Statistical Pattern Detection in Univariate Time Series of Intensive Care On-Line Monitoring Data

Michael Imhoff*, Marcus Bauer[†], Ursula Gather[†], Dietrich Löhlein*

* Chirurgische Klinik, Städtische Kliniken, Beurhausstraße 40, D-44137 Dortmund, Germany

[†] Fachbereich Statistik, Lehrstuhl Mathematische Statistik und industrielle Anwendungen, Universität Dortmund, D-44221 Dortmund, Germany

Abstract

Objectives: To determine how different mathematical time series approaches can be implemented for the detection of qualitative patterns in physiologic monitoring data, and which of these approaches could be suitable as a basis for future bedside time series analysis.

Design: Off-line time series analysis.

Setting: Surgical intensive care unit of a teaching hospital.

Patients: 19 patients requiring hemodynamic monitoring with a pulmonary artery catheter.

Interventions: None.

Measurements and results: Hemodynamic data were acquired in 1-minute intervals from a clinical information system and exported into statistical software for further analysis. Altogether, 134 time series for heart rate, mean arterial pressure and mean pulmonary artery pressure were visually classified by a senior intensivist into five patterns: no change, outlier, temporary level change, permanent level change, and trend. The same series were analyzed with low order autoregressive (AR) models and with phase space (PS) models. The resulting classifications from both models were compared to the initial classification. Outliers and level changes were detected in most instances with both methods. Trend detection could only be done indirectly. Both methods were more sensitive to pattern changes than they were clinically relevant. Especially with outlier detection, 95% confidence intervals were too close. AR models require direct user interaction, whereas PS models offer opportunities for fully automated time series analysis in this context.

Conclusion: Statistical patterns in univariate intensive care time series can reliably be detected with AR models and with PS models. For most bedside problems both methods are too sensitive. AR models are highly interactive, and both methods require that users have an explicit knowledge of statistics. While AR models and PS models can be extremely useful in the scientific off-line analysis, routine bedside clinical use cannot yet be recommended.

Key Words

1 Introduction

Today most of our bedside decisions are based on subjective judgment and experience, rather than on hard data analysis. Most of the changes of a variable in over time are more important than one pathological value at the time of observation. Over the past three decades mathematical methods have been developed that allow the assessment of single or multiple variables over time.

There are various approaches to describe time dependent data generated from dynamical systems reflecting the different natures of the underlying processes. Time-dependent data are either generated in controlled scientific and engineering experiments or observed in medical, biological, environmental, and econometric studies. Although in neighboring sciences that deal with dynamical systems some of the same phenomenona are investigated, different terminologies and interpretations of the data generating mechanism are applied. This also leads to the use of a different calculus for data analysis.

Methods used for describing these systems assume that the process under consideration is deterministic or stochastic or a combination of the two. In mathematics and theoretical physics, methods from the theory of dynamical systems [1] are often used to describe deterministic processes, whereas in physical and engineering applications approaches from the fields of systems theory [2] and digital signal processing are frequently applied [3]. Most of this work deals with fully deterministic processes as well as deterministic processes and additive noise. In that context noise is seen as measurement errors or additive random effects, which are superimposed on the deterministic signal. By contrast, in statistical time series analysis the models are constructed in a way that random effects also drive the processes themselves. This approach is suitable to model medical, biological, environmental, and econometric variables like blood pressure or stock prices, because many of these processes can be regarded as stochastic processes [4, 5, 6, 7, 8, 9, 10]. We, therefore, concentrate here on time series analysis techniques for analyzing intensive care data.

Only a few investigations have employed this methodology in the field of intensive care medicine. In general it has been shown that time series analysis techniques are suitable for retrospective analysis of physiologic variables [8, 9, 10]. A statistically similar, but methodologically more challenging task is the on-line analysis of intensive care monitoring data. Statistically, approaches to this problem are rare and not readily available for application in clinical practice [11, 12, 13].

For this study, two entirely different statistical methods were used to describe critical care time series: autoregressive models, which have been used in numerous applications since their introduction in the 1970s, and phase space models, which represent a new approach to time series data. In autoregressive models the current value of a process is expressed as a finite, linear aggregate of previous values of the process and a stochastic term [14, 15]. In phase space models, after a transformation, the time series data are regarded as a multivariate sample with dependent observations [16].

Several authors have applied autoregressive models in the field of critical care [8], in longitudinal physiological experiments [10], as well as in studies of laboratory data of the chronically ill [9]. Autoregressive models have also been successfully used to describe on-line time series from intensive care bedside monitors [8].

Clinical investigations with autoregressive models into therapeutic effects have recently been reported from the fields of cardiology [10], rheumatology [17], neurology [18], psychiatry [19], and nursing research [20].

The phase space approach originally came from the theory of nonlinear systems and is used for analysis of complex, deterministic, and especially chaotic systems. Several authors have applied measures which are based on phase space embeddings to judge the complexity of heart rate dynamics [21, 22, 23]. The application of these methods to heart rate dynamics is only feasible, when data are recorded with high sampling frequency (e.g., 50 ms [21]).

In clinical information systems data is often recorded in longer time intervals (e.g., 1 min) and, therefore, a different approach for analysis has to be selected. A new application of phase space models was recently introduced in the context of intensive care data, where phase space models are formulated in the statistical framework of linear stochastic systems [16]. This allows a meaningful application of these models to different physiologic variables, even with small sample sizes.

The detection of qualitative patterns in physiologic monitoring data (e.g., outliers, level changes, trends) is one of the basic applications of medical time series analysis. Traditional statistical methods for pattern identification like cluster or discriminant analysis are appropriate for time independent data [24] but cannot be used for identifying and discriminating time series patterns. One possible statistical approach to the identification of patterns in time series are state space models, first used in the engineering sciences [25]. Several applications of Kalman Filter techniques to intensive care data exist [12, 26]. But these procedures are not very reliable in the identification of patterns and they require significant computational power [26, 27].

Recently, neural networks have been used for describing and controlling dynamical systems [28, 29]. The transfer of neural networks to pattern recognition in intensive care data is difficult: Every patient has to be controlled individually, and it is not clear how to construct training phases. It is also difficult to decide with the help of neural networks whether or not the state of a critically ill patient is improving, because too little information about the health of the patient is available for the training phase [30, 31].

In the following investigation we chose autoregressive and phase space models, because they are suitable to model the underlying dynamic processes of intensive care variables and seem to be promising for the identification of patterns [8, 9, 16]. This study evaluates two aspects of the application of time series analysis to on-line monitoring data: (a) can all patterns be correctly identified with the applied statistical methods? (b) are the applied methods adequate for clinical use?

2 Patients and methods

On a 16-bed surgical intensive care unit in a 2000 bed teaching hospital on-line monitoring data were acquired from 19 consecutive critically ill patients (8 females, 11 males, mean age 65 years), who had pulmonary artery catheters for extended hemodynamic monitoring, in 1-minute intervals from a standard clinical information system. These data were transferred into a secondary SQL database and exported into standard statistical software for further analysis. The system configuration was comprised of the following: (a) Clinical Information System (CIS): Emtek Continuum 2000, Version 4.1M3, (b) Decision Support System (DSS): Sybase SQL server 4.9.2, (c) Statistical Software: SPSS version 6.1, SAS version 6.12 with additional programs in SAS/IML, and (d) System platform: Sun Sparc, Sun Solaris 1.1.2 (CIS, DSS) and 2.5 (statistical software).

From a total of 550,000 single observations of seven variables (heart rate and invasive pressures), segments of 150 to 500 observations for each variable were visually classified by a senior intensivist into five medically relevant patterns: no change, outlier, temporary level change, permanent level change, and trend. The intensivist did not have to define objective criteria to explain why he chose a specific classification. A total of 134 time series were included in the study. The classifications are listed in table 1. All these a priori classifications were done by one senior intensivist. After all time series were classified, they were presented again in a different order for reclassification by the same intensivist. There were no differences in the classification, which was attributable to the obviousness of the patterns as shown in the examples in figures 2 to 6.

The same segments were analyzed with second order autoregressive [AR(2)] models and with phase space [PS] models.

Each series of autoregressive models was split into a model estimation period, and a prediction period, where the pattern should be diagnosed. The average length of the estimation period was 173 minutes - that of the prediction period 123 minutes (table 1).

Table 1   Physiological time series included in the study:
                Variables and patterns, as identified a priori by an intensivist.
                Length of estimation and prediction periods for autoregressive models.

Variables and Patterns						
	No Change	**Outlier**	**Level Change**		**Trend**	**Total**
			Temporary	**Permanent**		
Variable						
HR	8	6	7	11	8	40
AP (mean)	5	24	5	7	0	41
PAP (mean)	10	5	12	24	2	53
Total	23	35	24	42	10	134
Estimation Period for AR-models						
Mean	208	163	144	177	184	173
Maximum	501	481	334	451	410	501
Minimum	81	51	41	50	80	41
Prediction Period for AR-models						
Mean	98	94	152	150	94	123
Maximum	299	270	470	337	150	470
Minimum	20	20	40	50	50	20

The actual measurements were compared to the confidence intervals (CI) for the prediction period. According to the percentage of values outside the CI, the variation was classified into the different patterns. Values outside the CI were classified as an outlier, if less than five consecutive observations (= minutes) were outside the CI. A level change was identified by five or more consecutive observations outside the CI, and was called a temporary level change with less than 50% of the prediction period outside the CI, and a permanent level change with 50% or more of the prediction period outside the CI.

Trend patterns were indirectly identified by deviations of the autocorrelation function (ACF) of the residuals and the Durbin-Watson statistics. In this case, the ACF of the original series was analyzed for trend typical patterns. Typically, the ACF plot of a time series without a trend declines exponentially (diagram 1). In the case of a trend (i.e., a nonstationary series) the ACF plot fades only slowly over a larger number of lags (diagram 1). In these cases, an AR(2) model of the first derivative of the time series was calculated. If this model showed sufficient goodness of fit, a significant trend of the time series was assumed.

Diagram 1   Top: Plot of the autocorrelation function for the first 16 lags of a series without a trend (prediction period of the series from diagram 3). Rapid decline of the coefficients for the ACF.
Bottom: Plot of the autocorrelation function for the first 16 lags of a series with a trend (prediction period of the series from diagram 6). Slow decline of the coefficients for the ACF.
Solid bars = Coefficients of the ACF or PACF.
Solid lines = 95% Confidence Limit for the ACF or PACF.

Because the correct CI width was unknown before the study, a posteriori adjustment of the CI was done and the analysis rerun with this CI. This was done in cases where an inexplicable difference between the visual classification and the percentage of outliers occurred. The margin of the CI determines the "sensitivity" of the model. In this study an autoregressive or phase space model was con-

sidered too sensitive if at least one outlier was detected by the model that was not described in the initial visual classification.

In order to compare the CIs between different time series as one measure of the sensitivity of the predicted model the relative confidence interval (CI_{rel}) was calculated as the difference between the upper (UCL) and lower confidence limits (LCL) in relation to the fitted model (FIT) for the entire prediction period expressed in percent:

$$CI_{rel}[\%] \; = \; 100 \cdot \frac{UCL - LCL}{FIT}$$

For PS models, the first 60 observations were taken and retrospectively analyzed (i.e., outlying regions were estimated and patterns in this time interval identified). After this, a time window of length 60 was moved through the data. That means, that at time point 61 it was determined if the phase space vector y_{61} was in a distant region. If not, then no pattern was present, and the estimated outlying region was replaced by a new one, that was estimated from the last 60 observations $y_2, \ldots, y_{61}$. This was repeated for every new observation as long as for the time point t the phase space vector $\vec{y}_t$ was in a distant region. Then the system was not in a steady state, and after analyzing the consecutive observations $y_{t+1}, y_{t+2}, \ldots$, it was decided if a pattern was present similar to the retrospective analysis.

A detailed description of the concepts of autoregressive and phase space models and the underlying statistics is given in appendix 1.

3 Results

With autoregressive models, both outliers and level changes could always be identified. All series that were classified as "without change" were also correctly identified by autoregressive models. Phase space models allowed the identification of series without any change and with outliers in evary instance, too. Temporary level changes were correctly identified in 20 out of 24 series, and permanent level changes in 37 out 42. In the instances where identification failed, the level changes were marked by a very slow decrease or increase of the observed values.

Diagrams 2 to 6 display typical examples of each pattern analyzed by the two different methods. Diagram 2 shows the same series for mean pulmonary artery pressure, which was clinically classified as "no change", while in the AR model the confidence interval was adjusted to 99% to allow the correct classification. The graphic representation of the series without a change is obvious in the PS model and is confirmed by the phase space procedure.

Typical graphs for outliers are shown in a time series for heart rate in diagram 3. As before, the 95% CI for the AR model was too close for the clinical situation. With a 99.9% CI, two outliers could be identified that extrude from the CI. The same outliers are graphically represented by two vectors in the PS model that protrude from the imaginary ellipsoid of the vector cloud.

Diagram 2 Top: Autoregressive model for a time series without significant change: AR model for mean PAP. All data points for the prediction period are within the 99% CI.
  Solid line: Measured series of mean pulmonary artery pressure (PAPm)
 Dashed line: fitted AR(2) model of PAPm
  Dotted lines: 99% confidence interval for AR(2) model of PAPm
  Time: time after start of first measure in minutes
Bottom: Phase space model for the same time series: PS-model for mean PAP. All vectors of the PS model are within an imaginary ellipsoid.
  d_n: Value for PAPm at time "n"
  $d\_(n-1)$: Value for PAPm at time "n-1" (i.e. one observation prior to d_n)

The analysis of a time series for heart rate with a temporary level change is shown in diagram 4. In the PS model the temporary level change is represented by a small secondary ellipsoid of vectors that extrude from and fall back into the main ellipsoid, when the series returns to baseline values. Similar-

Diagram 3 Top: Autoregressive model for a time series with an outlier: AR model for mean AP. Three data points (outliers) are outside the 99.9% CI.

Solid line: Measured series of heart rate (HR)
Dashed line: fitted AR(2) model of HR
Dotted lines: 99.9% confidence interval for AR(2) model of HR
Time: time after start of first measure in minutes

Bottom: Phase space model for the same time series with an outlier: PS model for mean AP. Few vectors extrude from the imaginary ellipsoid. Each vector representing an outlier.

d_n: Value for HR at time "n"
d_(n-1): Value for HR at time "n-1" (i.e. one observation prior to d_n)

ly, the time series lies outside the 95% CI of the AR for the time of the temporary level change.In the PS model of a permanent level change the vectors from the secondary ellipsoid do not fall back into

9

the main ellipsoid, as the time series in the AR model will not return to the baseline within the pre-diction period (diagram 5).

Diagram 4   Top: Autoregressive model for a time series with a temporary level change: AR-model for HR. A series of values is outside the 95% CI. Quantification with additional regressor.
  Solid line: Measured series of heart rate (HR)
  Dashed line: fitted AR(2) model of HR
  Dotted lines: 99.9% confidence interval for AR(2) model of HR
  Time: time after start of first measure in minutes
Bottom: Phase space model for the same time series: PS model for HR Several vectors extrude from the imaginary ellipsoid and form an additional ellipsoid, which falls back to the main ellipsoid. Changes cannot be quantified.
  d_n: Value for HR at time "n"
  d_(n-1): Value for HR at time "n-1" (i.e. one observation prior to d_n)