

Exploiting the Underlying Cepstral Coefficients for Large Scale and Fine-Tuned EKG Time-Imagery Analysis

Including R-R, P-R, R-T, and R-PVC Interval Imaging

James P. LaRue, PhD, *Jadco Signals*, Richard L. Tutwiler, PhD, *Professor Emeritus ARL / Penn State University/CEO LiveMotion3D LLC*, Dennison J. LaRue, *Senior, Biomedical Engineering, Clemson University*
James@JadcoSignals.com rlt1@psu.edu dlarue@g.clemson.edu

Abstract—This paper presents a three-part expansion for feature extraction of the electrocardiogram (EKG) based on a deeper exploitation of what cepstral processing can reveal, specifically in two distinct forms of time-varying-cepstral-power images. These two images have been reviewed for readability by qualified EKG technicians and have received positive reviews in light of their perception of having a tremendous gain in large-scale, yet fine-tuned, EKG strip interpretation ability. Our demonstration gives a brief review of EKGs in the time domain before moving to the image domain where we emphasize the pattern recognition aspects of our approach. Section II of this paper defines, and gives real biological applications for, the cepstrum and introduces our framing style for visualization. We give an example of a voice signal from one of the authors in this paper. Section III of this paper demonstrates our framing technique using two EKG records. Here, we lay the groundwork for imaging EKG as an image of EKG cepstral-feature terrain in the context of R-R, P-R, and R-T, as well as R-PVC intervals since this record contains premature ventricular contractions (PVCs) in addition to normal sinus rhythm (NSR). The first record (named record 105) was obtained from the Physionet MIT-BIH (MITDB); the second record was obtained from the Clemson University Biomedical Lab and is the EKG of another author of this paper. Section IV of this paper applies what we have learned from the first two parts to a variation of recurrence plots. In this instance, the recurrence plot itself is imbedded within a correlation matrix formed from a cepstral domain image. In this way, we point to our next stages of pattern recognition that involve yet another layer of exploitation of (dynamic) sub-matrices through 2-D segmentation and extraction which, when placed back in context of time domain, may further reveal clues to understanding the complex and seemingly acyclical nature of heart rate reset.

Index Terms—Autocorrelation Surface, Cepstral Coefficients, Cepstrum, Convolution, Electrocardiogram, Heart Rate, Normal Sinus Rhythm, Recurrence Plot, Shah Impulse Function, Three Dimensional Matrix Image of Cardiac Arrhythmia

I. INTRODUCTION

We hear that there is nothing new under the sun; and, this may certainly be true. However, with each passing day, the way in which the sun's new light captures the 'nothing new,' and reveals a fresh perspective about it, is. This paper takes the basic heartbeat and brings a fresh perspective to under-

standing it by incorporating the attributes of periodicity, periodicity within periodicity, and recurrence of periodicity by exploiting the tools of spectrum, cepstrum, and correlation in recurrence. The end result will be to offer a new perspective of how the modulation of beat, in this case of the heart, evolves in time by applying these few attributes to form images with recognizable patterns; and, how appropriate to do this in the setting of an AIPR Workshop. It may be said that these attributes we are utilizing are not new, and we agree, however, this paper will show that by taking these proven techniques in signal processing and applying them in a practical setting that being a bounded and discretely sampled time series, there is a tangible and new organizational intelligence that describes deviation patterns of the beat, or electrocardiogram (EKG), having been there all along, waiting to be exploited.

This paper is comprised of three parts. Section II will define the four main attributes we use to form the new EKG image analytic perspective: convolution, Discrete Fourier Transform (DFT), cepstrum, and recurrence. It will emphasize the role of strict adherence to properly measuring periodicity before meaningful analysis can take place. We will demonstrate these attributes in the classic context of a glottal response used with speech processing. Section III will introduce our techniques with a signal taken from the standard Physionet EKG MITDB data base. In addition to pointing out the commonly named parts of the EKG, in particular the P-QRS-T waveforms, we will show how to take a 1-D plot of an EKG time series, cut it up (justifying the size of the cut using the cepstrum), and then place the segments into contiguous columns of a matrix. It is the setting of a framed matrix, where we are able to take advantage of the 2-D array of pixels on the monitor, that our analysis of the EKG becomes a complete departure from current techniques that may involve a Pan Tompkins (that preceded Hamilton Tompkins) type of 1-D analysis. We will end section III with a look at what we call the cepstral terrain that, in part, was motivated once again by the context of the AIPR workshop. Section IV will then address our rendition of a recurrence matrix. Whereas most recurrence plots involve taking two subsequences of the same signal and plot them against each other (the first subsequence

on the x-axis against the second subsequence on the y-axis), our recurrence plots involve the matrix of framed contiguous segments of EKG we will describe in section III. In essence, we are developing recurrence objects by exploiting the fact that an EKG signal is not as stationary in heart rate (HR) as expected. Hence, our visualization of how the EKG evolves, which is most easily seen by the movement of the QRS complex across the columns, is a time-relativistic observation of how the EKG communicates (possibly bi-directionally). The autocorrelation surface of this matrix brings out the underlying communication in a localized view along the diagonal as well as a global view off the diagonal. The resulting image resembles a traditional single signal recurrence plot, but, in our case, there is an order of magnitude more information available, because we use multiple signals simultaneously. It is interesting to point out that our particular MITDB signal is described clinically in Physionet as having Normal Sinus Rhythm (NSR) and, while this paper does not in any way attempt to rebut clinical diagnosis, our novel framing and recurrence techniques may at least offer that fresh perspective we mentioned in the opening paragraph, to bring about a deeper insight.

In addition to forming a matrix of recurrence, we will form an image of cepstral coefficients that we refer to as the EKG cepstral terrain. It is completely driven by cepstral coefficients derived by taking short-time windows of the EKG in time domain; thus, this image is itself a time domain rendering and, in fact, is the driving source for how we choose to frame the data. We will be able to show the evolution of these coefficients through the power of a Matlab generated matrix image. Furthermore, we will point out in this EKG terrain how the P and T waves are coupled to the QRS complexes, thus showing a measure of self-similarity of the heartbeat.

In review, our main figures consist of a two-piece matrix system. The first matrix piece is a reframing of the time series where a specific and cepstral-motivated number of discrete samples is chosen followed by chopping the EKG into consecutive pieces of that number duration and then placing each piece in consecutive columns of a matrix. The second matrix piece is the autocorrelation surface of the first matrix piece; this is our rendition of a recurrence plot. It is from this second image matrix that a K-means clustering algorithm is utilized to identify and classify short-term patterns in the P-QRS-T coupling. Then, we will set these images in motion according to a series of ordered cepstral-based reframing sizes in the context of a physical time relativistic setting, i.e., keeping control of the moving local stationarity of periodicity of the EKG process.

II. OUR BASE OF FOUR ATTRIBUTES

We will start out by defining our base of four attributes of time domain analysis: convolution, DFT, cepstrum, and recurrence. We will start with the last concept of recurrence. Recurrence, for which we may think of as a dynamical repetition of state, does not necessarily emphasize the time it took to

repeat, but, rather, that it did. To build a recurrence plot, one starts with a given scalar time series $x(1), x(2), x(3), \dots$ from which we may construct an initial subsequence $x(n), x(n+L), x(n+2 \cdot L), x(n+3 \cdot L), \dots x(n+M \cdot L)$ where n is the discrete time sample, L is the time expanse between samples, and M is the chosen length for this subsequence. Then, by shifting n to the right (i.e., let n move to $n+1$ in each member of the initial subsequence), up to $L-1$, more subsequences may be formed. A 2-D recurrence plot (RP) would take any two of the subsequences and form a matrix of M^2 ordered pairs of points by taking an outer product of the two sequences. Credit goes to Eckmann, Kamphorst, and Ruelle [1] for introducing the concept of recurrence plotting where it is recognized that an RP is “a two-dimensional representation of a single trajectory.” For examples of this type of constructed RP, we also refer the reader to [2].

Now, in contrast to the above description for a single trajectory sourced RP, our RP is a study of recurrence between a series of (non-sub-sequenced) trajectories. The three main differences are: (1) Each trajectory is made of contiguous samples thereby eliminating the time expanse L ; (2) The length of a trajectory is no longer represented by the letter M , which may be chosen rather arbitrarily; instead, we represent trajectory (sequence) length by the letter P , which stands for the more deterministic concept of ‘periodicity;’ and (3) our resulting RP is an autocorrelation surface based on a matrix of individual trajectories framed in columns. For example, given a times series $\{x\}$ and length P (which we will show is derived by using the cepstrum), we form the first column of a matrix with the contiguous sequence of P points $x(1), x(2), \dots x(P)$. The second column is then formed with $x(P+1), x(P+2), \dots x(2 \cdot P)$. Additional columns may be added accordingly, thus framing the sequence $\{x\}$ into a P by N matrix where N is the number of contiguous frames that can be taken from $\{x\}$. Once the matrix of sequences is formed, we simply use the correlation coefficient command on pairs of columns to form the $N \times N$ autocorrelation surface matrix.

But, it must be emphasized that the surface matrix is a dynamic object as it is a function of P . For the physical example shown in this paper, namely the EKG, P is related to the very observable (but quasi-periodic) R-R intervals drawn from the QRS complexes in the EKG. It is through the remaining concepts of convolution, DFT, and cepstrum that the production of the (variable) P is realized.

Our next attribute to discuss is convolution, for which we think of as the measure in time of the reaction between two physical processes passing through each other like that of an impulse and a physical reaction to that impulse. We define the discrete convolution between two functions $x[n]$ and $h[n]$ as:

$$y[n] = x[n] * h[n] = \sum_{k=1}^T x[k] \cdot h[n - k] \quad (1)$$

where the indices $k=1:T$ represent the time the two processes of $\{x\}$ and $\{h\}$ interact.

Convolution is closely related to the next attribute of DFT, which is a measure of sinusoidal content for various frequencies in a (periodic) time domain process such as $y[n]$. A typical written formulation may be given as:

$$Y[f] = \sum_{m=1}^P y[m] \cdot \left(\cos\left(\frac{2\pi mf}{P}\right) - i \cdot \sin\left(\frac{2\pi mf}{P}\right) \right) \quad (2)$$

where Y represents the frequency transform of y for different integral values of (frequency) f , and P is the periodicity of discrete samples of $y[n]$, like that of an R-R interval from the EKG with m ranging from 1 to N .

Thus, this DFT is actually (and subtly) derived from a single sequence that is P elements long and herein lies the problem: the determination of P for the practical case of when the sequence $\{y\}$ is not a realization of pure sinusoids that necessarily are periodic in the sense of the least common multiple of the underlying frequencies f_i 's; anything else is at best an approximation. As my professor Dr. George Ioup, from the University of New Orleans, stated over and over and over again, "It comes down to the calculation of 'df' where df = sample rate/number of samples;" a seemingly simple calculation that is sometimes taken way too much for granted. 'df' is the driving factor in defining what a pure sinusoid is in the context of taking a DFT without ending up with side-lobes surrounding each instance of a particular f_i . In a practical setting where there may only be a single sampling rate available, the estimation of 'df' is then dependent on the number of samples in a period of signal that we derive easily but yet only approximately using the cepstrum.

The cepstrum, our staple attribute, is a measure of periodicity and is where we derive P . In the seminal paper by Bogert, Healy, and Tukey written in 1963 [3], the concept of cepstrum was introduced (also see [4] and [5]). The cepstrum was picked up by the domain of Speech Processing as it was a way to identify and classify the glottal response in voice. The glottal response is comprised of a basic waveform that is repeated in a time series [6]. Fig. 1 illustrates the glottal response with an author's solfege of 'Do' (as in Do Re Mi Fa) along with how it may be realized as a basic waveform that it convolved through a train of impulses spaced apart by the length of the waveform. To paraphrase Akay [7], "Cepstrum analysis is concerned with the deconvolution of two signal types: the fundamental (basic) waveform and a train of impulses (excitation function)."

It is at this point that we focus our attention to the signal processing of the train of impulses, as this is the source for amplifying certain cepstral coefficients that in turn help us to decide on a P value. Let us identify $\{x\}$ as the basic waveform sequence and $\{h\}$ as impulse train. Utilizing (1) we see that $\{y\}$ is written as the convolution of these two functions. In (2), the DFT of $\{y\}$ is given, but due to the remarkable convolution theorem for DFT [8], the DFT of $\{y\}$ can be written as the product of the separate DFT's of $\{x\}$ and $\{h\}$, that is, in the nomenclature of (1) and (2), the time-frequency pair




Fig. 1. We demonstrate the glottal idiolect response to 'Do.' Fig. 1A shows 2350 points of 'Do' sampled at 22050 samples/sec. There is a clear 10x repetition of a base waveform. Fig. 1B shows the overlaying of the first four base waveforms in 'Do.' Fig. 1C shows a train of impulses separated by $2350/10 = 235$ points. The 235 would be represented by 'a' in the usual definition of sampling function shown later in (5). Thus, we are showing that the convolution of 1B and 1C results in the waveform shown in 1A.

for $\{y\}$ and Y is given by:

$$y[n] = x[n] * h[n] \leftrightarrow X[f] \cdot H[f] \leftrightarrow Y = X \cdot H \quad (3)$$

The impulse function is also known as either the sampling function or the comb function, reasons for which should be obvious from Fig. 1C. A usual way to write the sampling function is by using the symbol for Shah, III , and is given by:

$$\text{III}(m) = \sum_{n=-\infty}^{\infty} \delta(m - n) \quad (4)$$

The Time/DFT pair for the sampling function is given by:

$$\text{III}\left(\frac{m}{a}\right) \leftrightarrow a \cdot \text{III}(a \cdot f), \quad (5)$$

where 'a' is a parameter that forces zeros in both the time and frequency domain outputs when m is not an integral multiple of 'a.' Thus, in the context of what has been written so far:

$$x[m] * h[m] = x[m] * \text{III}[m/a] \text{ transforms via the DFT to } Y(f) = X(f) \cdot a \cdot \text{III}(a \cdot f). \quad (6)$$

At this point, the DFT transform is in terms of complex numbers and so we now form the spectrum by taking the magnitude of the DFT product. At this point, we are half-way to deriving the cepstrum. To finish the cepstrum, as described historically, one simply takes the inverse DFT of this magnitude. In practice, we simply take the spectrum of the spectrum. In any event, it has been found that evidence of the underlying length between impulses appears in the cepstrum domain and we can now reveal that the constant 'a' in (5) is the very P we are looking for –the the diamond in the rough– so to speak.

This appearance of the P constant is due to the influence of the impulse response and the reflexive nature of (5) in passing the constant 'a' (or P) back and forth. Without doubt, there

are more in depth theoretical justifications for this observation. Our own somewhat naïve observation is that the magnitude of the spectrum forces the ‘harmonics’ of the Shah function to appear upright and so separate themselves from the wide-band response of the basic waveform, and so, by inducing a second spectrum (on this spectrum), the inverse P distance relationship is exploited more prominently than the effects of the waveform; but, all we are looking for is a number P to take back to the time domain in order to frame the time series $\{y\}$ into the $P \times N$ matrix we described above. Fig. 2 shows the evolution from the time domain to the DFT spectrum and then to the second order DFT cepstrum. Fig. 3 is a zooming in of Fig. 2 to emphasize the features associated with the cepstrum. It should be noted that the x-axis for each feature has been left in the sample number domain as we find it more helpful, if not more intuitive, to observe the spectrum as a time series itself.




Fig. 2. This shows how to arrive at the number 235, which expresses the timed periodicity of the arrival of waveforms. The top plot shows the time domain of a capture of ‘Do.’ The second plot shows the magnitude, or spectrum, of the DFT of the time series. Clearly, there is some evidence of a comb function on either end. And finally, the third plot is the magnitude of the DFT (could just as well have taken the inverse DFT) of the spectrum. Again, there is evidence of a comb but the ‘bristles’ are rounded, an effect due to the interaction between the glottal impulse and its vocal tract response.




Fig. 3. This is a zooming in of Fig. 2 to show the tight combing effect in Fig. 2 on the left side of the spectrum (the same response appears on the right as well) and the capture of the time between glottal impulses, or, the duration of the vocal tract response, that being 235. The number 235 is our P value we are looking for in order to begin our framing of the $P \times N$ matrix. For the remainder of this paper, and starting in section III, we will study the cepstral response for the P-QRS-T waveforms of the EKG.

In summary, section II has shown the relationship between convolution, DFT, cepstrum, and recurrence. Our example using voice, and our upcoming examples using EKG, depends on the realization that there is a unique waveform that is repeated contiguously and that the cepstrum can pull these features apart, so to speak, in order to exploit the length of the base waveform, i.e., the underlying periodicity of the time series. This length parameter is then used to form a matrix of contiguous intervals from the time series, and thus we have moved from time (initially without knowledge of periodicity) to spectrum, to cepstrum, and then back to time with an understanding of periodicity. It is this estimated periodicity that drives how the $P \times N$ matrix is formed. In section III, we will demonstrate the matrix formulation when we turn our attention to the EKG. Section IV will then study the recurrence relationship among the columns.

III. FRAMING A MATRIX FROM AN EKG TIME DOMAIN SERIES ACCORDING TO CEPSTRUM

In the context of modeling the EKG as a convolution of two functions, we may consider the sinoatrial impulse as one function and the system response to the atria and ventricle pathways as the other. And so, the observation by Bogart, Healey, and Tukey was that the impulse rate initially masked by the heart’s physical reaction to it is actually revealed because the logarithm separates the two processes in the frequency domain (recall that convolution in the time domain is multiplication in the (DFT) frequency domain) before returning them to the time domain, albeit a different type of time domain known as the cepstral domain. In the cepstral domain, the time axis tells us how many samples make up the periodic components of the original sequence. The y-axis is a measure of strength of the periodicity. And although these x-y measures, for our case of the EKG, can readily point to heart rate as being how many samples are likely to occur between heartbeats, there is more information within these plots due to the interaction of the pulse evolution throughout the P-QRS-T circuitry from atrial to ventricle.

To demonstrate the framing of a matrix, we will start by introducing our first EKG example. It comes from the MIT-BIH Arrhythmia Data Base (mitdb) found at <https://physionet.org/cgi-bin/atm/ATM>. The signal is labeled as MITDB 105. It was recorded from a 73-year-old female who had periods of NSR and Pre-ventricular Contractions (PVC’s) [9]. In Fig. 4A, we have plotted 5 sec x 360 samples/sec = 1800 samples of 105’s EKG. The basic waveform is composed of a P wave, QRS complex, and T wave. These have been marked for reference. The time series shown is a contiguous collection of this basic waveform; although, it should be pointed out, that each basic waveform is somewhat, but not exactly identical, to the others. A set of X-Y labels have been posted for ease in approximating the number of samples (our soon to be P) between QRS complexes and P and T components. It can be observed that there are about 260 and 250 samples between the first two corresponding sets of QRS

complexes. Fig. 4B shows the plot of the cepstrum with a main peak evident at 260. Leaning on what we learned in section II, we can think of the top plot as a convolution between seven impulses, separated by a distance of approximately 260, with one of the QRS complexes we see. Note that we skipped posting the DFT transform as it has little relevance for the remainder of this paper, only that we need it to access the cepstrum. What we will show next is how we can frame these seven EKG complexes into a matrix.




Fig.4. Fig. 4A shows 1800 samples of an EKG. XY labels are attached to physically show the distances between P, QRS, and T components. Fig. 4B is a cepstrum output showing high amplitude at the 260 mark. This coincides with the distances between QRS complexes as labeled in 4A. Note that the 260 is built with P-P, QRS-QRS, and T-T distances since they are all coupled together.

In Fig. 5, we list out the five steps to framing the EKG time series in Fig. 4A using the highest amplitude number 260 indicated in Fig. 4B. The 1st step is to take the first 260 points of the time series. The 2nd step is to move from an x-y plot to a y-x plot where sample number is listed on the new vertical axis (where we had to add the negative sign to trick Matlab) and the amplitude is now shown on the new horizontal axis. The 3rd plot is a matrix representation of the 2nd plot, albeit a one column matrix representation. The colorbar indicates the amplitude. The 4th plot is a replotting of the time series to be framed. It may be revealed at this time that the number of samples in this times series was purposely chosen to be an integral multiple of 260; more precisely $7 \times 260 = 1820$ samples. The 5th plot then is derived by taking the 7 contiguous sets of 260 points to form a 260×7 matrix where the top point of each column is the first point of each contiguous set. This is an example of framing the $P \times N$ matrix described in section II.

In Fig. 6A, we have plotted 200,000 samples from MITDB 105, which equates to a little more than nine minutes; it is hard to make out anything intelligible. Fig. 6B shows what 5 seconds of this EKG look like. In Fig. 6C, we have taken a three minute chunk of EKG, as indicated by the box in Fig. 6A,




Fig. 5. This is a five-step illustration of how we take a time domain series and produce an image matrix for a different view to better enable pattern recognition in an EKG.

and constructed 200 frames based on $P=260$. The QRSs under the NSR identification are the yellow pixels in each column. The seven red pixels are the PVCs that correspond to the seven spikes in the box. The advantage of looking at signals in this framing format is that there is a more efficient use of pixels on the monitor. There are two dimensions of time (column by row) and a third dimension of color to show wave intensity. Altogether, these dimensions yield yet another dimension of visual evolution of the NSR and PVC. We can readily observe that the NSR is producing a sinusoidal pattern and that the PVCs are followed by compensatory pauses.




Fig. 6. Fig. 6A is a 9.25 minute EKG recording stuffed on a 1-D plot. The x-axis ranges to 200,000 samples based on the 360 samples/sec sample rate. 6B shows a few EKG pulses. 6C is an extension of what we showed in the 5th step of Fig. 5 and shows the agility and power of spreading the samples in two dimensions to reveal an intelligible insight to the EKG, especially with the cyclic QRS patterns (shown in yellow in each column) and the PVCs interrupting the NSR.

There are two subtle observations to be pointed out. The first is that the T to P wave transitions are evident in the lighter blue sections that measure nearly 100 samples in each column. Notice how the waves are coupled to the QRS complexes. The second observation is that starting with the 140th column, the QRSs start to reappear in the columns at earlier times. This means that the framing at $P=260$ will need to be adjusted as the heart rate is changing.

If one was just concerned with heart rate, it may be enough to simply use the 1D R-R interval estimation given by a Pan-Tompkins type algorithm [10]. However, we believe there is an advantage in taking contiguous sets of short time cepstrum transforms and arranging them in yet another matrix form. To produce this new type of matrix, we took 125 contiguous snapshots of 1000 points across the first half of the data shown in Fig. 6A. The first snapshot is shown in Fig. 7A. This snapshot forms the first row of the matrix in Fig. 7B where we have kept 400 coefficients for each row. Looking down the rows, it is observed that the main cepstral components are in the 250-270 range and that there is a cyclical nature to them as we would expect by what was observed in Fig. 6. The effects of the PVC timing show up symmetrically at the 150 and 350 marks. (These are not PVCs; they are just the effect they have on the cepstral components). A subtle observation we shall now point out is the vertical column effects shown at the 50 and 100 marks. Although the attention is in general placed on the R-R peaks (plus the wrongly overlooked P-P and T-T peaks) as being responsible for the comb function producing the cepstral peaks around 260, there are other comb effects to consider, namely, the distances between the P and QRS, the QRS and T, and T and P wave components. The columns at 50 and 100 appear to reflect the coupling between these components and it is this (subtle) coupling that causes the rectangular appearance in deference to the shifting evident at the main cepstral components. A 3-D image of Fig. 7 is rendered in Fig. 8. We call this the Cepstral Terrain.




Fig. 7. Fig. 7A shows a cepstral output of the first window of 1000 points of an EKG. This is placed in row one of the matrix in Fig. 7B. The remaining 120+ rows in Fig. 7B come from taking the cepstrum of contiguous sets of 1000 pt. windows.




Fig. 8. What we refer to as a Cepstral Terrain based on its appearance. It is the 3-D view of Fig. 7B.

IV. RECURRENCE PLOT AND AUTOCORRELATION SURFACE

In section II of this paper, we began with a discussion of recurrence plotting. We differentiated ourselves from traditional RP's in that we are not binding ourselves to a single trajectory but, instead, we are using a series of trajectories simultaneously. In section III, we constructed a framework for $P \times N$ matrices where P was the periodicity of interest and its value was taken from cepstral components. In this section IV, we will define our RP as the autocorrelation surface of the $P \times N$ matrix where the N columns are used to find correlation coefficients to produce a symmetric $N \times N$ matrix such that the first row contains the coefficients of the first column correlated against itself (on the first diagonal entry) and the remaining $N-1$ columns, the second row contains the coefficients of the second column correlated against itself (on the second diagonal entry) and the remaining $N-1$ columns, and so on. We use the Matlab command 'corrcoef' on the $P \times N$ matrix. The reason RPs are utilized in this way is to exploit the subtle patterns in the $P \times N$ matrix and present them more clearly. We will present several examples with the $P \times N$ and $N \times N$ matrices side by side as they are different views of the same numbers.

The data we will use comes from same MITDB 105 and from an author of this paper utilizing the Clemson University Bioinstrumentation Lab. We will be constructing several matrices based on various P values. As we have discussed previously, P is related to heart rate and the heart rates for either subject are not stationary. Hence, the matrix images formed for our pattern recognition algorithm are relative to selected periodicity and evolution of the complexes in a neighborhood of the periodicity. For a more abstract approach see [11].

We begin with Fig. 9A, where we took all 200,000 points, or 9.26 minutes, of EKG from Fig. 6A. Once again, we framed by $P=260$. As a matter of reference, this P corresponds to 83.07 BPM ($360/260 \cdot 60$). Notice there are over 700 columns in our $P \times N$ matrix. Also, notice that there is a loss of resolution compared to what we saw in Fig. 6A and, in partic-

ular, the PVCs are not as evident. However, the good news is that in Fig. 9B, which is this paper's first example of an NxN autocorrelation surface RP, there are stand-out patterns along the diagonal due to NSR. On the diagonal of Fig. 9B, we have drawn five boxes of color and these directly correspond to the boxes in Fig. 9A. An interpretation of the localized EKG can be given best by looking at the RP. The red, yellow, and dark blue boxes (1st, 4th, and 5th) appear in solid rectangular form and may be interpreted as segments of a localized stationary heart rate. The purple and light blue boxes (2nd and 3rd) appear as X's and O's and may be interpreted as a heart rate that is changing in a relatively small neighborhood of time from 'slow to fast to slow' (or opposite) akin to a curve of a cubic polynomial. The point is that we are looking at over 9 minutes of contiguous EKG and learning to identify coherent structures of interest that we believe have not been demonstrated before and would seem to be challenging at best to draw from a traditional R-R interval plot. The three dashed dark blue boxes indicate a strong measure of cross-correlation between themselves and the 5th box on the diagonal. Back in section II, we mentioned a reference to a time expanse parameter labeled L. These off-diagonal components of the RP can now be reinterpreted as an expanse parameter where in this case L is used here to indicate the distance (in time) between similar complexes that are stationary across the rows of the PxN matrix.




Fig. 9. Fig. 9A is the framing matrix of the time domain series described in section III. Fig. 9B is the autocorrelation surface formed from Fig. 9A. The solid colored rectangles are given in pairs where the time domain structures incased in Fig. 9A are more pronounced along the diagonal in Fig. 9B. A heart rate that is stationary from a period of time will appear box-like in Fig. 9B, e.g., the red, yellow, and dark blue boxes. A heart rate that changes as a cubic polynomial forms what we call X's and O's, e.g., the purple and light blue boxes. The dashed boxes are indicators of off-diagonal correlation.

To demonstrate the sensitivity in choosing P, we now turn to Fig. 10. Figs. 10A and 10B are based on the same signal that formed Fig. 9, except that the PxN matrix is now framed at P=259, which yields an equivalent measure of 83.3 BPM. In Fig. 10B, we have imposed a binary threshold to bring out the contrast of the structures. Clearly, the structures have changed in both matrices with this very small relative change in frame time. (One may wonder if these relativistic changes are actually a form of parallel biological communications at the molecular level [12]). Once again, the solid boxes and X's and O's on the diagonal of the RP give a quick visual understanding and overview of how the QRS complexes are aligning in the PxN plot.




Fig. 10. This pair of plots uses the same data as that used to form Fig. 9; the only difference being the framing dimension change from 260 to 259. This indicates that timing is sensitive and structures are relative to framing dimension.

Our final data set comes from a 3-minute self-recorded EKG provided by the student author listed in this paper. The experiment was conducted at the Clemson Bioinstrumentation Lab. In this case, the sampling rate was set at 400 samples per sec. The experiment was carried out by intentionally breathing slow in the beginning before changing abruptly to breathing fast at the second half of the recording. The idea was to reproduce a cubic polynomial pattern, which we will show in Fig. 12, but as lagniappe (a little extra) we point out that our way of visualizing the EKG exposed a series of odd, but deterministic, set of QRS pulses.

In Fig. 11, we show the instrumentation used in the experiment along with the EKG time series, a few pulses, and two main modes of cepstral components corresponding to the two extreme breathing rates. A close examination of the raw times series in Fig. 11A reveals wide spacing between QRS complexes in the first half compared to the spacing in the second half. That the cepstral plot in Fig. 11C exhibits two modes is further motivation for producing the plot cepstral time plots in Figs. 7B and 8.




Fig. 11. 11A shows the equipment set-up. 11B shows 70,000 samples of college-level EKG. 11C shows a few pulses with some high amplitude QRS complexes. 11D shows a bi-modal cepstral component distribution of high amplitudes at 350 and 450. The other modes are evidence of coupling between the P and T waves.

Fig. 12A shows the PxN framing of Fig. 11A with $P = 454$. This is the higher of the two main Cepstral components in Fig. 11C, where we note that the larger components relates to more time between QRS complexes and hence the slower heart rate in contrast to the peak indicated at 350. The RP autocorrelation surface is in Fig. 12B. The yellow and green boxes point out a long and short set of cubic polynomial complexes thus corroborating with the results shown from the MITDB 105 plots.




Fig. 12. 12A shows the time series frames at 52.8 BPM where the coherence of a slow rate matches the how the experiment was conducted. 12B is labeled as a Cross Frame Correlation. It is the same as autocorrelation surface RP. The yellow box indicates the cubic polynomial pattern in time and recurrence.

Fig. 13 is our final picture involving EKGs. This particular framing was detected as we were running through several P values contiguously. The structures outlined in the yellow boxes were first detected on the autocorrelation surface RP (Fig. 13B) that commanded our attention back to the PxN framing matrix (Fig. 13A). The group set of pulses were then magnified to form the inset Fig. 13C. Fig. 13C makes obvious the presence of three nearly identical sets of quadruplets and, while we are not in any way presupposing to be in a position of clinical assessment as to what these quadruplets may mean, we believe that traditional R-R methods would not have identified such deterministic looking structure in the manner in which these framing plots do.




Fig. 13. In going through a series of framing dimensions, we uncovered a series of quadruplets. Whether or not this would have clinical meaning is secondary to understanding that such structures would hardly be evident in the usual 1-D R-R intervals that are usually given in clinical level evaluations.

V. NEXT STEPS

We have a feature identification algorithm in place for detecting the existence and sizing of the solid rectangles and the X's and O's complexes along the diagonal. The clustering algorithm is outlined in Fig. 14. Currently, we are working with work load restrictions as we expand the search off diagonal to our so-called off-diagonal L complexes. Currently, we form a movie of RP autocorrelation surfaces (as demonstrated at the AIPR 2016 Workshop) for neighborhoods of P values of interest and it is upon each movie frame that the algorithm is applied. The RP movie demonstrates the evolution of structural changes on the diagonal as well as off the diagonal as the algorithm progresses through a neighborhood of P values. This movie can be made available by contacting the authors of this paper.




Fig. 14. Simplified diagram for our approach to auto-detection of structures on the autocorrelation surface.

In closing, we leave the reader with one last figure, Fig. 15, which is that of another Cepstral Terrain like that shown in Fig. 8. Except, to produce this figure, we applied our technique to a five-minute clip of meditation music recorded at 44,100 samples per sec. The figure was constructed by keeping the first 1600 cepstral components for 300 contiguous frames of one-second clips (44,100 samples). It is interesting to observe the on/off switching of the cepstral components as we slide up the frames. We believe it is with this type of recording a similar plot for electroencephalograms (EEGs) can be obtained with our




Fig. 15. Similar to Fig. 6, this is a Cepstral-time plot of meditation music. Cepstral components are derived from 1 sec snippets at 44,100 samples/sec. Evolution of the components in time is with respect to the increasing frame number. This is merely a fun proof-of-concept for future work with EEGs and DNA sequencing.

particular interest being to address various neurological conditions. Our technique with compressing a long time series with cepstral analysis and matrix framing visualization techniques to uncover meaningful structures may have a significant impact with the ability to extract intelligence from EEGs recorded with high sample rates (on the order of 5000 samples/sec) across the set of available electrodes used to record over an extended period of time (hours instead of minutes), an amount of data that would be previously prohibited from analysis. To be clear, this technique allows for a high level view to detect structures that would be 'combing' for crucial parts of a time domain recording that may in turn lead to focusing the attention on a fewer number of electrodes and their recordings.

ACKNOWLEDGMENT

We thank Mr. Edward (Jim) Cupples (Rome, NY), Mr. George Akmon (Dallas, TX), and Mr. Randall King (Birmingham, AL), for idea, direction, and insistence on clarity.

REFERENCES

- [1] Eckman, Kamphorst, and Ruelle, "Recurrence Plots of Dynamical Systems," *Europhysics Letters*, 4,(9), pp 973-977, 1987.
- [2] Iwanski and Bradley, "Recurrence Plot analysis: To imbed or not to imbed?," *Chaos*, vol. 8, 1998, 861-871.
- [3] B. Bogert, M. Healy, and J. Tukey, "The Quefrequency Analysis of Time Series for Echoes: Cepstrum, Pseudo Autocovariance, Cross-Cepstrum and Saphe Cracking," *Proceedings of the Symposium on Time Series Analysis* (M. Rosenblatt, Ed), pp. 209-243. New York Wiley, 1963.
- [4] D. Childers, D. Skinner, and R. Kemerait, "The Cepstrum: A Guide to Processing," *IEEE*, vol. 65, No. 10, October 1977.
- [5] Castiglioni, Meriggi, Faini, Di Rienzo, "Cepstral Based Approach for Online Quantification of EKG Quality in Freely Moving Objects," *Computing in Cardiology* 2011:38, 625-628.
- [6] Drugman, Bozkurt, and Dutoit, "Complex Cepstrum-based Decomposition of Speech for Glottal Source Estimation," *Conference for International Speech Communication Association*, 2009.
- [7] Metin Akay, *Biomedical Signal Processing*, Academic Press, 1994, Chapter 6.
- [8] R. Bracewell, *The Fourier Transform and Its Applications*, 2nd ed., McGraw-Hill, 1986, pp. 192, 358.
- [9] Atwood, Stanton, and Storey, *Introduction to Basic Cardiac Dysrhythmias*, 2nd Edition, Mosby Lifeline, 1996.
- [10] Pan and Tompkins, "A Real-Time QRS Detection Algorithm," *IEEE Transactions on Biomedical Engineering*, Vol. BME-32, No.3, pp. 230-236, March 1985.
- [11] A. Winfree, *When Time Breaks Down, The Three-Dimensional Dynamics of Electrochemical Waves and Cardiac Arrhythmias*, Princeton University Press, 1987, Chapter 3.
- [12] Arno Liberles, *Introduction to Molecular Theory*, Princeton University Press, Princeton, New Jersey, 1987, Chapters 1 and 2.