Graphical Models for Multivariate Time Series from Intensive Care Monitoring

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

[1] Department of Statistics, University of Dortmund, 44221 Dortmund, Germany

[2] Surgical Department, Community Hospital Dortmund, 44137 Dortmund, Germany

SUMMARY

In critical care extremely high dimensional time series are generated by clinical information systems. This yields new perspectives of data recording and also causes a new challenge for statistical methodology. Recently graphical correlation models have been developed for analysing the partial associations between the components of multivariate time series. We apply this technique to the hemodynamic system of critically ill patients monitored in intensive care. We appraise the practical value of the procedure by reidentifying known associations between the variables. From separate analyses for different pathophysiological states we conclude that distinct clinical states can be characterised by distinct partial correlation structures.

* *Corresponding author's email address:* gather@statistik.uni-dortmund.de.

Key words: Multivariate time series; Partial spectral coherence; Graphical interaction model; Substantive research hypotheses; Hemodynamic variables.

1 Introduction

Graphical interaction models have become an important tool for investigating and modeling multivariate data. Here, hypothetical relations between the random variables are displayed in graphical form, where the vertices symbolise the variables and the edges illustrate the associations between them.

In most applications the statistical meaning of association is conditional dependence, i.e. a missing edge indicates *conditional independence* of the corresponding variables given all the remaining variables. This conditional structure is considered because in multivariate data usually a multitude of relations exists, but many of them are indirect, i.e. they are induced by others. In such cases the conditional structure may be less complex than the marginal one. The essential associations within the data become obvious after controlling the effects of other variables (e.g. confounders) which influence the variables considered. Such effects possibly induce spurious, indirect associations. Wermuth (1976) uses the concept of conditional independence to find parsimonious models in case of multivariate Gaussian or multinomial distributions while Dawid (1979) discusses conditional independence as a basic tool for statistical inference.

According to Wermuth and Lauritzen (1990) and Cox and Wermuth (1993) graphical models have some advantages in comparison to e.g. linear structural equation models, which are also used for generating and exploring substantial new research hypotheses, especially in econometrics and the social sciences (Goldberger and Duncan, 1973). The graphical visualisation is a simple and very helpful illustration of the model structure, which can also clarify structural differences between several models and identify equivalences among them (Lauritzen and Wermuth, 1989). Recently, Tritchler (1999) pointed out some examples of statistical reasoning in medical applications by directed graphs. The monographs by Whittaker (1990), Lauritzen (1996) and Cox and Wermuth (1996) give broad reviews on the topic.

Dahlhaus (2000) extends the concept of graphical models to multivariate time series. For describing the interactions between the components of a multivariate time series he uses the partial spectral coherence, which is a measure for the linear dependence between two components of a multivariate time series after the linear effects of the other series have been removed (Brillinger, 1981). By combining this tool with the concept of graphical models he proposes a method to detect (partial) associations between the variables of a multivariate time series in view of time-lagged dependences.

We apply this technique to time series of hemodynamic variables from randomly selected intensive care unit (ICU) patients. By comparing the results with known physiological relations we can appraise the value of this new method. Moreover, we make use of a given classification of observation periods for each patient into several clinical states to evaluate whether these states can be characterised also by different partial correlation structures.

We proceed as follows. In Section 2 we introduce graphical interaction models and review the idea of their extension to multivariate time series. In Section 3 we give a basic application of this method to the hemodynamic system and in Section 4 we use graphical models to distinguish between distinct clinical states of patients. Some conclusions are given in Section 5.

2 Graphical Interaction Models

A *graph* $G = (V, E)$ consists of a finite set of *vertices* V and a set of *edges* $E \subset V \times V$, that are ordered pairs of vertices. 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, then an undirected edge (*line*) is drawn. Both types of edges have different interpretations in the context of graphical models. An arrow specifies a direction for the influence between the variables, while a line stands for a symmetrical interaction. In more sophisticated graphical models further types of edges may be used, e.g. dashed and solid edges, and the vertices may have different forms, too (Lauritzen and Wermuth, 1989, Cox and Wermuth, 1993, Didelez and Pigeot, 1998).

In this paper we focus exclusively on *undirected graphs*, which have undirected edges only. Here, two variables a and b are conditionally independent given all other variables if they are not connected via an edge, i.e. $(a, b) \notin E$. This is the *pairwise Markov property* for undirected graphical models. Two variables a and b are related somehow (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 joint density of the variables has a positive continuous density w.r.t. a product measure, the pairwise Markov property is equivalent to the *global Markov property*. This property says, that two sets of variables A and B are conditionally independent given 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 associated if the effects of the separating subset C are controlled. The subset C may contain less than all the remaining variables which allows to identify important variables more clearly.

To exemplify the concept of graphical models we consider some of the variables which are analysed empirically in section 3. Take a look at the vital signs heart rate (HR), arterial mean pressure (APM), pulmonary arterial mean pressure (PAPM), central venous pressure (CVP), blood temperature (Temp) and pulsoximetry (SpO2). Figure 1 shows known physiological associations between these variables. There are associations between the distinct kinds of blood pressure (which are supposed to be positive) as well as between each of them and the heart rate (supposed to be negative). In contrast no association is supposed between pulsoximetry or blood temperature and any of the other variables.

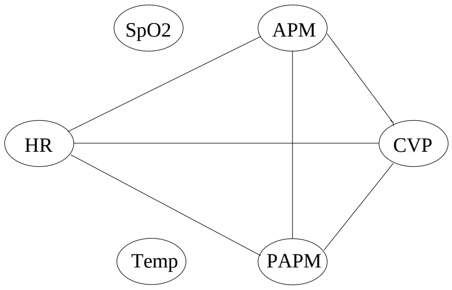

Figure 1: Graphical model for the hemodynamic system based on expert knowledge.

Having a closer look at the graph we also notice that there is no indirect association as any connected vertices are connected directly by an edge. Eventually this simple physiological model is correct and there really are no induced associations. On the other hand it is generally very difficult to distinguish between direct and induced associations by experience only and without a careful data analysis. One possibility might be to take the strength of the associations into account. We would like to give a naive, purely hypothetical example for this. If medical experience told us, that the positive associations of PAPM to both APM and CVP were much stronger than the association between APM and CVP, then we might conclude that the latter one is a consequence of the stronger relations. In that case the edge between APM and CVP in the graphical model in Figure 1 could be removed and this would illustrate the indirect character of the association between the two variables via PAPM and HR. More exactly, missing of this edge would mean that given the temporal course of PAPM and HR the course of APM does not add any further knowledge to explain the course of CVP and vice versa. Hence, there are parallels between the specification of a graphical model and the selection of variables in a regression model (Whittaker, 1990).

Indirect associations can result from subsequent direct influences. Continuing with our consideration we could ask if an increase of PAPM typically is caused by an increase of CVP, and then the increase of PAPM might result in an increase of APM. This would be an example of a causal mechanism between multiple variables. In graphical models such mechanisms can adequately be described by directed edges. Searching for such causal mechanisms is not the scope of this paper as we lack appropriate methods which have proven their suitability for multivariate time series data. Cox (1992) gives some general comments and outlines different views with respect to the problem of recognizing causality by data analysis. Granger causality (Granger, 1969) is a frequently used concept for variables measured subsequently in time since it gives an intuitive definition of causality by means of improved predictability of a variable when past observations of the causative variable are included in the model. Here however, analyzing the symmetric partial association structure between multivariate time series variables is used as a first step to get a better understanding of the data generating mechanism.

When analysing the associations between the vital signs of an individual we should actually consider the time series structure of the measurements, which are not independent of course. Dahlhaus (2000) defined graphical interaction models for multivariate time dependent data by using the partial spectral coherence, which is a well-known tool for analyzing the second-order properties of multivariate time series (Brillinger, 1981, section 8.3). Dahlhaus calls such a model "conditional correlation graph" since it achieves a graphical representation of the *linear relations* between the components of a multivariate time series. Hence, a missing edge implies a weaker property in this context than in conditional independence graphs, where a missing edge means conditional independence. Both concepts are equivalent under the assumption of joint multivariate normality of the variables.

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

$$\gamma_{ab}(h) = Cov(X_a(t+h), X_b(t)),\ 1 \leq a, b \leq k;\ h \in \mathbb{Z}$$

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

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

(see Brillinger, 1981, p. 232ff). 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 analogy to graphical models for multivariate data it is useful to distinguish between direct and induced linear relations between two series $X_a(t)$ and $X_b(t)$. Therefore the linear effects of the remaining components of $\mathbf{X}(t)$ on $X_a(t)$ and $X_b(t)$ have to be eliminated. Let $\mathbf{Y}(t) = (X_j(t), j \neq a, b)$, $t \in \mathbb{Z}$, denote the series of the other components. The optimal $\mu_a \in \mathbb{R}$ and the optimal univariate filter $d_a(h), h \in \mathbb{Z}$, have to be determined, such that the quadratic distance

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

is minimal. Let $\epsilon_a(t)$ be the residual

$$\epsilon_a(t) = X_a(t) - \mu_a - \sum_h d_a(h)\mathbf{Y}(t-h)\,.$$

In the same way we calculate $\epsilon_b(t)$ from $X_b(t)$. The *partial cross-spectrum* between $X_a(t)$ and $X_b(t)$ can then be defined as the cross-spectrum between $\epsilon_a(t)$ and $\epsilon_b(t)$

$$f_{X_a X_b \cdot \mathbf{Y}}(\lambda) = f_{\epsilon_a \epsilon_b}(\lambda)\,,$$

while the *partial spectral coherence* is a standardization hereof

$$R_{X_a X_b \cdot \mathbf{Y}}(\lambda) = \frac{f_{X_a X_b \cdot \mathbf{Y}}(\lambda)}{f_{X_a X_a \cdot \mathbf{Y}}(\lambda) f_{X_b X_b \cdot \mathbf{Y}}(\lambda)}\,. \tag{1}$$

In the same way the (partial) cross-spectrum between two vector time series can be defined. For assessing the properties of these complex-valued functions often their squared absolute value and their argument $\phi_{ab}(\lambda)$ or $\phi_{ab\cdot\mathbf{Y}}(\lambda)$ respectively, the *(partial) phase-spectrum*, is analysed, where $\phi_{ab}(\lambda)$ comes from the Eulerian representation

$$f_{ab}(\lambda) = |f_{ab}(\lambda)| \exp\{i\phi_{ab}(\lambda)\}\,.$$

The argument $\phi_{ab}(\lambda)$ (respectively $\phi_{ab\cdot\mathbf{Y}}(\lambda)$) of the (partial) cross-spectrum is called *(partial) phase-spectrum.*

Brillinger (1981, Theorem 8.3.1) shows that the partial cross-spectrum can be calculated using the formula

$$f_{X_a X_b \cdot \mathbf{Y}}(\lambda) = f_{X_a X_b}(\lambda) - f_{X_a \mathbf{Y}}(\lambda)[f_{\mathbf{Y}\mathbf{Y}}(\lambda)]^{-1} f_{\mathbf{Y} X_b}(\lambda)\,, \tag{2}$$

where the components of the vectors $f_{X_a \mathbf{Y}}(\lambda)$ and $f_{\mathbf{Y} X_b}(\lambda)$ and the matrix $f_{\mathbf{Y}\mathbf{Y}}(\lambda)$ are cross-spectra between the corresponding variables.

In a *conditional correlation graph* of a multivariate time series the vertices $a = 1, \ldots, k$ are the components of the time series and an edge between two vertices a and b is omitted whenever $f_{X_a X_b \cdot \mathbf{Y}}(\lambda) = 0$ for all frequencies $\lambda \in \mathbb{R}$. Hence, a missing edge indicates that the partial linear relation between the variables is zero. Dahlhaus (2000) proved the global Markov property for this kind of graphical models under the assumption that the spectral matrix is everywhere regular.

In applications, in order to investigate research hypotheses by empirical analysis of multivariate time series data, one can first estimate the cross-spectra from the data and then use versions of equations (1) and (2) for the empirical functions to estimate the partial spectral coherences. Thereafter a decision has to be made on whether the partial spectral coherence may equal zero because sampling variability always causes estimates to be distinct from zero. Dahlhaus and Eichler (2000) developed the program "Spectrum" which estimates the cross-spectrum by a nonparametric kernel estimator. In addition, the program 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. Thus an approximate 5%-test for the hypothesis of partial uncorrelatedness of two variables can be performed by comparing the estimated partial spectral coherence with this bound.

3 Application to time series from intensive care

In ICUs of modern hospitals more than 2000 physiological variables, laboratory data, device parameters etc. are observed for each critically ill patient in the course of time. The appropriate analysis and online monitoring of this enormous amount of dynamic data is essential for suitable bedside decision support in time critical situations (Bauer, Gather and Imhoff, 1999). Distinguishing between direct and induced associations between the observed variables and primary and secondary consequences of medical interventions is however difficult from experience only.

As we have pointed out in section 2, the associations between physiological variables may differ in strength. Figure 2 shows the opinion of an expert w.r.t. the associations between the same variables as shown in Figure 1. In the following, we give a data analysis based on classification criteria described below and compare the results to the expert's opinion.

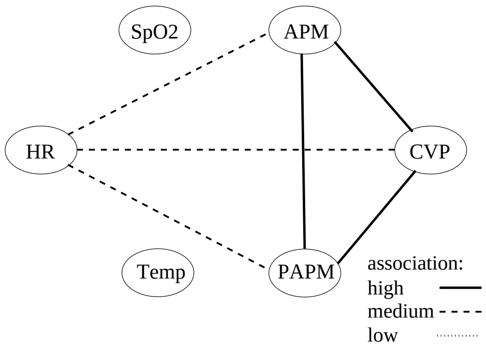

Figure 2: Graphical model for the hemodynamic system based on expert knowledge.

The data investigated here was acquired with a unix-based clinical information system (CIS) on the surgical intensive care unit of the Community Hospital Dortmund. This CIS (Eclipsys Emtek Continuum 2000, Version 4.1M3) has been in routine clinical use since 1992 and provides a high-quality acquisition of patient data at the bedside (Imhoff, 1995). Numerous variables from bedside medical devices can automatically be charted at 1-minute time intervals. In this way online-monitoring data was acquired in one minute intervals from 19 consecutive critically ill patients with pulmonary artery catheters for extended hemodynamic monitoring amounting to 102132 sets of observations. Time series from ten hemodynamic variables (heart rate HR, arterial diastolic pressure APD, arterial systolic pressure APS, arterial mean pressure APM, pulmonary arterial diastolic pressure PAPD, pulmonary arterial systolic pressure PAPS, pulmonary arterial mean pressure PAPM, central venous pressure CVP, blood temperature Temp, pulsoximetry SpO2) are included in the analysis. We use the program "Spectrum" of Dahlhaus and Eichler (2000) to estimate the partial spectral coherences between these variables.

We begin with a simple preliminary example to explain how we can derive graphical models from the partial spectral coherences. Figure 3 displays the squared estimated partial spectral coherence (below the diagonal) and the estimated partial phase spectra (above the diagonal) between arterial and pulmonary arterial pressures of one patient. Since we focus here on the partial spectral coherence, only the plots below the diagonal are considered. Dahlhaus, Eichler and Sandkühler (1997) used the partial phase spectrum for deriving the direction of the influence between two parameters in point processes. For time series this is not an established method yet. In our case there were no visible directions. The dashed lines in the plots represent the bounds of an approximate joint 5%-test of the hypothesis that the partial coherence equals zero for all frequencies. The strong associations between both the arterial and the pulmonary arterial pressures become obvious immediately, while the partial spectral coherence is close to zero between these groups in all cases.

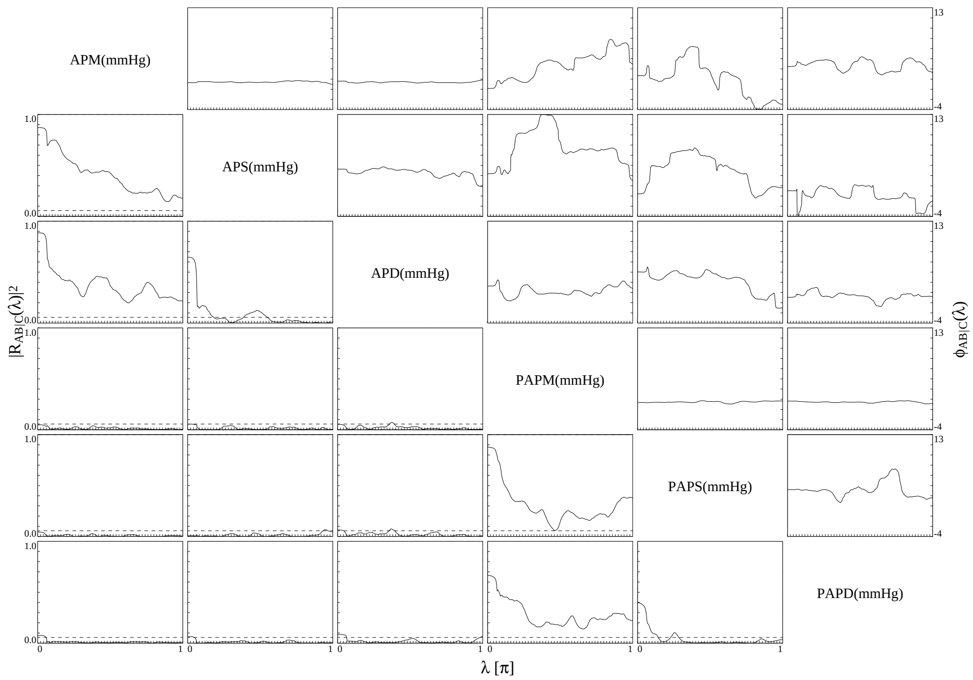

Figure 3: Typical example of the partial spectral coherences (below the diagonal) and partial phase spectra (above the diagonal) between the arterial and pulmonary arterial pressures of one patient. The partial correlations between APM, APS and APD as well as between PAPM, PAPS and PAPD are evident from the partial spectral coherence.

Dahlhaus (2000) suggests to construct a graphical correlation model for a multivariate time series by comparing the partial spectral coherence to the approximate test bound. This allows a "yes - no" judgement only, while it is well-known that different associations between physiological variables may have distinct strengths. We decide to classify the associations in high, medium, low and zero correlation on the basis of the area under the estimated partial spectral coherence $R_{X_a X_b \cdot \mathbf{Y}}$. This area can be measured by the partial mutual information (Granger and Hatanaka, 1964) between the time series $\{X_a\}$ and $\{X_b\}$, which is defined by

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

(Brillinger, 1996) or by variants of this. The graphical correlation model which results from the partial spectral coherences shown in Figure 3 is presented in Figure 4.

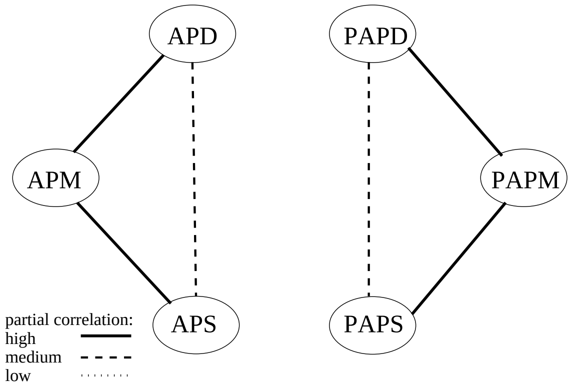

Figure 4: Graphical correlation model for arterial and pulmonary arterial pressures.

In a first step we try to analyse the associations between all vital signs mentioned above. However, an analysis of the partial spectral coherence between many variables affords a huge number of observations. Since we want to get a general impression, the estimated partial coherences are averaged over all patients. The resulting graphical correlation model is displayed in Figure 5. As expected strong correlations exist between the systolic, diastolic and mean arterial pressure (APS, APD and APM), and between the systolic, diastolic and mean pulmonary arterial pressure (PAPS, PAPD, PAPM). Additionally, a strong association between the heart rate and the arterial pressures becomes visible, while the system of arterial pressures is related to the system of pulmonary arterial pressures via the blood temperature only. The strong association between the blood temperature and SPO2 in not supposed to exist in reality, and the other associations are small. In fact, the unexpected association between Temp and SPO2 turned out to have been caused by a systematic error of the measuring instruments which could actually be detected by our analysis.

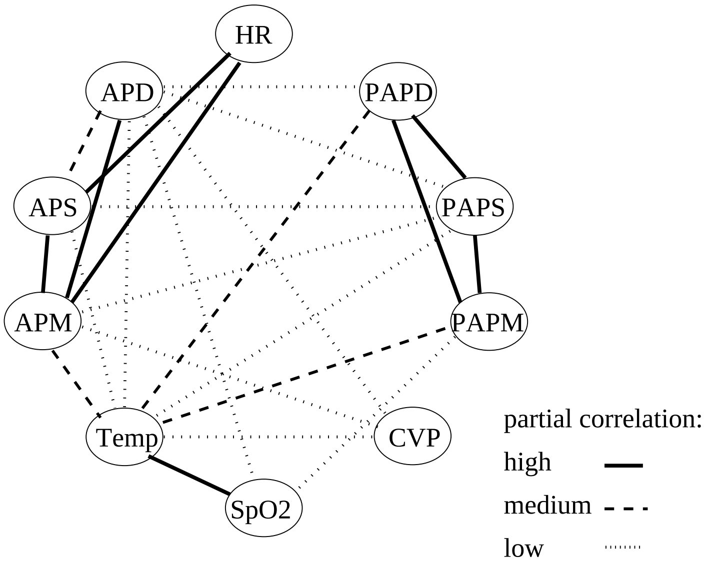

Figure 5: Graphical correlation model for all considered vital signs found by averaging partial coherences over all patients.

At first sight the non-existence of direct connections between the arterial and the pulmonary arterial system is not in the line with physiological knowledge. The reason for this unexpected finding is that the data analysis is accomplished by conditioning on *all* other variables. Hence, for analysing a possible association between APM and PAPM we condition on all remaining variables including APS, APD, PAPS and PAPD. However, APM is a (slightly non-linear) transform of APS and APD, and PAPM is a (slightly non-linear) transform of PAPS and PAPD. Since we eliminate the linear effects of the other variables, the remaining variability for both APM and PAPM is very low. The global Markov property does not apply since the corresponding regularity condition is not guaranteed. This means that we cannot identify known associations when a bunch of variables is included which almost measure the same quantity. In consequence, systolic and diastolic pressures are dropped in the following calculations, which is in line with medical reasoning. In this way a set of 'important variables' consisting of HR, SPO2, APM, PAPM, CVP and Temp is retained.

For these variables all observations without missing values are included in the analysis for each patient. Between 366 and 10929 observations can be used for the estimation of the partial coherences for each person. A representative example for the associations between the remaining variables is shown in the partial spectral coherence plot in Figure 6. The following list gives typical examples of partial correlations with different strengths.

- zero partial correlation: e.g. HR–SpO2
- low partial correlation: e.g. SpO2–PAPM
- medium partial correlation: e.g. HR–PAPM
- high partial correlation: e.g. HR–APM

Graphical models are generated using distinct types of edges corresponding to the classification of the partial spectral coherences described above. Using the clinical data we test whether this method can reliably identify

- known strong associations, e.g. between systolic, mean and diastolic blood pressures, and
- known likely associations, e.g. between HR, PAP, and CVP.

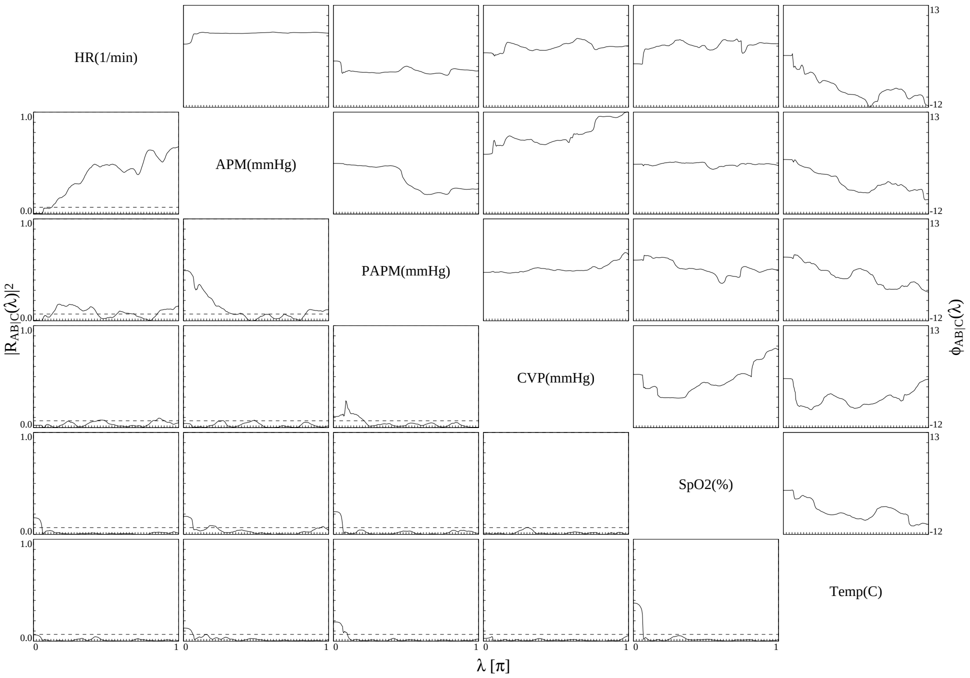

Figure 6: Partial spectral coherence and partial phase spectra of vital signs.

For all patients the known strong associations can be reidentified. The strongest partial correlations are observed between mean arterial and intrathoracic pressures, i.e. central venous and pulmonary pressures. Strong correlations can also be detected between heart rate and mean arterial and pulmonary arterial pressures. The weakest partial coherences are between body temperature and all other vital signs. Figure 7 provides the graphical correlation model resulting from the partial spectral coherences shown in Figure 6.

4 Distinguishing clinical states

Distinct clinical states such as pulmonary hypertension, septic shock, congestive heart failure and vasopressor support are accompanied by different pathophysiological responses of the circulatory system. These changes may be supposed to result in differ-

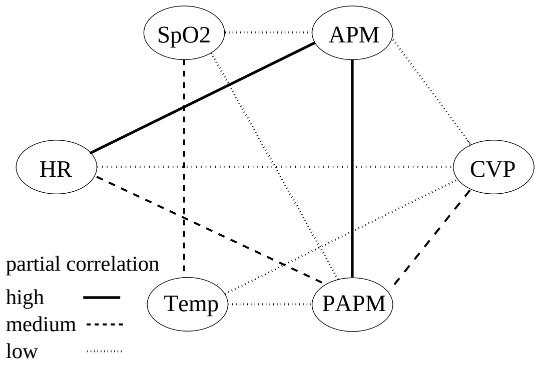

Figure 7: Graphical correlation model for vital signs.

ences in the interactions between the vital signs also. To give an example, Figure 8 illustrates the expert's opinion w.r.t. the associations when pulmonary hypertension is the clinical state. 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 diminishes the interactions between PAPM and APM as changes in PAPM will have a less than normal effect on left ventricular preload. In consequence, the expected associations of vital signs show another picture for the state of pulmonary hypertension than the associations under normal physiological conditions as displayed in Figure 2.

In the following we investigate, whether graphical correlation models can detect differences in the status of the circulatory system. Therefore, in a next step, for each patient the pathophysiological state is determined for every time point as scored from the medical record. Then the partial coherences between the vital signs are estimated separately for each of the states and for each individual patient.

Figure 9 shows the associations corresponding to the estimated partial coherences for each clinical state and for each patient. Although the number of samples is small,

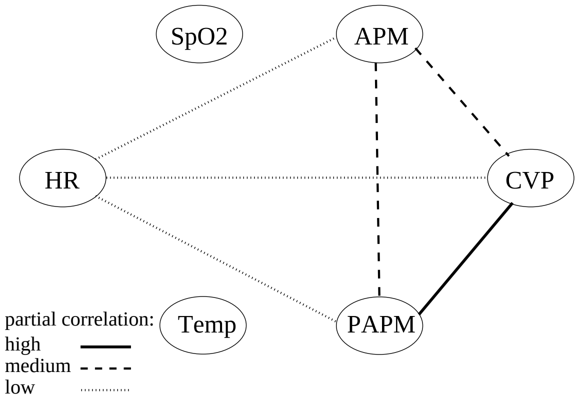

Figure 8: Graphical correlation model for pulmonary hypertension, expert opinion.

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 correlations could be found between heart rate and blood pressures as well as between intrathoracic and arterial pressures, the status of pulmonary hypertension is predominantly associated with strong correlations between intrathoracic and arterial pressures. For the clinical status of vasopressure support there are strong partial correlations between APM and PAPM, while some of the other associations are found to be weaker than for other states. 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. Hence, there are strong positive interactions between APM and PAPM, while the influence of CVP on the other variables is reduced. The strong partial correlation between APM and PAPM in the status of congestive heart failure can be explained by a failure of the left ventricle. This causes a decrease of APM as the heart is not able to push the blood forward (forward failure), and in consequence there is a build-up in front of the left ventricle and thus an increase of PAPM (backward failure).

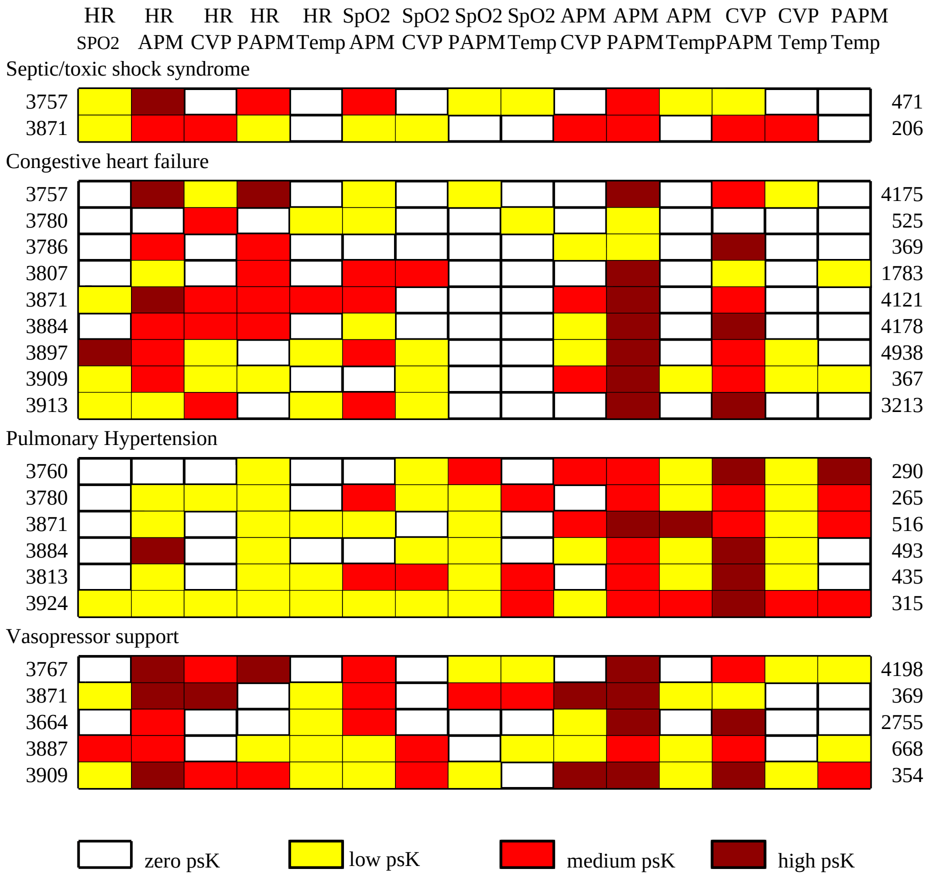

Figure 9: Classification of the partial correlation structure for distinct clinical states.

Figure 10 summarizes the impressions gained in Figure 9. Graphical correlation models for distinct clinical states based on data analysis are contrasted with the associations expected by expert knowledge for a non-critical state. The differences are in accordance with physiological knowledge. Further studies are projected to validate these preliminary findings with larger groups of patients.

5 Conclusion

We found that graphical correlation analysis based on partial coherences can reliably detect known physiological association patterns in hemodynamic time series. Our

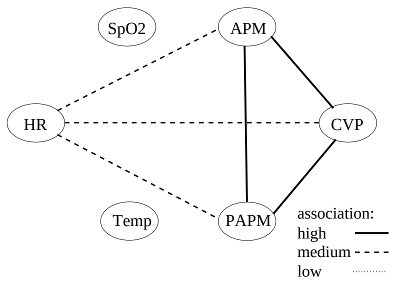

Non critical state (expert knowledge)

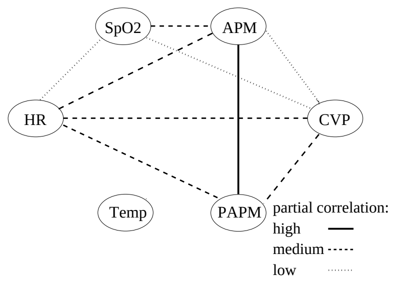

Congestive heart failure

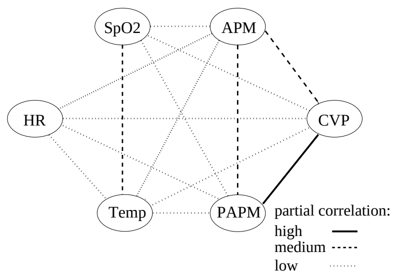

Pulmonary hypertension

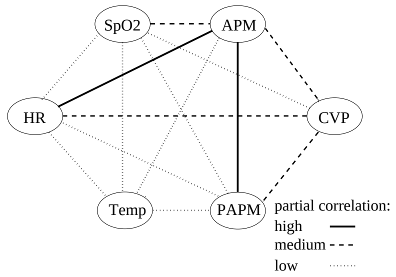

Vasopressure support

Figure 10: Graphical correlation models for distinct clinical states

results confirm that the method actually detects known physiological dependences between monitoring variables. The method can be considered as a model selection approach where the graph represents an appropriate model for the time series.

The insights gained by graphical partial correlation analysis can also be useful to improve the online-monitoring of vital signs, since such an analysis allows the intelligent application of methods for dimension reduction. Either a suitable subset of important, i.e., influential variables can be selected from the graphs or information w.r.t. the conditional association structure can be deduced from the graph to enhance methods such as principal component analysis and factor analysis for time series, which are accomplished via spectral analysis, too (Brillinger, 1981). Multiblock principal component analysis has been suggested to monitor high-dimensional time series (MacGregor et al., 1994). Here, the variables are organized in subsets for which principal component analysis is applied separately. Such subsets of closely related variables can also be found from a graphical model for past data.

Special correlation patterns may even represent specific pathophysiological states in the critically ill. This might be useful for getting further insights in the causes of clinical complications as well as for detecting such complications by additional data analysis.

Although it is too early for a final judgement, graphical partial correlation analysis seems to support the analysis of correlations in multivariate physiological time series, and further extensive work has to be done. As the final statistical analysis results in simple graphs, interpretation of the partial correlation patterns can be accomplished by physicians without further statistical knowledge.

ACKNOWLEDGEMENTS

The financial support of the Deutsche Forschungsgemeinschaft (SFB 475, "Reduction of complexity in multivariate data structures") is gratefully acknowledged.

REFERENCES

Bauer, M., Gather, U. and Imhoff, M. (1999). The Identification of Multiple Outliers in Online Monitoring Data. Preprint 29/1999, Department of Statistics, University of Dortmund, Germany.

Brillinger, D.R. (1981). *Time Series. Data Analysis and Theory.* San Francisco: Holden Day.

Brillinger, D.R. (1996). Remarks Concerning Graphical Models For Time Series And Point Processes. *Revista de Econometria* **16**, 1-23.

Cox, D.R. (1992). Causality: Some Statistical Aspects. *Journal of the Royal Statistical Society*, Series A **155**, 291-301.

Cox, D.R. and Wermuth, N. (1993). Linear dependencies represented by Chain Graphs. *Statistical Science* **8**, 204-283 (with discussion).

Cox, D.R. and Wermuth, N. (1996). *Multivariate Dependencies.* London: Chapman & Hall.

Dahlhaus, R. (2000). Graphical interaction models for multivariate time series. *Metrika*, to appear.

Dahlhaus, R. and Eichler, M. (2000). Spectrum. This program is available at http://www.statlab.uni-heidelberg.de/projects/graphical.models/

Dahlhaus, R., Eichler, M. and Sandkühler, J. (1997). Identification of synaptic connections in neural ensembles by graphical models. *Journal of Neuroscience Methods* **77**, 93-107.

Dawid, A.P. (1979). Conditional Independence in Statistical Theory. *Jornal of the Royal Statistical Society* B **41**, 1-31.

Didelez, V. and Pigeot, I. (1998). Maximum Likelihood Estimation in Graphical Models with Missing Values. *Biometrika* **85**, 960-966.

Goldberger, A.S. and Duncan, O. (eds.) (1973). *Structural Equation Models in the Social Sciences.* New York: Seminar Press.

Granger, C.W.J. (1969). Investigating causal relations by econometric models and cross-spectral methods. *Econometrica* **37**, 424-438.

Granger, C.W.J. and Hatanaka, M. (1964). *Spectral Analysis of Economic Time Series.* Princeton: Princeton Press.

Imhoff, M. (1995). A clinical information system on the intensive care unit: dream or night mare? In *Medicina Intensiva 1995*, J.A.G. Rubi (ed), XXX. Congreso SEMIUC, Murcia.

Lauritzen, S.L. (1996). *Graphical Models.* Oxford: Clarendon Press.

Lauritzen, S.L. and Wermuth, N. (1989). Graphical Models for Associations between variables, some of which are qualitative and some quantitative. *Annals of Statistics* **17**, 31-57.

MacGregor, J.F., Jaeckle, C., Kiparissides, C. and Koutoudi, M. (1994). Process monitoring and Diagnosis by Multiblock PLS methods. AIChE J. **40**, 826-838.

Tritchler, D. (1999). Reasoning About Data With Directed Graphs. *Statistics in Medicine* **18**, 2067-2076.

Wermuth, N. (1976). Model search among multiplicative models. *Biometrics* **32**, 253-263.

Wermuth, N. and Lauritzen, S.L. (1990). On Substantive Research Hypotheses, Conditional Independence Graphs and Graphical Chain Models. *Journal of the Royal Statistical Society, Series B* **52**, 21-50.

Whittaker, J. (1990). *Graphical models in applied multivariate statistics.* Chichester: John Wiley & Sons.