Patterns of Dependencies in Dynamic Multivariate Data

Ursula Gather[1], Roland Fried[1], Michael Imhoff[2], and Claudia Becker[1]

[1] Department of Statistics, University of Dortmund, Vogelpothsweg 87, 44221 Dortmund, Germany
{gather, fried, cbecker}@statistik.uni-dortmund.de
http://www.statistik.uni-dortmund.de/lehrst/msind/msind_e.htm
[2] Surgical Department, Community Hospital Dortmund, Beurhausstr. 40, 44137 Dortmund, Germany
mike@imhoff.de

Abstract. In intensive care, clinical information systems permanently record more than one hundred time dependent variables. Besides the aim of recognising patterns like outliers, level changes and trends in such high-dimensional time series, it is important to reduce their dimension and to understand the possibly time-varying dependencies between the variables. We discuss statistical procedures which are able to detect patterns of dependencies within multivariate time series.

1 Introduction

Modern technical possibilities allow for simultaneous recording of many variables at high sampling frequencies. Possibly there are interactions between the observed variables at various time lags and we have to treat the data as multivariate time series. Often the data contain strong dynamic structures, that are unknown in advance. In intensive care for instance, physiological variables, laboratory data, device parameters etc. are observed for each critically ill patient. The appropriate analysis and online monitoring of this enormous amount of dynamic data is essential for suitable bedside decision support in time critical situations [**?**]. Thus, methods for automatic abstraction of the dynamical information into clinical relevant patterns are needed. Physicians typically select some of the observed variables and base their decisions on patterns like level shifts and trends detected in them. Statistics offers alternatives such as dynamic factor analysis providing a few latent variables which describe the dynamical information in the data appropriately. However, to get reliable and interpretable results by any method for dimension reduction without substantial loss of information we need to understand the relations between the variables.

We analyse 11-variate time series describing the hemodynamic system as measured in intensive care, consisting of arterial and pulmonary artery blood pressures (diastolic, systolic, mean), denoted by APD, APS, APM, PAPD, PAPS, PAPM, central venous pressure (CVP), heart rate (HR), pulse (PULS), blood

temperature (TEMP), and pulsoximetry (SPO2). We first discuss graphical models, that reveal the partial correlation structure within multivariate time series [?], [?]. Performing such graphical analyses for different physiological states allows to characterise distinct health states by distinct correlation structures [?]. Next, we briefly explain how the information obtained from graphical models can be used to enhance dynamic factor modelling. Finally, methods for detecting non-linear dependence-patterns (SIR; [?]) are transferred into the context of high-dimensional medical time series.

2 Graphical Models

2.1 Partial Correlation Graphs

Graphical models are almost standard nowadays for investigating relations within multivariate data [?], [?], [?], [?]. A *graph* $G = (V, E)$ consists of a finite set of *vertices* V and a set of *edges* $E \subset V \times V$. A visualization can be accomplished by drawing a circle for each vertex and connecting each pair a, b of vertices for which $(a, b) \in E$ or $(b, a) \in E$. If only one of these pairs is included in E, e.g. (a, b), then a directed edge (*arrow*) is drawn from a to b. If both pairs are included in E an undirected edge (*line*) is used. An arrow specifies a directed influence, while a line stands for a symmetrical relation.

When analysing the relations between the vital signs of an individual we should consider the time series structure of the measurements, which are not independent. Partial correlation graphs for multivariate time series introduced by Brillinger [?] and Dahlhaus [?] visualize the (partial) *linear relations* between the components of a multivariate time series. Such an analysis of symmetric relations is a first step to get a better understanding of the underlying dependence structure. We just note that directed graphical models providing information on possible causalities have also been developed recently [?].

In the following, let $X(t) = (X_1(t), \ldots, X_k(t))'$, $t \in \mathbb{Z}$, be a multivariate stationary time series of dimension k. Suppose that the autocovariance function

$$\gamma_{ab}(h) = Cov(X_a(t+h), X_b(t)),\ h \in \mathbb{Z}, \tag{1}$$

is absolutely summable with respect to all time lags h for all $a, b \in \{1, \ldots, k\}$. Then the *cross-spectrum* between the time series $X_a = \{X_a(t), t \in \mathbb{Z}\}$ and $X_b = \{X_b(t), t \in \mathbb{Z}\}$ is defined as the Fourier-transform of their covariance function $\gamma_{ab}(h), h \in \mathbb{Z}$,

$$f_{ab}(\lambda) = f_{X_a X_b}(\lambda) = \frac{1}{2\pi} \sum_{h=-\infty}^{\infty} \gamma_{ab}(h) \exp(-i\lambda h)\,. \tag{2}$$

This defines a decomposition of the covariance function γ_{ab} into periodic functions of frequencies λ. The variables X_a and X_b are uncorrelated at all time lags h iff $f_{ab}(\lambda)$ equals zero for all frequencies [?].

In order to distinguish between direct and induced linear relations between two component series X_a and X_b, the linear effects of the remaining components $Y = \{(X_j(t), j \neq a, b), t \in \mathbb{Z}\}$ have to be controlled. For this the optimal $\mu_a \in \mathbb{R}$ and the optimal filter $d_a(h), h \in \mathbb{Z}$, have to be determined, such that the quadratic distance

$$E\left[X_a(t) - \mu_a(t) - \sum_h d_a(h) Y(t-h)\right]^2 \tag{3}$$

is minimized. Let $\epsilon_a = \{\epsilon_a(t), t \in \mathbb{Z}\}$ be the residuals obtained from this, and calculate ϵ_b from X_b in the same way. The correlations between these residual series at all time lags define the partial correlation structure between X_a and X_b after eliminating the linear effects of the other variables. The *partial cross-spectrum* $f_{X_a X_b \cdot Y}(\lambda)$ between X_a and X_b is then defined as the cross-spectrum between ϵ_a and ϵ_b, while the *partial spectral coherence* is a standardization hereof,

$$R_{ab \cdot Y}(\lambda) = \frac{f_{ab \cdot Y}(\lambda)}{\sqrt{f_{aa \cdot Y}(\lambda) f_{bb \cdot Y}(\lambda)}}\,. \tag{4}$$

In a *partial correlation graph* for a multivariate time series the vertices $a = 1, \ldots, k$ are the components of the time series and an undirected edge between two vertices a and b is omitted, $(a, b) \notin E$, whenever the correlations between the residual series ϵ_a and ϵ_b are zero for all time lags. This is equivalent to $f_{X_a X_b \cdot Y}(\lambda) = 0$ for all frequencies $\lambda \in \mathbb{R}$. This defining property is called the *pairwise Markov property*. Partial correlation graphs illustrate the indirect character of some marginal correlations as two variables a and b are marginally related (possibly via other variables) if they are *connected* by a *path*, i.e. if vertices $a = a_0, a_1, \ldots, a_k = b$, $k \geq 0$, exist, such that there is an edge between each subsequent pair of vertices.

If the spectral density matrix is regular at all frequencies the pairwise Markov property implies the *global Markov property* for this kind of graphical model, that is a stronger property in general [**?**]. The latter property says that two sets of variables $A \subset V$ and $B \subset V$ have zero partial correlations given the linear effects of a set of variables $C \subset V$ if C separates A and B in G, i.e. if any path between two variables $a \in A$ and $b \in B$ necessarily contains at least one variable $c \in C$. In other words, the variables in A and B are not related if the effects of the separating subset C are controlled. This allows to illustrate zero partial correlations by means of separation properties of a partial correlation graph. The subset C may contain less than all the remaining variables which allows to identify important variables more clearly.

2.2 Application to Physiologic Time Series

In the following we apply partial correlation graphs for multivariate time series to physiologic variables representing the hemodynamic system. Online-monitoring data was acquired in one minute intervals from 25 consecutive critically ill patients with pulmonary artery catheters for extended hemodynamic

monitoring, amounting to 129943 sets of observations altogether, i.e. about 5200 observation times are available for each patient on the average. We compare "empirical associations" found by statistical analysis to "physiological associations" based on medical knowledge. Physiological association means that a change in one physiological variable leads to a change in another physiological variable. This term does not imply any causal, linear or non-linear relation. Not even a direction or an ordering in time is expected for it as e.g. during volume depletion an increase of CVP and PAPM will typically lead to an increase in APM and a decrease in HR, while in congestive heart failure under high doses of inotropes an increase in HR may typically lead to an increase in APM.

The program "Spectrum" developed by Dahlhaus and Eichler [**?**] estimates the cross-spectra by a nonparametric kernel estimator. The partial spectral coherences are estimated from these cross-spectra using an inversion formula derived by Dahlhaus [**?**]. Then a decision has to be made on whether the partial spectral coherences may be identical to zero because sampling variability always causes estimates to be distinct from zero. Spectrum also constructs an approximate bound for the 95%-percentile of the maximal squared estimated partial spectral coherence under the assumption that the true partial spectral coherence equals zero. This allows one to perform an approximate 5%-test for the hypothesis of partial uncorrelatedness of two variables by comparing the estimated partial spectral coherence with this bound. However, it is well-known that different relations between physiological variables may have distinct strengths. The strength of a relation can be understood as an expected relative change in one of the variables when the other one changes by a certain relative amount. Therefore, we classify the relations as high, medium, low and zero partial correlation on the basis of the area under the estimated partial spectral coherence $R_{X_a X_b \cdot Y}$. This area can be measured by the partial mutual information between the time series X_a and X_b, which is defined by

$$-\frac{1}{2\pi}\int \log\{1-|R_{ab\cdot Y}(\lambda)|^2\}d\lambda, \tag{5}$$

(see Granger and Hatanaka [**?**] and Brillinger [**?**]) or by variants of this. Then we construct partial correlation graphs using gradually distinct edges to represent the strength of the empirical relation.

Figure 1 displays the resulting partial correlation graph for the hemodynamic system derived from one patient. Neglecting the many low relations we can identify some groups of strongly related variables from the graph. High partial correlations exist between the systolic, diastolic and mean arterial pressure, between the heart rate and the pulse, as well as between the systolic, diastolic and mean pulmonary artery pressure. The central venous pressure is mainly related to the pulmonary artery pressures. The blood temperature and the pulsoximetry seem to be rather isolated. These findings are similar for all patients analysed in our case-study, particularly the subgroups of highly partially correlated variables could always be identified. Some differences were found w.r.t. the relations of the other variables (CVP, Temp, SPO2) when performing such a

one-step model selection, but the partial correlation graphs obtained by a stepwise search strategy [14] closely resembled the one shown here in most of the cases. For some of the patients medium to high partial correlations were found between temperature and pulsoximetry, which were not expected by medical knowledge, but turned out to be artifacts as the measuring instrument used for pulsoximetry is influenced by the temperature.

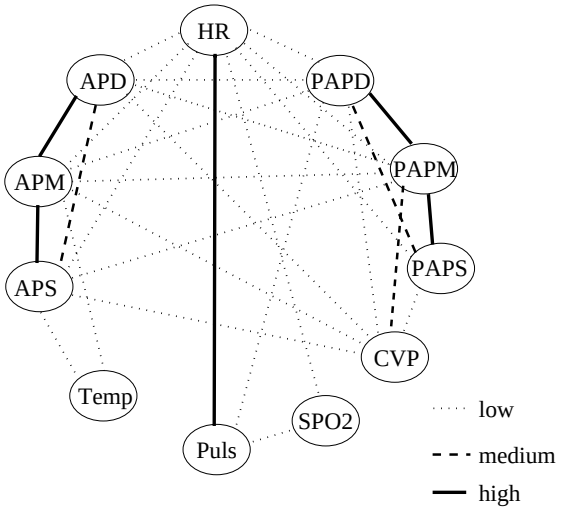

Fig. 1. Partial correlation graph for the hemodynamic system of one patient.

Such a partitioning of the variables into strongly related subgroups can be used to reduce the number of variables which have to be considered for online monitoring, i.e. for variable selection. From Figure 1 we see that APM provides a lot of information (measured via the partial mutual information) on both APD and APS, while both APD and APS provide less information on each other. Selecting APM from this strongly related subgroup and neglecting APD and APS for clinical monitoring is therefore meaningful from a statistical point of view. In the same way we see that choosing one variable out of PAPD, PAPM and PAPS should be sufficient. Indeed, physicians often select the mean pressures APM and PAPM as well as the heart rate. Thus, they choose exactly that variable out of each subgroup identified by the partial correlation model which has the strongest relation to the other variables. This is due to the nature of the mean pressures, which are "in between" the diastolic and systolic pressures. Hence, the statistical results agree with medical knowledge.

2.3 Distinguishing Clinical States

A patient can suffer from several clinical states like pulmonary hypertension, septic shock, congestive heart failure and vasopressor support during his visit at an intensive care unit. Distinct clinical states are accompanied by different

pathophysiological responses of the circulatory system. These changes may be supposed to result in differences in the interactions between the vital signs, particularly in the interactions between the groups of closely related variables identified above. Next we investigate, whether partial correlation graphs can detect differences in the status of the circulatory system. We drop systolic and diastolic pressures in the following calculations, which is in line with medical reasoning as systolic, diastolic and mean blood pressures are different representations of the same physiological process, i.e. a pulsatile blood pressure. The reason for dropping some of the variables is that in the previous data analysis the linear influences of *all* other variables have been subtracted. Elimination of the influences of APD and APS from APM e.g. may hide some relations of APM to other variables as the remaining variability is very low.

For each patient the predominant clinical state is determined for every time point as scored from the medical record using a set of rules used at the Hospital Dortmund. In the following we estimate the relations within a set of 'important variables' consisting of HR, SPO2, APM, PAPM, CVP and Temp for each state and for each patient separately.

Figure 2 summarizes the results of our data analysis using partial correlation graphs. Although the number of samples is small, there are obvious differences between the partial correlation patterns for distinct clinical states. Most of these differences can be explained by known physiological mechanisms. While for most states strong partial correlations could be found between heart rate and blood pressures as well as between the blood pressures, the status of *pulmonary hypertension* is predominantly associated with strong partial correlations between PAPM and CVP. This corresponds to the fact that this state can be characterized by an elevated PAPM. Since CVP also influences the right ventricle, one expects strong interactions between CVP and PAPM. On the other hand, the higher resistance within the pulmonary bloodstream attenuates the interactions between PAPM and APM as changes in PAPM will have a less than normal effect on left ventricular preload. In the status of *congestive heart failure* there are strong partial correlations between APM and PAPM. This can be explained by a failure of the left ventricle, that causes a decrease of APM as the heart is not able to push the blood forward (forward failure). In consequence there is a build-up in front of the left ventricle and thus an increase of PAPM (backward failure). For the clinical status of *vasopressure support* there are strong partial correlations between APM and PAPM, too. However, there are also medium to high partial correlations between HR and APM. This is due to the therapy the patient gets in this status, which affects the heart as well as the bloodstream and puts off the usual autonomous regulation of the circulatory system.

3 Factor Analysis

The detection of patterns like outliers, trends and level shifts in multivariate data becomes increasingly complex with increasing number of dimensions. Therefore, physicians usually base their decision on a subset of the variables which they

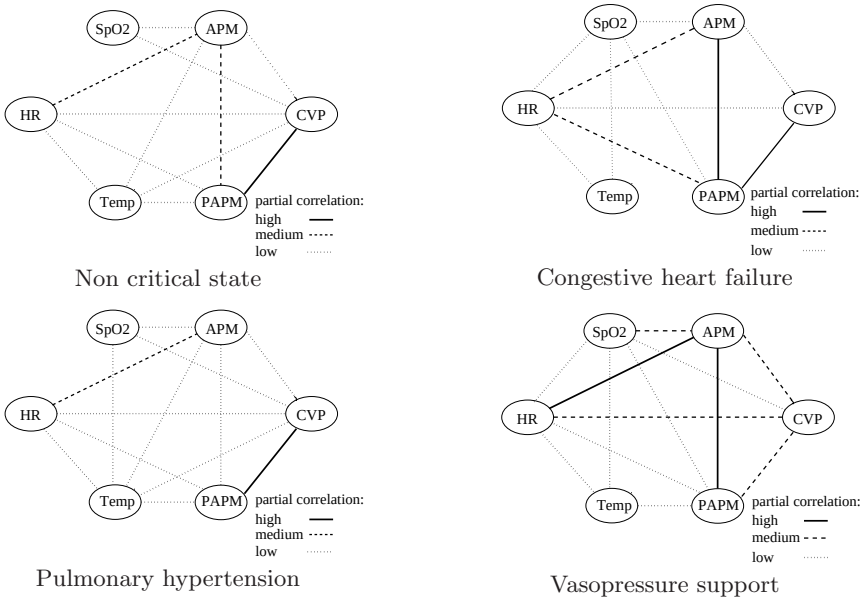

Fig. 2. Classification of the partial correlation patterns for distinct clinical states.

regard as particularly important. Similarly, factor analytic methods allow to compress the data into a few latent variables which capture the essential information in the observed data as good as possible in a certain (statistical) sense. In this way a reduced number of factor series is obtained, that can be more easily overlooked by the physician than the observed series, the number of model parameters is reduced, and clinical relevant patterns are easier to detect also by an automatic data analysis. More precisely, in factor analysis it is assumed that there are a few, say l, latent variables called factors which drive the series and cause the correlations between the observable variables. For dynamic factor analysis of a multivariate time series $\{X(t), t \in \mathbb{Z}\}$ Peña and Box [**?**] propose the model

$$X(t) = \Lambda Z(t) + \varepsilon(t), \quad t \in \mathbb{Z}, \tag{6}$$

where Λ is a $k \times l$-matrix of loadings, $Z(t)$ are l-dimensional vectors of latent factors following a VARMA(p,q)-model, and $\{\varepsilon(t), t \in \mathbb{Z}\}$ is a k-dimensional process of Gaussian white noise with zero mean and arbitrary covariance matrix Σ_ϵ, which is independent of $\{Z(t), t \in \mathbb{Z}\}$. To get identifiability of the parameters, $\Lambda'\Lambda$ can be restricted to be the identity. If the factor series are independent, i.e. $Z(t)$ follows a VARMA(p,q)-model where all coefficient matrices are diagonal, then the time-lagged autocovariance matrices $\Gamma_X(h)$ of $\{X(t), t \in \mathbb{Z}\}$ are symmetrical for $h \geq 1$ and the columns of Λ can be chosen as the common eigenvectors of $\Gamma_X(h)$ while the eigenvalues $\gamma_i(h)$, $i = 1, \ldots, l$, are the diagonal elements of the autocovariance matrices $\Gamma_Z(h)$ of $\{Z(t), t \in \mathbb{Z}\}$ then. These findings can be used to identify a factor model for a given time series [**?**].

For the construction of an intelligent alarm system based on a dynamical factor analysis of the vital signs it is important that the factors are interpretable for the physician and that they reveal the important patterns found in the observable variables. Then we can simply monitor the factors and a summary measure of the model errors like the sum of the squared residuals for each time point.

One possibility to achieve better interpretability of the factors is to apply a rotation in the l-dimensional space. The previous findings obtained using graphical models justify separate factor analyses for the groups of strongly related variables, cf. Figure 1. Treating the subsets consisting of the arterial pressures, of the pulmonary artery pressures including the central venous pressure, and of the heart rate and the pulse simplifies the task of factor extraction from the hemodynamic system. We treat the blood temperature and the pulsoximetry separately as these do not show strong relations to the other variables. Since the variables of each group are measured on the same scale, we use the sample covariance matrices and calculate the eigenvalues and the eigenvectors for each group for the covariance matrices up to the time lag 4. We find one factor to be sufficient for each group. The resulting factor loadings are shown in Table 1. Very similar loadings are obtained if we analyze the variables jointly and use VARIMAX rotation to get better interpretable results. Hence, the factor loadings calculated in a factor analysis of all variables "identify" each of the rotated factors to belong to one of the subgroups. Thus, the results of both analyses nearly coincide. Analysing the variables in groups as is done here, however, affords less observations to get stable eigenvectors and provides better interpretable results as the loadings of the other variables not included in the respective group are exactly identical to zero.

Table 1. Factor loadings calculated for grouped variables.

Variable	factor 1	factor 2	factor 3
PAPD	0.3671	0	0
PAPM	0.5414	0	0
PAPS	0.6849	0	0
CVP	0.3211	0	0
APD	0	0.2449	0
APM	0	0.4776	0
APS	0	0.8438	0
HR	0	0	0.6964
PULS	0	0	0.7177

A closer inspection reveals that the calculated factors represent structural changes found in the corresponding component series better than any single variable [**?**]. Thus, dynamic factor modelling may facilitate the detection of patterns like level shifts and trends in multivariate physiological time series. Moreover, one may speculate that clinical relevant patterns of dependencies may be

detected from an analysis of the partial correlations between the factor series similarly as in Section 2.3. For this aim, methods for estimating time-varying multivariate spectra [**?**] have to be further developed and adapted to the online monitoring context. Alternatively, moving window techniques can be applied to detect changes in the (partial) correlations between the factors [**?**].

4 Sliced Inverse Regression

4.1 A Dynamic Version of Sliced Inverse Regression

Restricting to linear relations as in the previous sections is not always appropriate. For this reason, Becker et al. [**?**] transfer sliced inverse regression (SIR) into the time series context to investigate possibly non-linear dependencies within multivariate time series. In its original form, SIR is a tool for dimension reduction in non-dynamic regression problems [**?**]. To explain the basic idea assume that $X = (X_1, \ldots, X_d)'$ is a set of explanatory variables (predictors), Y is a dependent variable, and ε an additional error variable. Instead of relating the whole set of predictors to Y via an unknown link function, i.e. taking $Y = g(X, \varepsilon)$, one assumes that it is sufficient to consider a lower-dimensional space, the so-called central dimension reduction (dr) subspace [**?**] $\mathcal{B}$ of dimension $r < d$, such that there is a function $f : \mathbb{R}^{r+1} \mapsto \mathbb{R}$ with

$$Y = f(\beta_1' X, \ldots, \beta_r' X, \varepsilon), \tag{7}$$

$$\text{where } \mathcal{B} = \text{ span}[\beta_1, \ldots, \beta_r] \tag{8}$$

is a space of smallest possible dimension r such that (7) is satisfied. Hence, a reduction of the regressor space from d to r dimensions is supposed to be possible. SIR then estimates $\mathcal{B}$ using information contained in an estimate of the inverse regression curve $E(X|Y)$. The basic theorem underlying SIR states that under certain regularity assumptions the appropriately standardised inverse regression curve almost surely falls into a linear transform of $\mathcal{B}$. Reversely, the space in which $E(X|Y)$ is mainly spread out yields information on $\mathcal{B}$. To identify this space, the inverse regression curve $E(X|Y)$ is roughly approximated by a simple step function, and the space itself is estimated by means of a principal component analysis of the steps. Note that estimating the function f itself is not part of SIR but has to be performed afterwards. Hence, the SIR procedure is an "intermediate" in between projection pursuit regression (yielding an estimate of f together with a reduced space of projections) and unsupervised principal component analysis of the regressor X (yielding a reduced space without considering any information on f).

Becker et al. [**?**] suggest the following modification of the original SIR procedure to incorporate time series structure: Various lagged measurements of the variables are bound together to construct the regressor space, forming higher dimensional observations. Then the original SIR method is applied to this higher dimensional regressor. Formally, let $(Y(t), X(t)')'$ denote the observation of $(Y, X')'$

at time t. Then we can search for a dr subspace of the regressor space using a modified version of equation (7):

$$Y(t) = f(\beta_1' \widetilde{X}(t), \ldots, \beta_r' \widetilde{X}(t), \varepsilon(t)),\ t \in \mathbb{Z}, \tag{9}$$

where $\widetilde{X}(t) = (X(t)', Y(t-1), X(t-1)', \ldots, Y(t-p), X(t-p)')'$ if we want to take p time lags into account. The order p has to be chosen using preliminary information or by successive fitting of increasing orders. Experience shows that for online monitoring data observed in intensive care $p = 2$ is sufficient [**?**]. Effectively, by applying this dynamic version of SIR we assume a nonlinear transfer function model

$$Y(t) = g(X(t), \ldots, X(t-p), Y(t-1), \ldots, Y(t-p), \varepsilon(t)) \tag{10}$$

where the innovations $\varepsilon(t)$ form a white noise process. Here, it is usually supposed that feedback is not present in the system, meaning that Y does not influence the explanatory variables X and that the processes $\{\varepsilon(t), t \in \mathbb{Z}\}$ and $\{X(t), t \in \mathbb{Z}\}$ are independent. Appropriate dimension reduction in this model class is crucial to get an impression of g as finding a suitable transfer function and estimating its parameters is difficult even in case of a single explanatory variable [**?**]. For more details on non-linear models see [**?**].

4.2 Application to Physiologic Time Series

Dynamic SIR as described above is a powerful exploratory tool in data situations, where we expect that there is some dependence structure between the variables without knowing it in detail. We provide a single illustrative example of applying dynamic SIR, concentrating on the relations of APD to the other vital signs. This variable is typically neglected by the physician. Therefore we are interested in knowing which combinations of the other variables provide important information on the course of APD. We use the data from the same patient as in Section 2.2 (also see [**?**]). Applying dynamic SIR to the standardized variables yields a three-dimensional reduced regressor space which is mainly spanned by APM and APS, where APD at time t depends strongly on APM at the same time, less but still clearly on APS at time t and APM at times $t-1$, $t-2$, and only slightly on all other variables except for CVP, HR and pulse.

Figure 3 depicts a graphical illustration of these findings. We connect APD to any of the other variables by a line if the simultaneous observation of the latter has a clear impact in the dimension reduction, and draw an arrow towards APD if this is true for a past observation. The strength of the relation (low, medium, high) is interpreted according to the weight the corresponding explanatory variable gets in the dimension reduction step. The exact classification done here is to some extent subjective as we use dynamic SIR as an exploratory tool only. The results are rather similar to those obtained from the partial correlation graph. Both methods identify APM and APS as the most important variables when we want to explain APD.

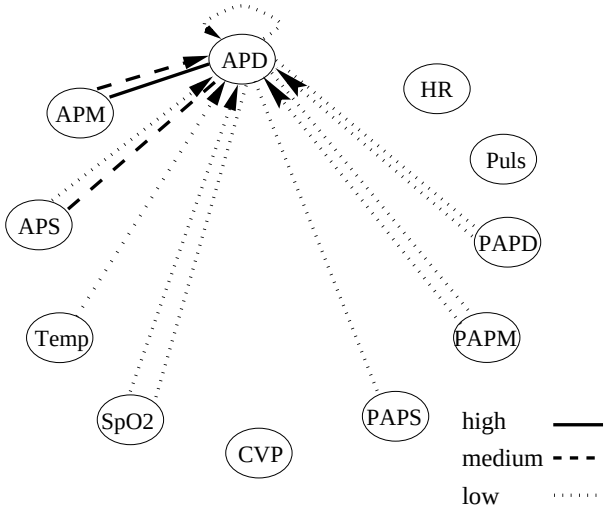

Fig. 3. Relations of APD to other vital signs for one patient derived from dynamic SIR.

Comparing Figures 1 and 3, we see that some low relations of APD to other variables identified by graphical modelling are not found by dynamic SIR and vice versa. This is due to the fact that partial correlation graphs treat both variables symmetrically, whereas dynamic SIR uses an asymmetric regression approach assuming that one of the variables can be explained by the others. For a discussion of the differences between these frameworks in a non-dynamic setting see [**?**]. For the dynamic setting we just note that dynamic SIR additionally provides some information about the dynamics since it identifies the relevant lags for the relations [**?**].

Finding a three-dimensional reduced subspace points at nonlinearities since for a linear link function only one direction is needed. This was empirically validated via simulations even in the presence of feedback [**?**]. In the clinical context, nonlinearities can be due to therapeutic interventions influencing all variables in different ways. A linear factor model as used in Section 3 can therefore only provide a local description of the relations in the steady state. This can be totally satisfactory if detection of relevant changes like intervention effects is the main goal. The possible non-linear nature of the dependence patterns could be analysed further, investigating plots of the variables against the dimension reducing directions [**?**].

5 Conclusion

Partial correlation graphs and dynamic SIR are useful tools to detect patterns of dependencies in multivariate time series by statistical data analysis. The

insights gained by these methods can be useful to improve online monitoring of vital signs since they allow an improved application of methods for dimension reduction. Either a suitable subset of important variables can be selected based on the data analysis, or the information obtained can be used to enhance methods such as principal component analysis [**?**] or factor analysis for time series if we are able to identify e.g. subsets of closely related variables. In our experience the results of the statistical data analysis agree with medical knowledge. Therefore, we expect to gain new insights when applying these methodologies to other variables, for which we have less medical knowledge so far.

In particular, we have found evidence that specific partial correlation patterns represent specific clinical states in the critically ill as the findings presented here are in good agreement with medical knowledge on the causes and symptoms of distinct states. This can be useful for reaching deeper insights into the causes of clinical complications. In view of the high sampling rates of modern equipment partial correlation analysis could even be useful to detect these complications online in the time critical situations on the intensive care unit.

ACKNOWLEDGEMENTS

We appreciate the suggestions by a referee as they made the paper more readable. The financial support of the Deutsche Forschungsgemeinschaft (SFB 475, "Reduction of complexity in multivariate data structures") is gratefully acknowledged.