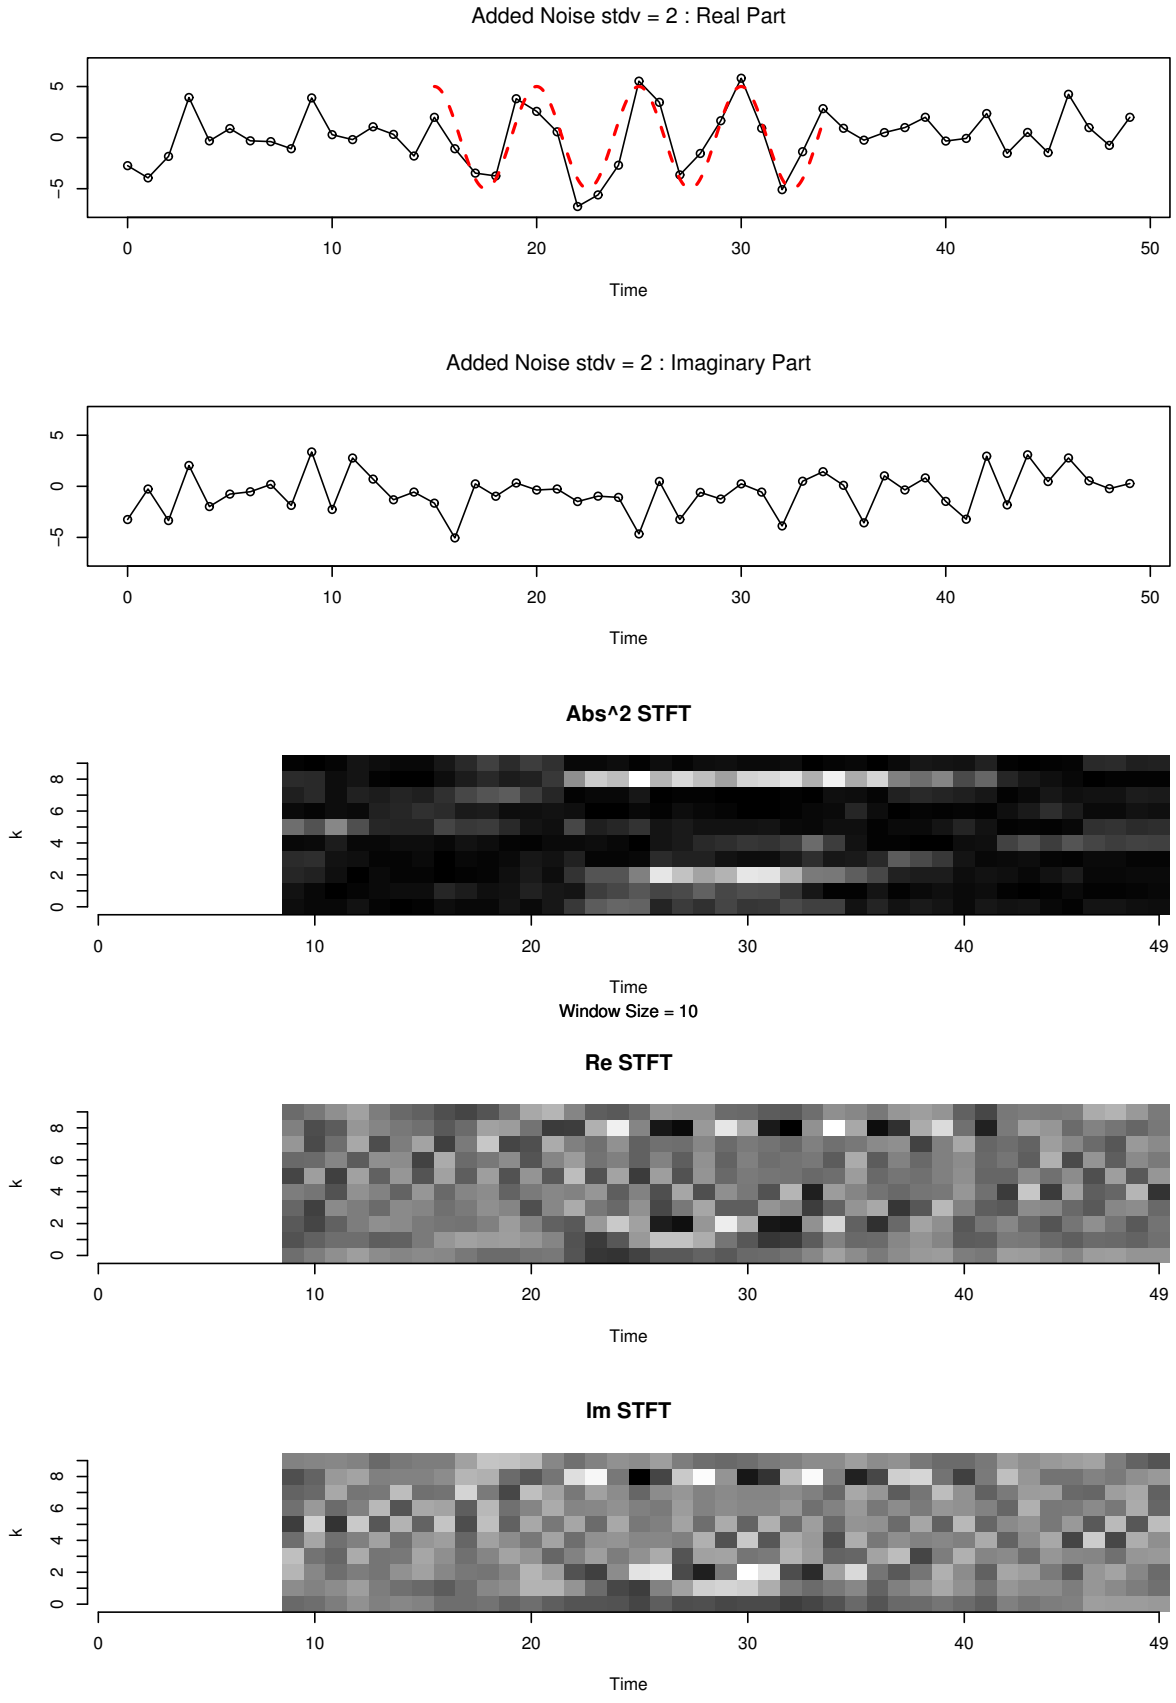

Figure 5: A local signal with a white noise series in background. As the noise level increases, the time series plot becomes noisier and the STFT less similar to Figure 3.

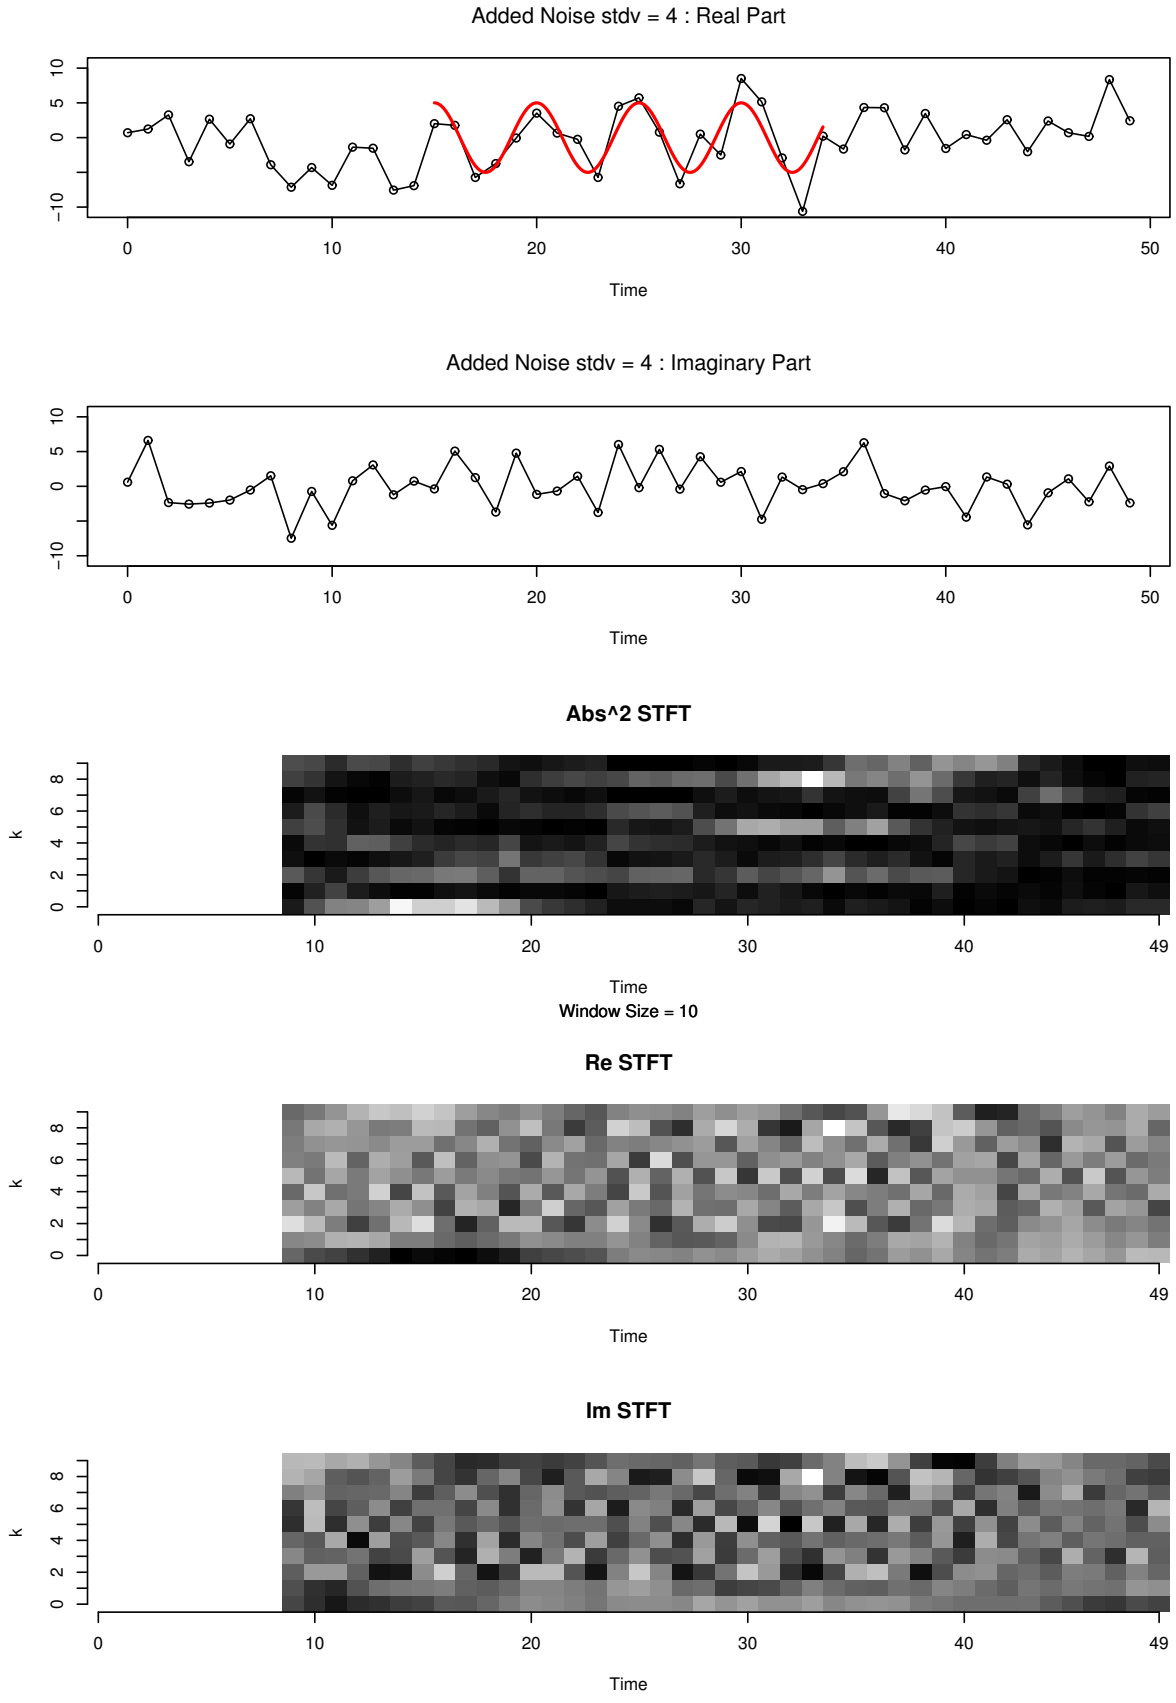

Figure 6: A local signal with a white noise series in background. As the noise level increases, the time series plot becomes noisier and the STFT less similar to Figure 3.

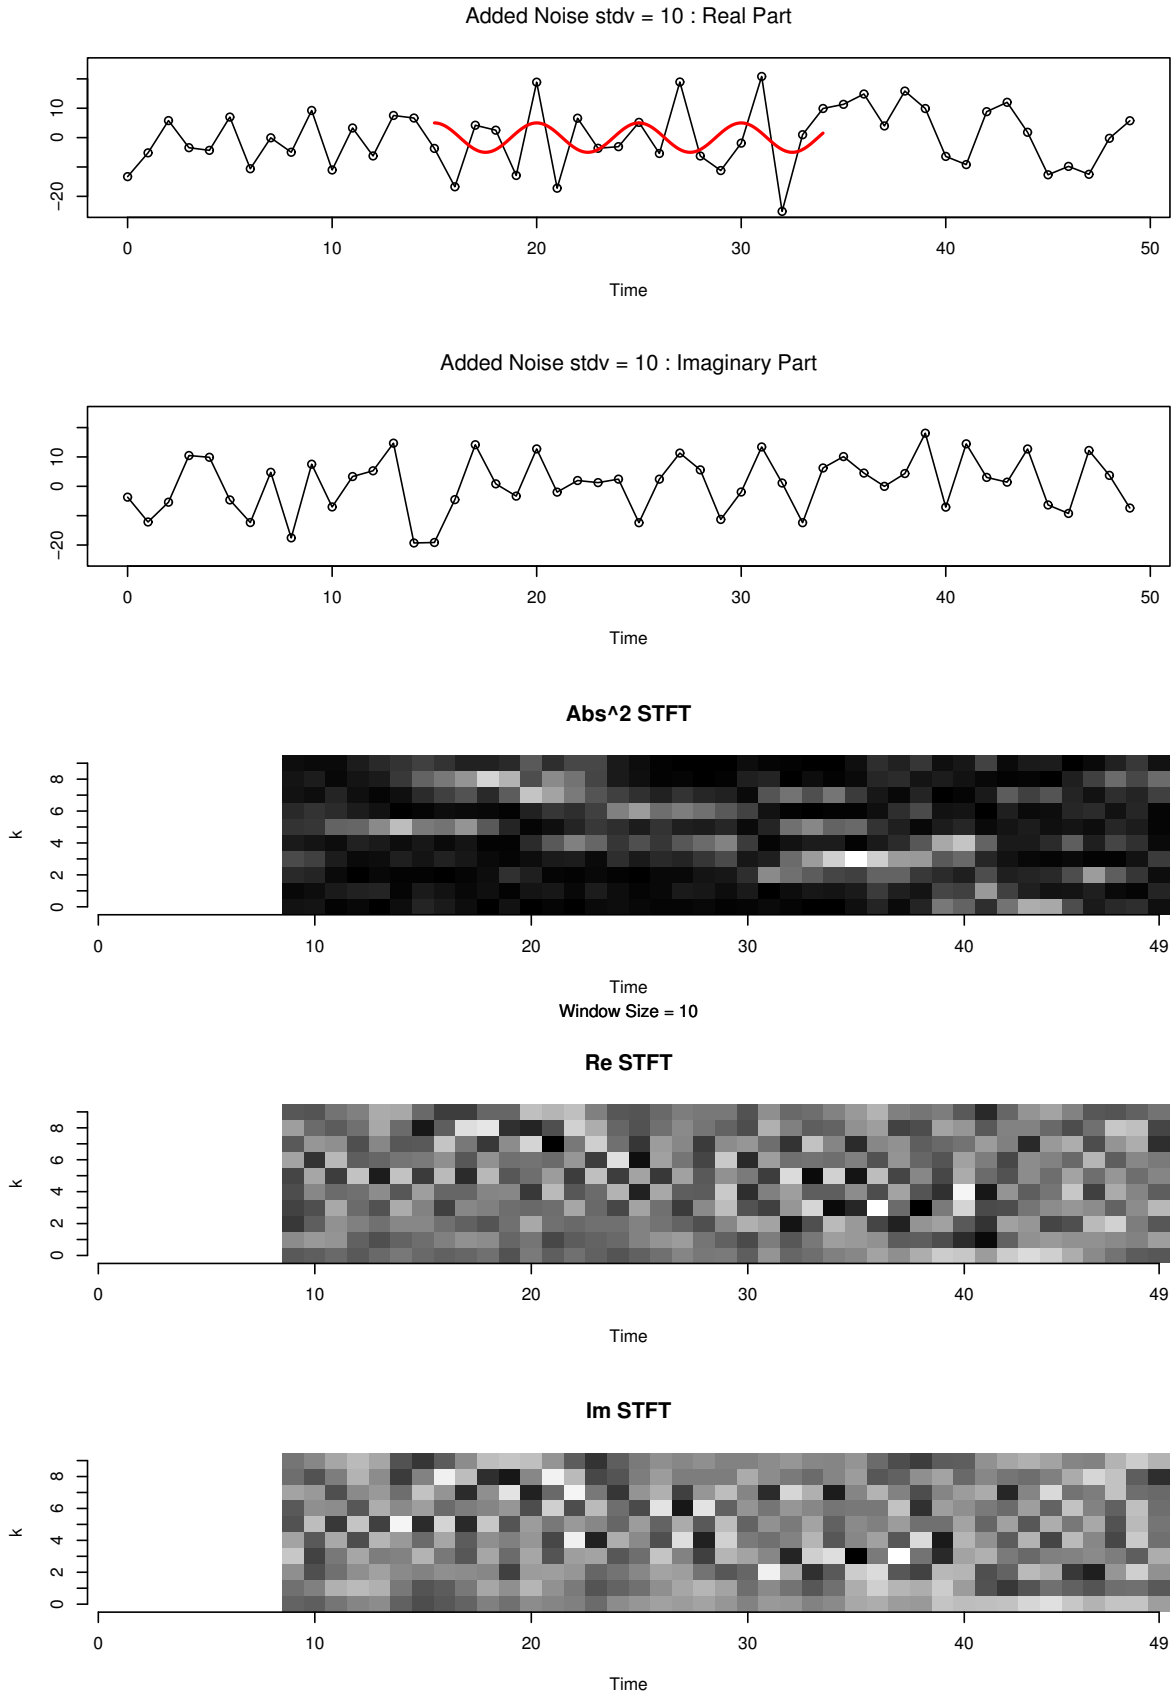

Figure 7: A local signal with a white noise series in background. As the noise level increases, the time series plot becomes noisier and the STFT less similar to Figure 3.