

Using the Memories of Multiscale Machines to Characterize Complex Systems

Nick S. Jones

Oxford Centre for Integrative Systems Biology, Departments of Physics and Biochemistry, Oxford University, Oxford OX1 3PU, United Kingdom

(Received 8 April 2007; published 22 May 2008)

A scheme is presented to extract detailed dynamical signatures from successive measurements of complex systems. Relative entropy based time series tools are used to quantify the gain in predictive power of increasing past knowledge. By lossy compression, data is represented by increasingly coarsened symbolic strings. Each compression resolution is modeled by a machine: a finite memory transition matrix. Applying the relative entropy tools to each machine's memory exposes correlations within many time scales. Examples are given for cardiac arrhythmias and different heart conditions are distinguished.

DOI: [10.1103/PhysRevLett.100.208702](https://doi.org/10.1103/PhysRevLett.100.208702)

PACS numbers: 89.75.-k, 05.45.Tp, 87.19.Hh, 89.70.Cf

Before understanding a complex system, one often needs to interpret the complex signals it generates. While it is easy to find correlations between arbitrarily separated pairs of points in a signal time series [1], highly correlated signals can lack such pairwise correlations (Section I). Instead, at a higher resolution, one can estimate the joint probabilities of sequences of events and, e.g., find the order of corresponding Markov chains [2]. Unfortunately the set of possible event sequences will generically increase exponentially with sequence length while becoming proportionately harder to estimate; i.e., what these methods gain in resolution over pairwise statistics, they lose in range. We must, however, expect nature to show correlations which are both long ranged and more than pairwise.

This Letter suggests a tool, akin to the autocorrelation, which is intuitive, sensitive to more than pairwise correlations and yet is long-ranged enough to capture the long-time correlations shown by some complex systems. It combines two concepts: (i) a natural measure of the predictive power one gains as one has an increasingly long symbolic string (Sec. II–V); (ii) use of lossy compression to express the dynamics of a complex system as a set of Markov sources (transition matrices) with each one representing the dynamics on a different time scale (VI–VIII). The following considers a system which, at any time t , can be in state x_t chosen from an alphabet (finite set) \mathcal{A} . The system passes through states x_1, x_2, \dots, x_T at fixed intervals and the data is ergodic and stationary. From now on x_j, \dots, x_k will be represented by x_j^k . Having addressed pairwise correlation measures in Section I, Sections II–VI develop a new means of mapping correlations in strings and VII–IX consider physiological examples and continuous time series.

I. Pairwise statistics can fail to capture structure.—Any approach which investigates the time structure of a data string must be compared with conventional methods like the autocorrelation function. Pairwise measures, which compare a symbol at one point in a string with a (possibly

different) symbol at another point, fail to capture conditional behavior on other intervening symbols. The following one-parameter, order-two transition matrix with alphabet $\mathcal{A}_4 = \{A, B, C, D\}$ should remind the reader of this phenomenon; it can create strings without pairwise correlations. Using the notation that X is a random variable and x is a particular instantiation of that variable, it is of the form $p(X_3|X_1X_2)$ with $(x_i \in \mathcal{A}_4)$

$x_1x_2 \setminus x_3$	A	B	C	D
AA	$\frac{1}{4}$	$\frac{1}{4}$	Λ^+	Λ^-
AB	$\frac{1}{4}$	$\frac{1}{4}$	Λ^-	Λ^+
BA	$\frac{1}{4}$	$\frac{1}{4}$	Λ^-	Λ^+
BB	$\frac{1}{4}$	$\frac{1}{4}$	Λ^+	Λ^-
pq	$\frac{1}{4}$	$\frac{1}{4}$	$\frac{1}{4}$	$\frac{1}{4}$

where $\Lambda^\pm = \frac{1}{4} \pm d$, the free parameter d is $0 \leq d \leq \frac{1}{4}$ and where pq is any ordered pair, p preceding q , other than AA, BB, BA, AB . One can prove that, for strings generated by this matrix, the probability of obtaining the symbol $u \in \mathcal{A}_4$ an interval $l \neq 0$ after $v \in \mathcal{A}_4$ is independent of both l and v . The autocorrelations or cross-correlations of such a series are thus absent (for explanations of symbolic autocorrelations, see Voss [1]). However, the string is highly structured; measures like $\mathbf{r}H_{n||m}$ below can expose this. Simple measures which reveal correlations between points, without having to create data objects which scale exponentially with order, can be very useful. The approach in Section VI yields data objects that both increase slowly with order or time and illuminate more than pairwise correlations.

II. The transition entropy.—If a string is sufficiently long, one can estimate transition probabilities $p(X_i|X_{i-m}^{i-1})$. It should be stressed that this paper is not, in the first instance, about the estimation of these probabilities; we will assume that they are given to us exactly [3]. The following entropies investigate the structure of this order m transition matrix.

Call $H(X_t) = \sum_{x_t \in \mathcal{A}} -p(x_t) \log p(x_t)$ the Shannon entropy of X_t . The transition, or conditional, entropy $\mathbf{t}H_m$ is defined as follows:

$$\mathbf{t}H_m = \sum_{x_{t-m}^{t-1} \in \mathcal{A}^m} p(x_{t-m}^{t-1}) H(X_t | x_{t-m}^{t-1}). \quad (1)$$

This transition entropy, $\mathbf{t}H_m$, measures the entropy of predictions one makes when equipped with a length m string, when one does not know what the string is. If each length m string that occurs exactly predicts the next state, then $\mathbf{t}H_m = 0$ (for the series $\dots ABABA \dots$ each length 1 string uniquely determines its ensuing state: $\mathbf{t}H_1 = 0$). If no string imparts a predictive advantage then $\mathbf{t}H_m = H(X_t) \forall m$. Convexity arguments [4] show that $\mathbf{t}H_m \geq \mathbf{t}H_n$ ($n > m$).

III. The relative transition entropy.—The *relative transition entropy*, $\mathbf{r}H_{n||m}$, defined below is a measure of gain in predictive power as one moves from knowledge of a length m string to a length n string ($n > m$). The relative entropy or Kullback-Leibler divergence [4] between the distribution $Q(X)$ and $P(X)$, where X can take $|\mathcal{A}|$ different values, is: $D(Q || P) = \sum_{x \in \mathcal{A}} Q(x) \log \frac{Q(x)}{P(x)}$. It is often described as the average disbelief in a model's predicted distribution P , when observing random outcomes X from real data Q . By contrast, we will use it to capture the degree that predictions made when equipped with more knowledge of the past, represented by Q , are inconsistent with those made with reduced knowledge, P .

One can compare predictions about a symbol at time t given knowledge of a particular set of preceding n symbols (x_{t-n}^{t-1}) with predictions made when given only the preceding m (x_{t-m}^{t-1} with $n > m$). The divergence, $D(p(X_t | x_{t-n}^{t-1}) || p(X_t | x_{t-m}^{t-1}))$, measures the information lost if one loses the knowledge that the sequence x_{t-m}^{t-1} was preceded by x_{t-n}^{t-m-1} . Averaging over all strings x_{t-n}^{t-1} yields the relative transition entropy $\mathbf{r}H_{n||m}$:

$$\mathbf{r}H_{n||m} = \sum_{x_{t-n}^{t-1}} p(x_{t-n}^{t-1}) D(p(X_t | x_{t-n}^{t-1}) || p(X_t | x_{t-m}^{t-1})). \quad (2)$$

This quantifies the predictive power lost when one moves from having a length n string of prior information to the shorter length m , for a randomly selected string x_{t-n}^{t-1} [5]. Let us now establish a few properties of $\mathbf{r}H_{n||m}$. Using (1) and (2) one can readily prove that $\mathbf{r}H_{n||m} = \mathbf{t}H_m - \mathbf{t}H_n$, $n > m$. If $\mathbf{t}H_m = \mathbf{t}H_n$ then $\mathbf{r}H_{n||m} = 0$. Since $D(Q || P) \geq 0$, with equality only when $Q = P$, we further know that if $\mathbf{r}H_{n||m} = 0$ then $p(X_t | x_{t-n}^{t-1}) = p(X_t | x_{t-m}^{t-1}) \forall X_t, x_{t-n}^{t-1}$. The length n and m predictions are *exactly* the same. In general $D(Q || P)$ can be unbounded [4], but here, some thought shows that $0 \leq \mathbf{r}H_{n||m} \leq H(X_t)$. We can now formulate a hierarchy of differential quantities. Defining the Shannon entropy for strings of length n as $H_n = H(X_{t-n+1}^t)$, one readily finds that $\mathbf{t}H_n = H_{n+1} - H_n$ and $\mathbf{r}H_{n||m} = \mathbf{t}H_m - \mathbf{t}H_n$. Authors have noted that the way that H_n and $\mathbf{t}H_n$

decrease with n , reveals structure in the string [6,7]: in Section VI we will use $\mathbf{r}H_{n||m}$ to map these correlations [5].

IV. Example: $\mathbf{r}H_{n||m}$ for the distribution in the table.— $p(X_3 | X_1 X_2)$ yields the stationary state, $p(X_1 X_2) = \frac{1}{16}$ so $p(X_2 | X_1) = \frac{1}{4}$. By Eq. (2) one finds $\mathbf{r}H_{1||0} = 0$ (comparing $p(X_2 | X_1)$ and $p(X_2)$). Comparing $p(X_3 | X_1 X_2)$ and $p(X_3 | X_2)$ shows that $\mathbf{r}H_{2||1} = \frac{1}{4} \log_2 [2(\Lambda^+)^{\Lambda^+} (\Lambda^-)^{\Lambda^-}]$. I.e. knowing the current state is of no help in predicting the next state ($\mathbf{r}H_{1||0} = 0$) but knowing the current and preceding state does help ($\mathbf{r}H_{2||1} > 0$ and so $\mathbf{r}H_{2||0} > 0$). Since, for $n > 2$, $m > 1$, $\mathbf{r}H_{n||m} = 0$ one concludes that knowing more than the preceding and current state gives no further predictive advantage.

V. Introducing a measure to detect concealed structure.—We noted that if the predictions of length m and $m+1$ strings differ then $\mathbf{r}H_{m+1||m} > 0$; however, since Eq. (2) is an average, small changes in this quantity can hide dramatic changes between the structure of order $m+1$ and order m transition matrices. One can readily construct examples where there exists a string x_{t-m-1}^{t-1} such that $D(p(X_t | x_{t-m-1}^{t-1}) || p(X_t | x_{t-m}^{t-1})) \gg \mathbf{r}H_{m+1||m}$. It is thus useful to introduce the quantity: $\mathbf{r}H_{m+1||m} = \max_{x_{t-m-1}^{t-1}} D(p(X_t | x_{t-m-1}^{t-1}) || p(X_t | x_{t-m}^{t-1}))$ the *maximum relative transition entropy* over all strings x_{t-m}^{t-1} in \mathcal{A}^m . This measures when knowledge of a *particular* extra symbol imparts a large predictive advantage.

VI. Introducing multiscale Markov sources.—This section introduces a method for describing data from complex systems by fitting finite state machines with memories to each of their different time scales. Consider coarsening time series to lower and lower time resolutions. For each resolution one might estimate a small, order m , transition matrix (Markov source is another name for transition matrix [4]). Let us call these matrices, one for each resolution, a set of multiscale Markov sources. Suppose the real data was generated by a high order Markov source of order $l \gg m$. An order l source (alphabet \mathcal{A}) has $|\mathcal{A}|^l (|\mathcal{A}| - 1)$ parameters. By contrast, we will see that a corresponding set of multiscale Markov sources requires only $\sim \log l$ parameters. The sources thus form a compact multiscale representation of the data.

Let us now examine more details of the coarsening. We first break the symbolic series, y_i^T , $y_i \in \mathcal{A}$, into consecutive nonoverlapping blocks, each c^r symbols long. We fix c , the basic block size, and let r vary to give different block sizes c^r (increasing r increases the block size and we will see that this lowers the resolution). Then we coarsen by mapping each possible block (of which there are $|\mathcal{A}|^{c^r}$) onto a single symbol from a smaller set \mathcal{C} ($|\mathcal{C}| < |\mathcal{A}|^{c^r}$). The manner of this map will be discussed below. The new coarsened string at resolution r has T/c^r elements ${}^r x_1^{T/c^r}$ with ${}^r x_i \in \mathcal{C}$. From this string one can estimate an order m Markov source, ${}^r M$. Supposing the raw data was generated by a source of order l , one might fix $|\mathcal{C}|$, c and m and vary r

to give a set of Markov sources ${}^r_m M$ with $r \in \{1, 2, \dots, \lceil \log_c \frac{l}{m} \rceil\}$. By choosing this range of r values the set of sources has a similar memory to the order l transition matrix. While the order l source needs $|\mathcal{A}|^l (|\mathcal{A}| - 1)$ parameters, the total number of parameters in the multi-scale sources is $|\mathcal{C}|^m (|\mathcal{C}| - 1) \lceil \log_c \frac{l}{m} \rceil$. The set of Markov sources ${}^r_m M \forall r \leq \lceil \log_c \frac{l}{m} \rceil$ thus gives a compact multi-scale dynamic model for the correlations at each timescale [8] (see top diagram in Fig. 2).

Such lossy compression lies broadly within rate-distortion theory [4]. Distortion measures capture the lossiness of maps from blocks to single symbols. The following motivational example uses the crude Hamming distortion. Blocks of c^r symbols, each symbol in \mathcal{A} , can be viewed as coming from an alphabet $\mathcal{B}_{c,r}$ of size $|\mathcal{A}|^{c^r}$. Call a compression a map $f: \mathcal{B}_{c,r} \rightarrow \mathcal{C}$. A map is optimal if a version reconstructed from the compressed string (using an inverse map $g: \mathcal{C} \rightarrow \mathcal{B}_{c,r}$) and the original string are as close as possible with respect to a given measure. The Hamming distortion is $d(g(f(X_t)), X_t) = \delta_{g(f(X_t)), X_t}$ for $X_t \in \mathcal{B}_{c,r}$. Given $p(X_t)$, the optimal map minimizes the expected symbol-by-symbol distortion between the reconstructed and original letters $\langle d(g(f(X_t)), X_t) \rangle$. Here, some thought shows that the optimal f : (i) takes each of the $|\mathcal{C}|$ most probable symbols in $\mathcal{B}_{c,r}$ to a distinct symbol in \mathcal{C} and (ii) takes all other symbols to an arbitrary symbol in \mathcal{C} .

VII. Example: sudden cardiac death.—This section applies the above tools to heart arrhythmia. A simplified view of the heart is that, in any short interval, it can have a normal (N) or ventricular beat (V) or no beat at all (\emptyset). The raw data is a list of times of beats labeled as N or V [9]. By discretizing time into blocks of $\tau = 32$ ms this list was converted into a symbolic string of the form ‘... $\emptyset N \emptyset V \emptyset \emptyset N$...’. The following uses 24 hours of heart data for a patient with many V beats. Given this three-letter alphabet one can attempt to estimate a transition matrix of order m : $p(X_t | x_{t-1}^{m-1}) \forall X_t, x_{t-1}^{m-1}$ [10]. Since $\tau \ll$ interbeat interval (~ 0.85 s) and the system is very structured, the size of the transition matrix grows slowly with m . Accommodation of finite size effects in the estimation of such transition matrices, is delicate [3] and a crude approach was used here (partly justified by the wealth of data). The transition matrices can only give information about the data as a whole (rather than the behavior of the heart at any one time) as, alongside the presence of multi-scale nonstationarities [11], this patient was particularly unhealthy. Figure 1(a) shows $\mathbf{r}H_{m||m-1}$ for the patient. As $m\tau$ (the duration of string one is given) increases towards the $N - N$ beat interval [see Fig. 1(c)], one’s ability to make good predictions increases markedly. But, when $m\tau$ nears the heart beat interval, further knowledge gives less predictive advantage (because one is already equipped with knowledge of a characteristic time period of the process). As a result $\mathbf{r}H_{m||m-1}$ begins to fall around 0.7 s, mirroring the distribution in Fig. 1(c). Beats with intervals >1 s are

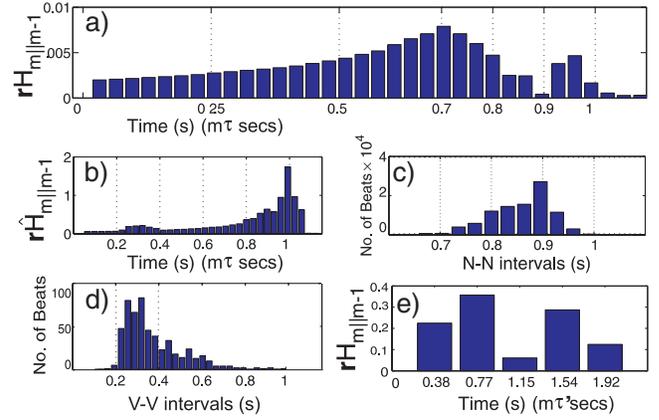


FIG. 1 (color online). A patient with cardiac arrhythmia: (a) The m th bar gives the predictive advantage of knowing $m\tau$ seconds of past activity over knowing $m\tau - \tau$ seconds of activity ($\mathbf{r}H_{m||m-1}$). $\tau = 32$ ms and $t \leq 35\tau$ (more than 99% of all length 35τ strings, picked uniformly at random from the data, occurred more than 300 times). (b) $\widehat{\mathbf{r}H}_{m||m-1}$ for the original data. (c),(d) The intervals between successive normal and ventricular beats, respectively. (e) $\mathbf{r}H_{m||m-1}$ for a coarsened string $r = 1$, $c = 12$, $|\mathcal{C}| = 10$ each time step is thus $\tau' = 0.384$ s. For length 4 and 5 strings in the coarsened data, 97% and 90%, respectively, of all such strings picked uniformly at random occurred more than 300 times.

rare so $\mathbf{r}H_{m||m-1}$ for $m\tau > 1$ s is small. Figure 1(b) plots $\widehat{\mathbf{r}H}_{m||m-1}$ (with the strong promise that all strings considered occurred more than 300 times in the data). It reveals hidden structure between 0.2 and 0.4 s. This peak is the compound effect of short $V - V$ events and misannotations in the uncorrected record [see Fig. 1(d)]. The coarsened data, Fig. 1(e), reveals structure on another time scale: one sees that a large part of the predictive knowledge is contained in the first second of activity but another characteristic time scale, open to physiological interpretation, appears in the range 1.19–1.54 s. Plots like Fig. 1 might distinguish between heart conditions, since these can depend on dynamics of a few seconds [12].

VIII. Continuous time series.—Multiscale Markov sources can also be found for continuous time series (e.g., wind speed) as well as symbolic strings. The data, sampled at T times, is again broken up into blocks of c^r consecutive points and coarsened. The alphabet we compress to is now a set of $|\mathcal{C}|$ letters with each one representing a different motif of c^r consecutive reals. We allocate each block of raw data to its closest motif using a mean-square distance. For example, suppose we want to compress blocks of two data points ($c = 2$, $r = 1$) to one of three symbols and we are given a three-letter code book with three motifs: $\mathcal{C} = \{N = (1, 1), V = (0.1, 2), U = (10, 10)\}$. Using the mean-square distance, the sequence of continuous data ${}^0y_1^T = \dots |11.1|0.21.5|1.11.01|\dots$ is optimally represented by ${}^1x_1^{T/2} = \dots |N|V|N|\dots$. See Fig. 2. The optimal set of motifs, for a fixed $|\mathcal{C}|$, allow the com-

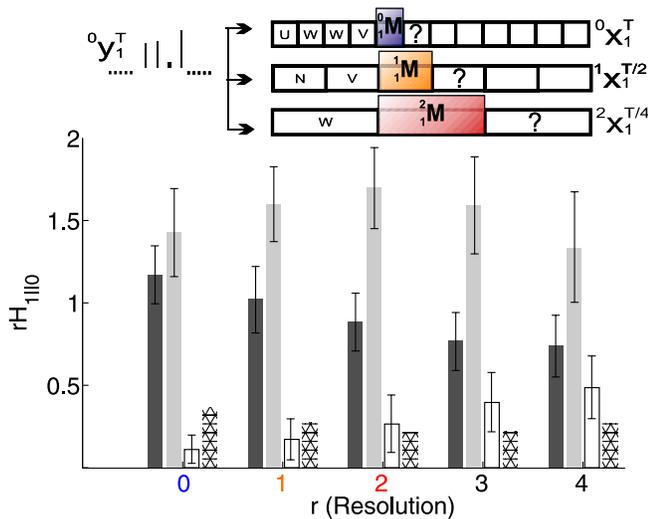


FIG. 2 (color online). Top: schematic of vector quantization (${}^1y_1^T \rightarrow {}^r x_1^{T/c^r}$) and multiscale Markov sources ${}^r M$. Main: The black, gray, and white bars are the means, \pm standard errors, of groups of patients who were healthy, experiencing congestive heart failure and atrial fibrillation, respectively. Unphysiological and ectopic beat intervals were filtered (as in [14]) and hourly trends removed. The fourth, hatched, columns are for random phased $1/f$ noise. The x axis considers the beat interval data at five different resolutions $r = 0, \dots, 4$ $c = 2$. At each resolution $rH_{1||0}$ is found; this measures the extra predictive power from knowing one symbol over knowing none. Gaussian white noise has $rH_{1||0} \rightarrow 0 \forall r$ and is not plotted (Brownian noise has $rH_{2||1} \rightarrow 0 \forall r$).

pressed sequence to reconstruct the original with minimum total mean-square error [13]. A selection of algorithms exists for finding optimal motifs (vector quantizers [4]). In our case, such algorithms have as input the set of T/c^r blocks of length c^r and the value of $|C|$. They output the set, C , of motifs of length c^r which minimizes the total mean-square error for this block size. Given these motifs one can then convert the continuous time series into its closest symbolic equivalent (see the above example: ${}^0y_1^T \rightarrow {}^1x_1^{T/2}$; see Fig. 2). Given this string of symbols generated from blocks of c^r real valued data points, one again determines ${}^r M \forall r$.

IX. Example.—Fig. 2 applies this idea to three different groups of cardiac patients [9], using the generalized Lloyd algorithm to find the appropriate set of motifs, C , for each resolution, r [4]. The raw cardiac data was an “interval series”: each data point being the time interval between successive heart beats. The matrix ${}^r M$ was found for $r = 0, \dots, 4$, $c = 2$, $|C| = 10$ and the extra predictive power from knowing one symbol was estimated: $rH_{1||0}$. The three different heart conditions can be distinguished (for comparable results see [14]); healthy hearts show a slow loss in predictability, the disordered beats that occur in atrial fibrillation yield a low degree of predictability whereas congestive heart failure shows an increase in predictability at some scales.

X. Conclusion.—This Letter presents a means of producing a one-parameter map of predictive knowledge acquired as one is equipped with increasingly long substrings of a symbolic data set. It suggests that lossy compression allows this short range mapping technique to be compactly extended to the study of longer ranged correlations. Examples are given where different heart conditions are distinguished and characterized using these methods. Underlying this work is the view that dynamical signatures of some systems can be found by treating them as sets of Markov sources with each source characterizing dynamics on a different time scale.

Thanks to M. Costa, A. Goldberger, C.-F. Lee, C.-K. Peng, and the NIH Research Resource for Complex Physiological Signals.

- [1] Symbolic examples: Fourier: R. F. Voss, Phys. Rev. Lett. **68**, 3805 (1992); Walsh-Fourier, etc.: D. Stoffer, J. Am. Stat. Assoc. **86**, 461 (1991); D. Stoffer *et al.*, Biometrika **80**, 611 (1993); Mutual information: A. M. Fraser and H. L. Swinney, Phys. Rev. A **33**, 1134 (1986).
- [2] From three fields: P. Billingsley, Ann. Math. Stat. **32**, 12 (1961); N. Merhav *et al.*, IEEE Trans. Inf. Theory **35**, 1014 (1989); M. J. van der Heyden *et al.*, Physica (Amsterdam) **117D**, 299 (1998).
- [3] H. Herzel *et al.*, Chaos Solitons Fractals **4**, 97 (1994).
- [4] T. Cover and J. Thomas, *Elements of Information Theory* (J. Wiley and Sons, NY, 1991); T. Berger, *Rate Distortion Theory* (Prentice-Hall, Englewood Cliffs, NJ, 1971); A. Gersho and R. Gray, *Vector Quantization and Signal Compression* (Kluwer Academic, Boston, MA, 1992).
- [5] T. Schreiber, Phys. Rev. Lett. **85**, 461 (2000) introduces the “transfer entropy” to reveal causal links between two time series. $rH_{n||m}$ reveals causal connection in the same series, and can be seen as the *self* transfer entropy.
- [6] P. Grassberger, Int. J. Theor. Phys. **25**, 907 (1986); W. Ebeling and G. Nicolis, Europhys. Lett. **14**, 191 (1991).
- [7] Ch. Bandt and B. Pompe, J. Stat. Phys. **70**, 967 (1993).
- [8] This can be connected with multiresolution source coding: e.g., D. Dugatkin, PhD. thesis Caltech, 2004.
- [9] Databases: www.physionet.org. Fig. 1 Sudden Cardiac Death Holter Fig. 2 the entire MIT-BIH Normal Sinus Rhythm and BIDMC Congestive Heart Failure; Fibrillation Data from [14].
- [10] M. Courtemanche *et al.*, Am. J. Physiol., Heart. Circ. Physiol. **257**, H693 (1989) extracts transition matrices from arrhythmias but discards interbeat intervals. Many papers have applied symbolic dynamics to hearts but, to my knowledge, exclude V beats and tend to record the change in beat intervals symbolically, e.g., J. Kurths *et al.*, Chaos **5**, 88 (1995).
- [11] P. Bernaola-Galvan *et al.*, Phys. Rev. Lett. **87**, 168105 (2001).
- [12] V. Schulte-Frohlinde *et al.*, Phys. Rev. E **66**, 031901 (2002).
- [13] A generalization of the coarsening in Costa *et al.* [14].
- [14] M. Costa *et al.*, Phys. Rev. Lett. **89**, 068102 (2002).