On-Line Identification of a Patient-Disease Model for Mechanical Ventilation

Anders Hansson*
Information Systems Laboratory
Durand 101A
Stanford University
Stanford, CA 94305–4055
USA

Silvia Miksch
Institute of Software Technology
Vienna University of Technology
Resselgasse 3/E 188
A–1040 Vienna
Austria

Abstract

Monitoring and therapy planning in real-world environments highly depend on good patient-disease models. The improvement of the technical equipment in modern intensive care units enables a huge number of on- and off-line data, which results in an information overload of the medical staff. Additionally, the underlying medical structure-function models are poorly understood or not applicable due to incomplete knowledge. We have developed an on-line identification scheme, which utilizes a priori knowledge as well as on-line measurements to identify the parameters of a disease model for mechanically ventilated newborn infants. The scheme benefits from an exponential weighting function to classify more recent measurement values as more important. We have evaluated our identification scheme with real medical data sets showing the benefits and drawbacks of our approach.

1 Introduction

How do we obtain good procedures for patient care? There are many answers to this question. However, when the therapy consists of decisions for how to tune the knobs of a machine connected to the patient based on on-line measurements from the patient, it is natural to think of the problem as a multi-stage decision problem. Depending on your background you may attack this problem e.g., using ideas from Artificial Intelligence (AI), and/or Control Theory, which both have their roots in Cybernetics [Wiener, 1948]. The main differences between the two approaches is how the measurement data is mapped into models, what types of models are used, and how the decisions are made based on these models.

*This author has been partly sponsored by the Swedish Research Council for Engineering Sciences under contract No. 95–838.

What is then the most appropriate type of model? This highly depends on

- *What is the model going to be used for?* In case the model is going to be used for qualitative reasoning it may not be necessary to have as a detailed model as may be necessary if the model is going to be used to implement a quantitative Kalman-filter or Wiener-predictor.
- *How accurate can the model be made?* This usually depends on the quality of the measured data. The more information is present in the data the more detailed the model can be made. This often depends on what experiments can be made, which in medical applications sometimes are constrained by safety requirements for the patient.
- *How time-consuming is it to obtain the model?* This often depends on what technical equipment is available to record the data and to process it. It also depends on how sophisticated algorithms are used for processing the data.

The latter two considerations usually imply that only certain sub-sets of what need to be modeled can be made detailed, e.g., in the form of differential equations or difference equations. How these models connect to one-another can then be modeled in a qualitative way applying particular approaches from AI, e.g. knowledge-based techniques and qualitative reasoning. It should be made clear that the challenge from medicine on computations for data analysis and therapy planning is much greater that the challenges from many other fields [Bellman, 1983]. Hence it is not surprising that a multi-disciplinary approach is a good idea.

The scope of this paper is to show how system identification techniques, see e.g. [Johansson, 1993], which have their roots in Control Theory, can be used to identify a so called "patient-disease model" for mechanical ventilation. The model is in terms of a difference equation that describes the real world with some degree of accuracy. This type of models has been used extensively

in medical applications of mathematics [Bellman, 1983]. It is especially useful in case Control Theory is going to be used to synthesize the therapy planning. However, it can also be used to derive qualitative models suitable for AI approaches. For example, a qualitative simulation is based on the derived qualitative descriptions to predict the outcome of therapeutic actions; or the qualitative descriptions are used to structure skeletal plans, in the context of medicine named clinical protocols, which are used to support plan generation and execution. The experience gathered by physicians can be expressed as inequality constraints on the parameters in the model. The only thing that remains to do in order to obtain a good patient-disease model is to decide what parameter values are the correct ones for a certain patient. The idea is to do this based on on-line measurements by solving a convex optimization problem at each sample time. The quality of the obtained model highly depends on the quality of the measurement data. System identification with no constraints or just equality constraints on the parameters has been treated in textbooks [Johansson, 1993], but the case with more complicated constraints, such as inequality constraints, is still an active research area, see e.g., [Boyd *et al.*, 1994]. The data used in the identification scheme is weighted exponentially in time in order to take into account that more recently observed data are more important for the model than old data. This is a standard approach in system identification.

The remaining part of the paper is organized as follows. In Section 2 background material and related approaches are described. In Section 3 the model will be given and rewritten to fit into the identification scheme to be described in Section 4. Then in Section 5 the identification scheme will be evaluated on real data, and in Section 6 concluding remarks and perspectives on future research will be given.

2 Background and Related Approaches

Medical monitoring and therapy planning at modern intensive care units (ICUs) have been refined by the technical improvement of the equipment. Nevertheless, the care of critically ill patients with respect to real-world medical environments entails non-trivial data analysis problems: first, the number of continuously and discontinuously assessed data acquired by different devices creates a rising information-management problem at ICUs. The available data occur at different observation frequencies and in various types such as qualitative and/or quantitative data. Second, the underlying medical structure-function models—the patient-disease models—are poorly understood or not applicable due to the incomplete knowledge and complexity of the pathophysiological mechanisms. Third, the projections of various signals are context-dependent, e.g., according to the degree of a signal's abnormality, and they are difficult to represent in a formal way. Therefore it is not possible to rely only on classical theories for data analysis [Avent and Charlton, 1990; Kay, 1993].

However, the need to retrieve adequate patient-disease models to treat critically ill patients is evident. During the last years, different approaches were proposed to derive domain-dependent models in general [Hamscher *et al.*, 1992]. Fewer encouraging approaches are available for identification of durative models, e.g., the temporal-abstraction module in the *M-HTP* project [Larizza *et al.*, 1992], the temporal resource management in the *Guardian* project [Hayes-Roth *et al.*, 1992], the trend detecting mechanism based on trend templates in the $TrenD_x$ project [Haimowitz *et al.*, 1995], the temporal-abstraction module in the *VIE-VENT* project [Miksch *et al.*, 1996], the *RÉSUMÉ* project [Shahar and Musen, 1996], and the *T-IDDM* project [Bellazzi *et al.*, 1996]. Currently, the problem of automating the data analysis has grown steadily under the labels knowledge discovery in databases (KDD) and data mining [Fayyad *et al.*, 1996]. Nevertheless, these approaches are based on the assumption that there exists a particular domain model, which only need to be customized for the particular purpose. This assumption is not always applicable in a real-world environment.

We are particularly interested in the treatment of mechanically ventilated newborn infants. Mechanical ventilation has greatly contributed towards the improvement of the mortality and morbidity of prematurely newborn infants [Goldsmith and Karotkin, 1993]. Enhanced knowledge about the pathophysiological mechanisms of barotrauma and oxygen toxicity led to the development of patient-tailored strategies of mechanical ventilation and helped to reduce harmful side effects of respirator therapy. However, the bulk of available data and the vague definitions of causal and functional dependencies about the interactions of signals over a period of time still cause problems to retrieve an appropriate durative patient-disease model.

Knowledge-based system technology may appropriately represent and organize the practical and theoretical knowledge of experienced specialists [Shortliffe, 1991]. During the past decade, several projects were initiated to support clinicians with their information processing needs [Uckun, 1993]. These systems range from simple, intelligent alarming systems to sophisticated systems for knowledge-based monitoring and therapy management. The pioneering work was the Ventilator Manager (VM), [Fagan *et al.*, 1980], which was designed to manage post-surgical mechanically ventilated patients. On the one hand, there are applications which assist patients with adult respiratory distress syndrome (ARDS),

like VM [Fagan *et al.*, 1980], VentPlan [Rutledge, 1995], and COMPAS [Sittig *et al.*, 1990]. On the other hand, projects were developed to support the treatment of newborn infants, like VMS [Boyarsky, 1987] and SIMON [Uckun *et al.*, 1993]. From the practical point of view, the usability of all these systems is limited concerning their data used and their reaction time. They mainly used invasively determined blood gas measurements for monitoring and therapy planning, which are only discontinuously and infrequently available. Such approaches are applicable to adults, but are not suitable for neonates. Newborn infants still have premature organs and the treatment needs to be based on more frequently and continuously sampled data, namely the trans-cutaneously assessed measurements, to guarantee immediately in-time reactions to life-threatening health conditions of the infants.

Developing VIE-VENT [Miksch *et al.*, 1996], an open-loop, monitoring and therapy planning system for mechanically ventilated newborn infants, we tried to overcome these limitations. VIE-VENT was especially designed for practical use under real-time constraints at neonatal intensive care units (NICUs). Experiences from the development of the VIE-VENT system [Miksch *et al.*, 1996] indicated non-trivial data analysis problems. However, the performance and the acceptance of VIE-VENT could be improved significantly by incorporating a more sophisticated on-line identification of the patient-disease model, in particular, a model for mechanically ventilated newborn infants.

In a first step, the experiences with VIE-VENT and medical textbook knowledge were incorporated in an "open-patien" model [Gutmayer, 1995]. The main purposes of this approach were to structure the vague causal and functional domain knowledge and to develop a training system for inexperienced physicians. The primary drawback was that the available knowledge was still too incomplete to apply qualitative reasoning techniques for therapy planning. In this paper it will be shown how this model can be improved using on-line measurements. In the long term, the final objectives will be a system, which will regularly obtain data from a patient, customize its disease model to the individual patient, and deliver appropriate advice about management of the patient's health condition.

3 Model

In this section the model for the ventilation will be given and rewritten to fit into the identification scheme of the next section. The model utilizes the experiences and the knowledge gathered from the development of VIE-VENT [Miksch *et al.*, 1996] and of the "open-patient" model [Gutmayer, 1995]. This a priori knowledge was acquired in cooperation with expert physicians. We discussed different medical histories and visualized particular sample cases to structure the necessary elements. The a priori knowledge describes the relationship between the blood-gas measurements, e.g. PO_2 (partial pressure of oxygen), PCO_2 (partial pressure of carbon dioxide), and the ventilator settings, e.g. FiO_2 (fraction of inspired oxygen), $PEEP$ (positive end expiratory pressure). The blood-gas measurements can be sampled at different sites, namely arterial, capillary, venous, and trans-cutaneous. Studies showed a piecewise linear relationship between the different kinds of blood-gas measurements depending on the sampling site [Marsden *et al.*, 1995; Horn *et al.*, 1996]. The model is formulated using the arterial measurements, e.g. P_aO_2 (arterial partial pressure of oxygen). In the evaluation we are using the on-line, continuously-assessed measurements, e.g., $P_{tc}O_2$ (trans-cutaneous partial pressure of oxygen). Beceause of the piecewise linear relationship it is possible to preprocess the data to take this into account.

It is fairly straightforward to see from [Gutmayer, 1995] that the partial pressure of oxygen, P_aO_2 obeys the difference equation

$$\begin{aligned} P_aO_2(k+1) &= W_{Q_s} W_{D_{const}} W_K \left[P_aO_2(k) \right. \\ &+ \left. W_{FiO_2}(k) + W_{PEEP}(k) + W_{VT}(k) \right] \end{aligned}$$

where

$$\begin{aligned} W_{FiO_2}(k) &= 500\,[FiO_2(k) - FiO_2(k-1)] \\ W_{PEEP}(k) &= g_P(PEEP(k)) \\ W_{VT}(k) &= g_V(VT(k)) multVT(f(k), T_i(k)) \\ VT(k) &= C\,[PIP(k) - PEEP(k)] \end{aligned}$$

and where g_P and g_V are known piecewise linear functions, $multVT$ is a known piecewise bilinear function, and where $VT(k)$ is the tidal volume, $f(k)$ is the ventilatory frequency, $T_i(k)$ is the inspiratory time, and $PIP(k)$ is the peak inspiratory pressure. The constants W_{Q_s}, $W_{D_{const}}$, W_K, and C are, based on physicians knowledge, known to be in the intervals

$$\begin{aligned} &W_{Q_s} \in [0.1, 0.97]; \quad W_{D_{const}} \in [0.25, 0.9] \\ &W_K \in [0, 1]; \quad C \in [0.05, 1.5] \end{aligned}$$

Specifically, it should be noted that for x in a small interval

$$g_V(x) = \alpha x + \beta$$

where

$$\alpha \in [0.8, 4]; \quad \beta \in [-5, 12]$$

Defining

$$a = W_{Q_s} W_{D_{const}} W_K$$

$$\begin{aligned} b_1 &= b_2 = a \\ b_3 &= a\alpha C \\ b_4 &= a\beta \end{aligned}$$

and

$$\begin{aligned} y(k) &= P_a O_2(k) \\ u_1(k) &= W_{FiO_2}(k) \\ u_2(k) &= W_{PEEP}(k) \\ u_3(k) &= multVT(f(k), T_i(k))\,[PIP(k) - PEEP(k)] \\ u_4(k) &= multVT(f(k), T_i(k)) \end{aligned}$$

it follows that

$$y(k+1) = ay(k) + b_1u_1(k) + b_2u_2(k) + b_3u_3(k) + b_4u_4(k)$$

with the following constraints on the parameters

$$\begin{aligned} &a = b_1 = b_2 \in [0, 0.873]; \quad b_3 \in [0, 5.238] \\ &b_4 \in [-4.365, 10.476] \end{aligned}$$

For measured data the model will not fit exactly. To this end the model is augmented with the so called regression error $\varepsilon(k)$ to

$$\begin{aligned} y(k+1) &= ay(k) + b_1u_1(k) + b_2u_2(k) \\ &+ b_3u_3(k) + b_4u_4(k) + \varepsilon(k) \end{aligned}$$

This equation is now in a form such that it is possible to apply the identification scheme described in the next section. Notice that the knowledge of g_V is not included in this model. However, all other knowledge is. This is what has to be sacrificed in order to rewrite the model in the above form.

4 Identification Scheme

In this section the identification scheme will be described. It will be seen how it can be cast as a so called Linear Matrix Inequality (LMI) problem, see [Boyd *et al.*, 1994].

Before defining the optimization problem that will yield the identified parameters, the model will be rewritten as a linear regression. With

$$\begin{aligned} \varphi(k) &= \begin{pmatrix} y(k-1) + u_1(k-1) + u_2(k-1) \\ u_2(k-1) \\ u_3(k-1) \end{pmatrix} \\ \theta &= \begin{pmatrix} a \\ b_3 \\ b_4 \end{pmatrix} \end{aligned}$$

it holds that the model reads

$$y(k) = \varphi^T(k)\theta + \varepsilon(k)$$

with constraints $\theta_l. \leq \theta. \leq \theta_u$, where

$$\theta_l = \begin{pmatrix} 0 & 0 & -4.365 \end{pmatrix}^T; \quad \theta_u = \begin{pmatrix} 0.873 & 5.238 & 10.476 \end{pmatrix}^T$$

It is natural to try to fit the parameters of the model such that they obey the above constraints and such that they minimize the quadratic loss function

$$V(\theta, N) = \sum_{i=1}^{N} \lambda^{N-i}\varepsilon^2(i)$$

where $\lambda \in (0, 1]$ is the so called forgetting factor. Typically λ is taken greater than 0.9. For these values, with $M = 1/(1-\lambda)$, it holds that $\lambda^M \approx 0.35$. Hence the rule of thumb is that only the last $1/(1-\lambda)$ samples influence the loss function. Now define

$$Y = \begin{pmatrix} y(1) \\ \vdots \\ y(N) \end{pmatrix}; \quad \Phi = \begin{pmatrix} 0 \\ \varphi(1) \\ \vdots \\ \varphi(N-1) \end{pmatrix}; \quad E = \begin{pmatrix} \varepsilon(1) \\ \vdots \\ \varepsilon(N) \end{pmatrix}$$

Then $E(\theta) = Y - \Phi\theta$, and with $\Lambda = \mathrm{diag}_{k=N-1}^{0}(\lambda^k)$ it holds that $V(\theta, N) = E^T(\theta)\Lambda E(\theta)$, and hence the optimization problem can be written

$$\min_{\theta_l. \leq \theta. \leq \theta_u} E^T(\theta)\Lambda E(\theta)$$

It can be shown, see [Boyd *et al.*, 1994], that this problem is equivalent to

$$\begin{aligned} &\min t \\ \text{s.t.} \quad &\begin{pmatrix} I & \sqrt{\Lambda}(Y - \Phi\theta) \\ (Y - \Phi\theta)^T\sqrt{\Lambda} & t \end{pmatrix} \geq 0 \\ &\theta_l. \leq \theta. \leq \theta_u \end{aligned}$$

The first inequality constraint says that the matrix should be positive semi-definite. Notice that the matrix is symmetric and linear in the variables θ and t. Hence this constraint is an LMI. The second constraint is also linear in the variable θ, and it can also be rewritten as an LMI. This shows that the optimization problem is a so called LMI-problem that can be solved with highly efficient algorithms, see e.g., [Boyd *et al.*, 1994].

In case there are no constraints on θ it can be shown that there is an explicit solution to minimizing $V(\theta, N)$ in terms of a set of linear equations:

$$\Phi^T\Lambda\Phi\theta = \Phi^T\Lambda Y$$

Notice that there is a unique solution if and only if $\Phi^T\Lambda\Phi$ has full rank. Typically this condition is easier to fulfill

the more the signals u_i are varying [Johansson, 1993]. Hence good identification without any a priori knowledge about the parameters is only possible if the input signals are varying enough. The other extreme is that one has very good knowledge about the parameters before the identification starts, and then less variations in the input signals are needed in order to get a good model. This a prior knowledge is of crucial importance to apply on-line identification of patient-disease models, because highly varying ventilator settings might be harmful to the newborn infants.

5 Evaluation

In this section the identification scheme developed in the previous section will be evaluated on real data. The continuously-assessed data available did not contain measurements of the partial pressure of oxygen but the trans-cutaneous ditto. However, the relationship is piecewise linear [Marsden *et al.*, 1995; Horn *et al.*, 1996] and it was possible to pre-process the data to compensate for this. The sample interval used was $h = 10$ min.

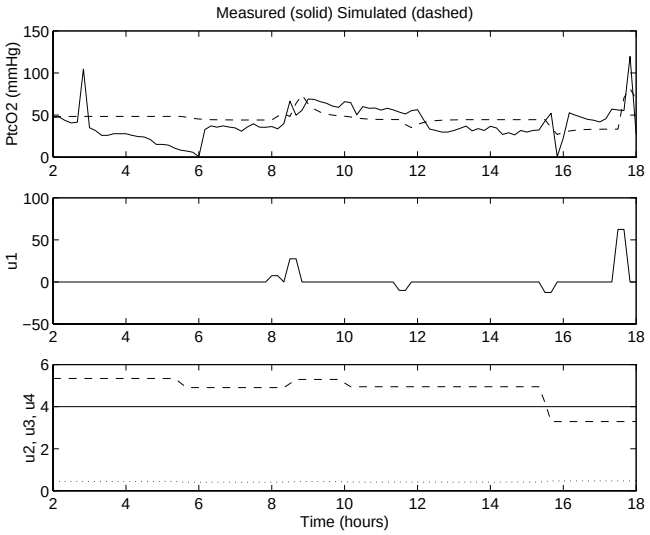

Figure 1: Measured signals and simulated signal.

In Figure 1 is seen how the measured data together with simulation results for the identified model look like. The parameters in the identification scheme were $\lambda = 1$ and $N = 1092$. This means that all N samples made an impact in the loss function. The optimal estimated parameters were $a = b_1 = b_2 = 0.5956$, $b_3 = 3.1044$, and $b_4 = 10.4760$. These are the parameter values used in the simulation, which was performed by iterating the model equation with $y(0) = 0$ and $\varepsilon(k) = 0$ for $k \geq 0$. It should be noted that the input signals u_i do not vary very much, and hence it is not to be expected that the model obtained should be very good. However, this is not the only contribution to the difference in the simulated signal and the measured signal. The peak at about 3 hours is most likely related to a measurement error, which the model cannot capture. Actually anything that cannot be related to changes in the ventilator settings cannot be expected to turn up in simulations based on the model, such as e.g., the negative trend between 3 and 6 hours. However, it is seen that the model does capture some of the behavior in the measured signal, and especially what is related to changes in the ventilator settings.

Since the residuals are very big it would be tempting to try to adapt the model as time goes on, i.e. to do on-line identification and to forget old data using a value of λ less than one. With on-line identification is meant that the model is re-identified whenever new measurement data is available, i.e. $V(\theta, N)$ is minimized at every sample instant. In this way a sequence of optimal estimated parameters $\theta(N)$ is obtained. If old data is forgotten using $\lambda < 1$ the model is said to be adapted. The result is presented in Figure 2 for $\lambda = 0.9$, which implies that only the last $1/(1-0.9) = 10$ samples are remembered in the loss function. Here prediction results are shown, i.e.

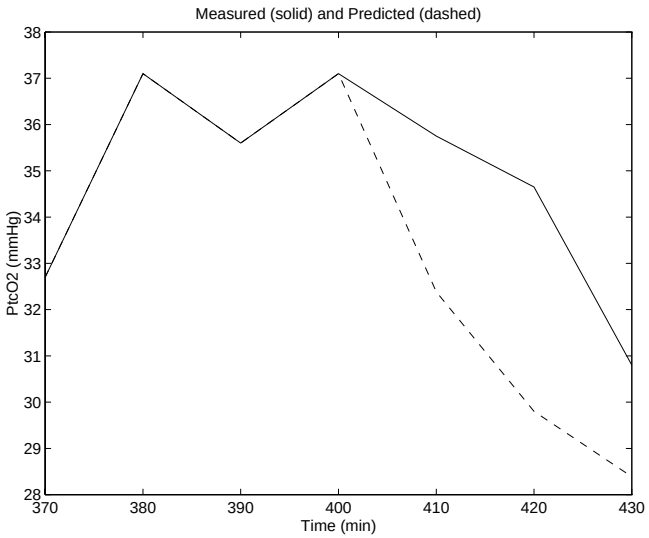

Figure 2: Measured signal and predicted signal with exponentially weighted forgetting ($\lambda = 0.9$, 10 samples are remembered)

the model equation is iterated using $\theta(N)$ as parameter starting at $k = N = 400/h = 40$ with $y(N)$ equal to the measured value at that time and $\varepsilon(k) = 0$ for $k \geq N$. It is seen that the prediction captures the change in the trend of the measured signal, and predicts it fairly correctly for 30 minutes.

6 Benefits, Drawbacks, and Future Developments

Applying our identification scheme to real medical data sets showed particular benefits and problems. As illustrated in Figure 1, our identification scheme fits very satisfactorily to the real developments of the signals, except for measurement errors and effects which are not caused by changes of input signals. The exponential forgetting of old signals makes it possible to adapt the model to changes in the environment. The problems mainly lie in the vague a priori knowledge and the bad quality of the data sets. The identified model could be improved by

- *High quality data sets:* The data sets should include all signals of the model and with accurate temporal correlation. For example the signal VT was not represented in the available data set, and we really did not know exactly at what time any of the ventilator settings ware changed, since the data logging was done partly manually. Also the more the ventilator settings are varying, the more information they contain, and the better the model can be made.
- *Good a priori knowledge:* More off-line experiments (using continuously-assessed data for analysis in retrospect) should be designed to get better general causal and functional dependencies about the patient-disease model. Ideally the only thing that should be estimated on-line is what makes different patients different.

Our on-line identification scheme shows the first step to perform effective monitoring and therapy planning. In the next step, we will customize the patient-disease model to an individual patient and prove how to apply control techniques to perform therapy planning.

Acknowledgment

We are very grateful for the discussions and the contributions of Werner Horn, Walter Gutmayer, Franz Paky, and Christian Popow as well as for the allowance to use their medical data sets to proceed with our experiments.

References

[Avent and Charlton, 1990] R.K. Avent and J.D. Charlton. A critical review of trend-detection methologies for biomedical monitoring systems. *Critical Reviews in Biomedical Engineering*, 16(6):621–59, 1990.

[Bellazzi *et al.*, 1996] R. Bellazzi, C. Larizza, A. Riva, A. Mira, S. Fiocchi, and M. Stefanelli. Distributed intelligent data analysis in diabetic patient management. In *Proceedings of the 1996 AMIA Annual Fall Symposium (formerly SCAMC)*, pages 194–8, Philadelphia, 1996. Hanley-Belfus Inc., Medical Publishers.

[Bellman, 1983] R.E. Bellman. *Mathematical Methods in Medicine*. World Scientific Publ.Co., Singapore, 1983.

[Boyarsky, 1987] A. Boyarsky. Computerized ventilation management system for neonates. *Journal of Perinatolog*, 7, 1987.

[Boyd *et al.*, 1994] S. Boyd, L. El Ghaoui, E. Feron, and V. Balakrishnan. *Linear Matrix Inequalities in System and Control Theory*. SIAM, Philadelphia, 1994.

[Fagan *et al.*, 1980] L.M. Fagan, E.H. Shortliffe, and B.G. Buchanan. Computer-based medical decision making: from mycin to vm. *Automedica*, 3, 1980.

[Fayyad *et al.*, 1996] U. Fayyad, G. Piatetsky-Shapiro, and P. Smyth. The kdd process for extracting useful knowledge from volumes of data. *Communication of the ACM*, 39(11):27–34, 1996.

[Goldsmith and Karotkin, 1993] J.P. Goldsmith and E.H. Karotkin. *Assisted Ventilation of the Neonates*. Saunders, Philadelphia, third edition, 1993.

[Gutmayer, 1995] W. Gutmayer. Entwicklung eines offenen patientenmodells am beispiel der künstlichen beatmung. Master thesis, Medizinische Kybernetik und Artificial Intellegence der Univeristät Wien, 1995.

[Haimowitz *et al.*, 1995] I.J. Haimowitz, P.P. Le, and I.S. Kohane. Clinical monitoring using regression-based trend templates. *Artificial Intelligence in Medicine*, 7(6):473–96, 1995.

[Hamscher *et al.*, 1992] W. Hamscher, L. Console, and J.de Kleer, editors. *Readings in Model-Based Diagnosis*. Morgan Kaufmann, San Mateo, CA, 1992.

[Hayes-Roth *et al.*, 1992] B. Hayes-Roth, R. Washington, D. Ash, R. Hewett, A. Collinot, A. Vina, and A. Seiver. Guardian: A prototype intelligent agent for intensive-care monitoring. *Artificial Intelligence in Medicine*, 4(2):165–85, 1992.

[Horn *et al.*, 1996] W. Horn, S. Miksch, G. Egghart, C. Popow, and F. Paky. Effective data validation of high-frequency data: Time-point-, time-interval-, and trend-based methods. *Computer in Biology and Medicine, Special Issue: Time-Oriented Systems in Medicine*, page to be published, 1996.

[Johansson, 1993] R. Johansson. *System Modeling & Identification*. Prentice Hall, Englewood Cliffs, NJ, 1993.

[Kay, 1993] S.M. Kay. *Fundamentals of Statistical Signal Processing*. Prentice Hall, Englewood Cliffs, NJ, 1993.

[Larizza *et al.*, 1992] C. Larizza, A. Moglia, and M. Stefanelli. *M-HTP*: A system for monitoring heart transplant patients. *Artificial Intelligence in Medicine*, 4(2):111–26, 1992.

[Marsden *et al.*, 1995] D. Marsden, M.C. Chiu, F. Paky, and P. Helms. Transcutaneous oxygen and carbon dioxide monitoring in intensive. *Arch. Dis. Childhood*, 60:1158–61, 1995.

[Miksch *et al.*, 1996] S. Miksch, W. Horn, C. Popow, and F. Paky. Utilizing temporal data abstraction for data validation and therapy planning for artificially ventilated newborn infants. *Artificial Intelligence in Medicine*, 8(6):543–76, 1996.

[Rutledge, 1995] G. W. Rutledge. *Dynamic Selection of Models*. PhD thesis, Stanford University, 1995.

[Shahar and Musen, 1996] Y. Shahar and M.A. Musen. Knowledge-based temporal abstraction in clinical domains. *Artificial Intelligence in Medicine*, 8(3):267–98, 1996.

[Shortliffe, 1991] E.H. Shortliffe. Knowledge-based system in medicine. In *Proceedings Medical Informatics Europe 1991*, Berlin, 1991. Springer.

[Sittig *et al.*, 1990] D.F. Sittig, N.L. Pace, R.M. Gardner, A.H. Morris, and J Wallace. Clinical evaluation of computer-based respiratory care algorithms. *Int. Journal of Clinical Monitoring and Computing*, 7:177–85, 1990.

[Uckun *et al.*, 1993] S. Uckun, B.M. Dawant, and D.P. Lindstrom. Model-based diagnosis in intensive care monitoring: the yaq approach. *Artificial Intelligence in Medicine*, 5(1):31–48, 1993.

[Uckun, 1993] S. Uckun. Intelligent systems in patient monitoring and therapy management. Technical Report Report KSL 93-32, Stanford University, Knowledge Systems Laboratory, 1993.

[Wiener, 1948] N. Wiener. *Cybernetics or Control and Communication in the Animal and the Machine*. John Wiley & Sons, Inc., New York, 1948.