Robust and Adaptive Methods in Analysing Online Monitoring Data

Ursula Gather and Karen Schettlinger
University of Dortmund, Department of Statistics
44221 Dortmund, Germany
E-mail: gather@statistik.uni-dortmund.de

1. Introduction

Current alarm systems used for critical care monitoring are clinically unsatisfactory. Variables such as heart rate or blood pressure are recorded at high frequency and compared to thresholds set by the medical staff for critical situations. These automatic alarm systems produce many false alarms due to measurement artifacts, patient movements, or transient fluctuations around the chosen alarm limit.

Physiologic time series monitored by current alarm systems exhibit trends, abrupt level changes and large spikes (outliers) as well as periods of relative stability. The measurements are overlaid with a high level of noise and many artifacts. Also, strong dynamic dependence can occur between the variables. It is important to construct alarm rules based on the true underlying signal, i.e., on a noise-free and artifact-free time series.

We present filtering procedures for robust online signal extraction and discuss their merits for preserving clinically relevant patterns such as trends, abrupt shifts and extremes and for the removal of irrelevant spikes or outliers. Our approach uses high-breakdown linear regression methods in moving time windows. Increased computational power and fast algorithms allow for applying these robust regression filters even in real time.

In the multivariate case, simply generalizing univariate robust regression methods does not result in affine equivariant procedures: if the error terms of the variables are highly correlated, efficiency is lost for multivariate signal extraction. On the other hand, multivariate affine equivariant regression methods with high breakdown usually assume that the data are in general position. However, for intensive care data this is often not the case since the measurements are recorded on a discrete scale.

We present procedures for univariate and multivariate signal extraction which are fast, efficient and robust, and can be used for discretely measured data with low variability as well as in situations with many outliers.

2. Univariate Signal Extraction

In the following, we consider a real valued time series $(\mathbf{y}_t)_{t\in\mathbb{Z}}$ observed at time points $t = 1, \ldots, N$. For this time series a simple decomposition model can be assumed, namely

$$x_t = \mu_t + \epsilon_t + \nu_t \,, \qquad t \in \mathbb{N}$$

where μ_t represents the underlying signal, ϵ_t symmetric, obervational noise and ν_t an outlier generating process, causing 'spiky' noise.

A simple approach for extracting the signal is to use a moving average: In time windows of fixed length n the average of the observations is calculated for estimation of the signal in the window centre. Moving averages are popular since they trace trends and are very efficient for Gaussian samples. However, sudden level shifts are 'smeared' and outliers can cause a considerable bias.

A running median, as suggested by Tukey (1977), is robust against outliers and capable of tracing level shifts. This filter can resist up to $[n/2]$ outliers within one time window but it deteriorates in trend periods.

Since intensive care time series often show trends and shifts but also contain outliers, methods applied to ICU data need to able to do both, trace trends and resist outlying values.

In view of the weakness of the running medians in trend periods, a better adaptation to temporal trends is achieved by assuming the signal to be locally linear instead of locally constant (Davies, Fried, Gather, 2004). This means, within a time window centred at time point t, the following model is assumed:

$$x_{t+i} = \mu_t + \beta_t\, i + \varepsilon_{t,i}, \quad i = -m, \ldots, m\,,$$

where μ_t again denotes the underlying level of the signal and β_t the slope at time t; the $\varepsilon_{t,i}$ are independent error terms with median zero.

The current level can either be represented by the level μ_t in the centre of the time window or by the last fitted value $\hat{\mu}_{t+m}^{online} = \hat{\mu}_t + \hat{\beta}_t\, m$ within one window (Gather, Schettlinger, Fried, 2006).

The choice of the window width n (or m respectively) can have a large impact on the extracted signal: Larger window widths assure robustness against a larger number of outlying values and result in smaller variability of the estimations, which yields a smoother extracted signal. However, they also come with a larger bias and larger computation times. In the following, the considered window widths of $n = 21$ or $n = 31$ seem to offer a compromise for the bias variance trade-off, and they are considered acceptable for the physiological data we have to analyse.

Further, standard methods for the estimation of μ_t and $\hat{\beta}_t$ such as least squares regression are not suitable in the presence of outliers.

It is rather advisable to apply robust regression methods which are able to deal with a certain amount of contamination without becoming strongly affected. Denoting the residuals in a window by $r_{t+i} = x_{t+i} - (\hat{\mu}_t + \hat{\beta}_t\, i)$ (for $i = -m, \ldots, m$), we compare the regression methods listed below for application to intensive care time series.

Apart from the well-known **$\mathbf{L_1}$−Regression** (Edgeworth, 1887), **Least Median of Squares** regression (LMS), introduced by Hampel (1975) and investigated in more detail by Rousseeuw (1983), was one of the first highly robust regression methods and hence shall be of special interest in the following. The LMS is defined by

$$(\tilde{\mu}_t^{LMS}, \tilde{\beta}_t^{LMS})' \quad = \quad \arg\min_{\tilde{\mu}_t, \tilde{\beta}_t} \left\{ \operatorname*{med}_{i=-m}^{m} \; \{r_{t+i}^2\} \right\}.$$

Replacing the sum in the least squares approach by a median is not only an intuitively nice way of achieving robustness but also results in an estimator exhibiting several merits: The LMS is regression-, scale- and affine equivariant, and if the data are in general position it has a breakdown point of $\lfloor n/2 \rfloor$/n, which is the maximum breakdown point among all regression equivariant estimators. Further, it possesses the exact fit property and has minimax bias in the class of 'residual admissible' estimates (Yohai, Zamar, 1993), a class which contains M-, S-, τ-, the LMS, the LTS and some R-estimates.

However, the Gaussian efficiency of the LMS amounts only to $20 - 25\%$ for small sample sizes and it gets even worse for increasing n. Also, the LMS has a rather low convergence rate of $n^{-1/3}$ and peaks or discontinuities around the quartiles of the influence function and sensitivity curve (Davies, 1993; Sheather, McKean, Hettmansperger, 1997). This causes unbounded influence and sensitivity w.r.t the effects of centrally located x-values and could also account for instabilities of the LMS w.r.t. small changes of centrally located data (e.g. Hettmansperger, Sheather, 1992).

Further, its exact solution is sometimes hard to compute as the objective function can exhibit several local minima; a straightforward implementation would thus need $O(n^4)$ computation time (Stromberg, 1993). Using the results from Edelsbrunner and Souvaine (1990), Bernholt (2005) was able to improve the LMS algorithm to $O(n^2)$ taking up $O(n^2)$ storage space.

Repeated Median (RM) regression (Siegel, 1982) is another highly robust regression technique with regression estimates defined by

$$\tilde{\beta}_t^{RM} = \underset{i=-m}{\overset{m}{\text{med}}} \left\{ \underset{j \neq i}{\text{med}} \frac{y_i - y_j}{i - j} \right\}$$

$$\tilde{\mu}_t^{RM} = \underset{i=-m}{\overset{m}{\text{med}}} \left\{ x_{t+i} - \tilde{\beta}_t^{RM} i \right\}.$$

Like the LMS it has a breakdown point of $\lfloor n/2 \rfloor / n$, possesses the exact fit property, and it is regression-and scale-equivariant. Further it is equivariant w.r.t. affine transformations of the response but lacks equivariance w.r.t. affine transformations of the explanatory variable.

An exact straightforward algorithm would need $O(n^2)$ computation time (Siegel, 1982), but applying the RM to moving time windows an update of the estimates can be achieved in linear time (Bernholt, Fried, 2003). Here, the term 'update' means that estimation takes place by using the stored information from the last time window - only inserting the new information given by the most current data point and deleting that of the oldest data point. Thus, update algorithms save computation time as the estimates do not have to be calculated for each window from scratch.

Another regression technique included in the comparisons below is **Least Trimmed Squares** (LTS) regression (Rousseeuw, 1983), defined by

$$(\tilde{\mu}_t^{LTS}, \tilde{\beta}_t^{LTS})' = \arg\min_{\tilde{\mu}_t, \tilde{\beta}_t} \sum_{k=1}^{h} (r_t^2)_{k:n},$$

where $(r_t^2)_{k:n}$ denotes the kth ordered squared residual for the current time window.

Letting $h = \lfloor n/2 \rfloor + 1$, the LTS yields the same finite sample breakdown point as the LMS, i.e. $\lfloor n/2 \rfloor / n$. The LTS is also regression-, scale- and affine equivariant and possesses the exact fit property.

A straightforward implementation of the LTS algorithm would take $O(n^3 \log n)$ computation time and $O(n^2)$ memory space (Hössjer, 1995). However, for the comparisons, an algorithm requiring $O(n^2)$ time and $O(n^2)$ storage space is used.

The last method to be included in the comparisons is **Deepest Regression** (DR) (Rousseeuw and Hubert, 1999), defined by

$$(\tilde{\mu}_t^{DR}, \tilde{\beta}_t^{DR}) = \arg\max_{\tilde{\mu}_t, \tilde{\beta}_t} \left\{ rdepth\Big((\tilde{\mu}_t, \tilde{\beta}_t), \mathbf{x}_t \Big) \right\}$$

with the regression depth *rdepth* of a fit $(\tilde{\mu}_t, \tilde{\beta}_t)$ at a sample $\mathbf{x}_t = (x_{t-m}, \ldots, x_{t+m})'$ defined as

$$\Big((\tilde{\mu}_t, \tilde{\beta}_t), \mathbf{x}_t \Big) = \min_{-m \leq i \leq m} \left\{ \min\{ L_t^+(i) + R_t^-(i), R_t^+(i) + L_t^-(i) \} \right\}$$

$$\text{with } L_t^+(i) = L_{t,\tilde{\mu}_t,\tilde{\beta}_t}^+(i) = \#\Big\{ j \in \{-m, \ldots, i\} : \; r_{t+j}(\tilde{\mu}_t, \tilde{\beta}_t) \geq 0 \Big\}$$

$$\text{and } R_t^-(i) = R_{t,\tilde{\mu}_t,\tilde{\beta}_t}^-(i) = \#\Big\{ j \in \{i+1, \ldots, m\} : \; r_{t+j}(\tilde{\mu}_t, \tilde{\beta}_t) < 0 \Big\},$$

with $L_t^-(i)$ and $R_t^+(i)$ defined similarly.

In contrast to LMS, LTS and RM regression, this method only has a breakdown point of about 1/3, but it is also regression-, scale- and affine equivariant. Straightforward computation would need $O(n^3)$ time, but we make use of an update algorithm which only needs $O(\log^2 n)$ time and $O(n)$ storage.

3. Comparisons for Univariate Signal Extraction - Simulations

All of the methods described above are robust w.r.t. outlying values. Tab. 1 lists the values for $k^\star$, the smallest number of contaminated observations within a window of size $n = 21$ or $n = 31$ which can cause a spike of any size in the extracted signal,

$$k^\star = \min\{k : \sup\{\parallel T(\boldsymbol{z}) - T(\boldsymbol{x}) \parallel, \boldsymbol{z} \in U_k(\boldsymbol{x})\} = \infty\}$$

with $U_k(\boldsymbol{x}) = \{\boldsymbol{z} = (z_1, \ldots, z_n) : \sharp\{i : z_i \neq x_{t+i}\} = k\}$.

Clearly, the methods with maximum breakdown point, namely LMS, LTS and RM are preferable to *Least Squares* regression where one outlier can cause already arbitrary large bias.

$k^\star$	$\mathbf{L}_2$	$\mathbf{L}_1$	**LMS**	**LTS**	**RM**	**DR**
$n = 21$	1	7	10	10	10	≥ 6
$n = 31$	1	10	15	15	15	≥ 10

Tab. 1: Smallest number of contaminated observations which can cause a spike of any size in the extracted signal.

Comparisons of the methods by their finite sample efficiencies at standard normal data with some trend ($\beta_t = 0.1$) and no trend ($\beta_t = 0$) show that the LMS has with about $20 - 25\%$ the lowest efficiency of all methods. Further it can be seen that the regression based methods are much more efficient than a simple median filter in trend periods.

Tab. 2 displays the finite sample efficiencies of the different online estimates $\hat{\mu}_t^{online}$ relative to the least squares estimate at situations with no trend ($\beta_t = 0$) and standard normal errors, heavy-tailed errors from a rescaled t_3-distribution with unit variance and shifted lognormal errors with median zero and unit variance.

		LMS	**LTS**	**RM**	**DR**
standard normal	$n = 21$	23.0	22.4	70.5	61.8
errors	$n = 31$	21.1	20.1	69.5	61.0
rescaled t_3	$n = 21$	57.7	56.2	143.7	136.5
errors	$n = 31$	57.9	56.6	150.3	142.4
shifted lognormal	$n = 21$	44.2	43.1	99.7	92.8
errors	$n = 31$	35.1	33.4	85.0	78.4

Tab. 2: Finite sample efficiencies for the online estimates $\hat{\mu}_t^{online}$ relative to least squares.

We see quite clearly in Tab. 2, that in terms of efficiency, RM performs best while DR is not much worse. Particularly in the non-normal situations, these robust methods can even outperform least squares while LMS and LTS are still much worse.

Further simulations compare the performance of the different methods at situations with standard errors and (a) an increasing number of outliers at random time points within a time window of size $n = 21$ and (b) a patch of outliers, i.e., an increasing number of outliers of the same size at subsequent time points at the end of a time window ($n = 21$), indicating a so called 'level shift'. For online monitoring, it is of special importance to trace sudden shifts of the signal because it may point at an abrupt change of the patient's state of health.

The performance of the estimates is measured by their estimated standard deviation, the bias w.r.t. the true signal level (i.e. zero) and the root mean squared error.

For the situations with outliers at random time points within the window, one would wish all curves to stay as close to zero as possible (small bias and small variability). Fig. 1(a) shows first that the

difference in the outcomes of LMS and LTS regression is negligible, and also that the results for RM and DR are quite similar.
LMS and LTS generally have the smallest bias and also the smallest standard deviation and root MSE for a large number of medium-sized or large outliers. On the other hand, RM and also DR have the smallest standard deviation and root MSE when the outliers are small in number or magnitude.

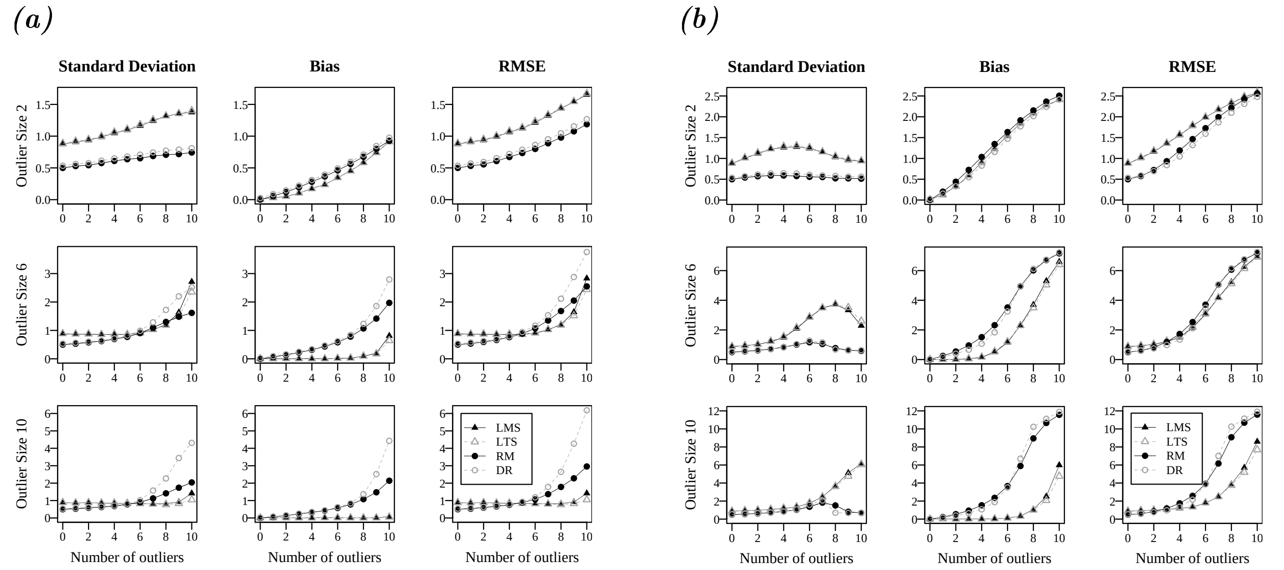

Fig. 1: Standard deviation, bias and root MSE of the online estimates $\hat{\mu}_t^{online}$ at standard normal data with outliers (a) at random time points and (b) at subsequent time points at the end of a window of size $n = 21$.

For the situations with consecutive outliers at the end of the window, it is desirable to find a method which is as little influenced by the outlying pattern as possible, if the number of outliers is small. However, since a larger patch might indicate a relevant change in the data structure, a larger bias does not indicate worse performance of the method but rather the ability to trace sudden shifts. A method which searches a compromise between the 'good' or the 'old' values respectively and the outlying pattern will not be able to track the sudden shift but will 'smear' the transition from one state to the other.
Fig. 1 (b) shows again the similarities between LMS and LTS as well as between RM and DR. Here, RM and DR show the best performance in terms of variability for any number of outliers, while LMS and LTS have a smaller bias. For a larger number of medium-sized or large outliers LMS and LTS have the smallest root MSE. In all other situations, RM and DR outperform them.

(a)

	LMS	**LTS**	**RM**	**DR**
$n = 21$	0.161	0.161	0.035	0.747
$n = 31$	0.323	0.324	0.049	0.956

(b)

	LMS	**LTS**	**RM**	**DR**
time	$O(n^2)$	$O(n^2)$	$O(n)$	$O(n \log^2 n)$
memory space	$O(n^2)$	$O(n^2)$	$O(n^2)$	$O(n)$

Tab. 3: (a) Simulated mean computation time of an update in milliseconds and (b) asymptotic computation time and memory space.

In terms of computation times, the simulated as well as the asymptotic values shown in Tab. 3 (a) and (b) show the advantage of the RM update algorithm by Bernholt and Fried (2003).

4. Comparisons for Univariate Signal Extraction - Applications

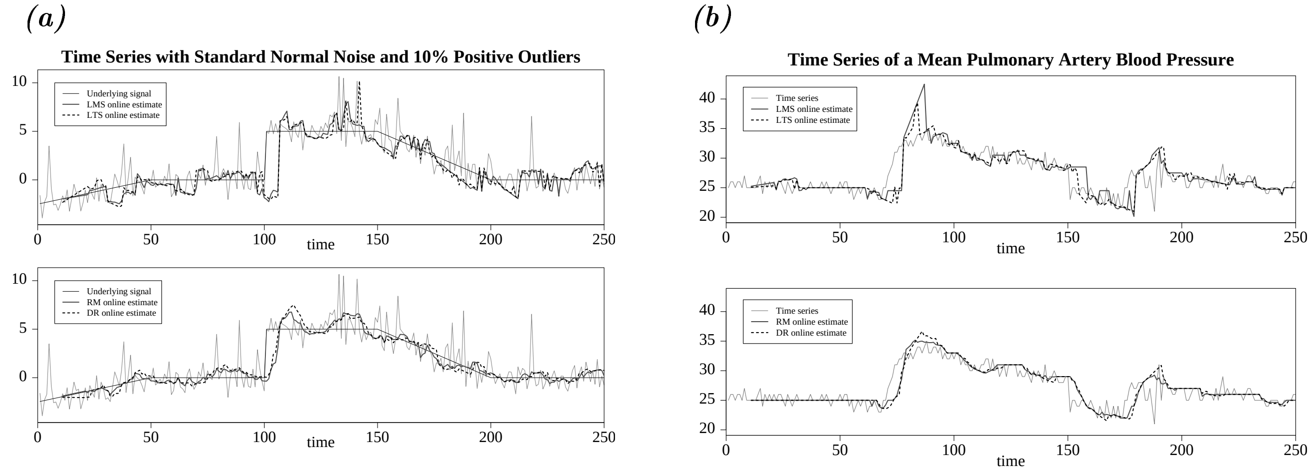

Fig. 2: Comparison of the regression filters at (a) a simulated and and (b) a real physiological time series.

Fig. 2 (a) shows a simulated time series with a signal which contains constant periods, trends, trend changes as well as a level shift. The signal is overlaid with standard normal noise and 10% positive additive outliers of size 5 at random time points.

Fig. 2 (b) depicts an extract of 250 measurements of a mean pulmonary artery blood pressure from an intensive care patient observed at a frequency of once per minute (in mmHg).

Both panels in Fig. 2 show that the differences in the outcomes between LMS and LTS regression are negligible, and also that there is little difference between the RM and DR filters. All filters estimate the signal well in constant and trace trend changes and shifts.

The LMS and LTS filters capture the suddenness of a shift but over-estimate the signal right after the shift. The RM and DR filters have a shorter time delay for the shift detection but therefore 'smear' shifts.

While the LMS and LTS estimates also tend to be more influenced by moderate data variation, the RM and DR filters yield smoother online signal estimations.

In conclusion, an LMS (or LTS) filter is considered the best choice for univariate online signal extraction if the series is to contain many large outliers, many shifts and it is feasible to estimate the current signal with some (small) time delay. For fully online signal estimation at trended time series with maximally moderately sized outliers, an RM filter is to be recommended, taking into account its fast computation.

5. Multivariate Signal Extraction

Extracting the signals for each time series observed at one patient separately disregards the dependence between physiological variables. Therefore, the univariate approaches should be generalised to the multivariate case.

Koivunen (1996) proposes a highly robust procedure which is based on a gliding minimum covariance determinant (MCD) location estimator. However, ICU data often show periods with strong trends and therefore regression based methods should be preferred to location based methods, because they produce smaller bias in such periods.

Analogously to the univariate case, a local linear model within a time window of length $n = 2m+1$ is assumed:

$$\boldsymbol{x}(t+i) = \boldsymbol{\mu}(t) + \boldsymbol{\beta}(t)\, i + \boldsymbol{\varepsilon}(t,i) + \boldsymbol{\eta}(t,i)\,, \quad i = -m, \ldots, m,$$

with $\boldsymbol{\mu}(t) = (\mu_1(t), \ldots, \mu_k(t))^T$ denoting the $k-$variate signal, $\boldsymbol{\varepsilon}(t) \in \mathbb{R}^k$ the errors, and $\boldsymbol{\nu}(t) \in \mathbb{R}^k$ an outlier generating process. Thus, for each time window a multivariate regression problem has to be solved.

Simply generalising the univariate regression techniques to the multivariate case generally results in methods which are not affine equivariant. If there are correlations between the error components such estimators lose some efficiency.

Robust multivariate regression techniques such as MCD regression (Rousseeuw, Van Aelst, Van Driessen, Agulló, 2004) or multivariate LTS regression (Agulló, Croux, Van Aelst, 2007) are based on the MCD functional. However, currently existing algorithms for the MCD use heuristics and need much computation time, but for application in intensive care fast and exact algorithms are required for reliable real-time signal extraction. Furthermore, the requirement that the data are in general position is often not fulfilled for intensive care data.

To overcome these difficulties, Lanius (2005) proposes a new multivariate filtering procedure assuming that every single component $\mu_j(t)$, $j = 1, \ldots, k$ of the multivariate time series is locally linear.

Within each time window $\{\boldsymbol{x}(t+i),\ i = -m, \ldots, m\}$ the signal is first approximated by k lines using a univariate robust regression technique to determine $\hat{\mu}_j(t)$ and $\hat{\beta}_j(t)$ for each $j = 1, \ldots, k$. Resting upon the results for univariate signal extraction, either the RM or the LMS seems to be a suitable choice here. For the resulting multivariate residuals $\boldsymbol{r}(t+i) = \boldsymbol{x}(t+i) - \hat{\boldsymbol{\mu}}(t) - i\hat{\boldsymbol{\beta}}(t)$ a modified orthogonalised Gnanadesikan-Kettenring estimator (Maronna, Zamar, 2002) is used for a robust estimation of $\boldsymbol{\Sigma}(t)$, the covariance matrix of the error terms. Using the estimated covariance matrix, the Mahalanobis distance $\boldsymbol{r}(t+i)^T\hat{\boldsymbol{\Sigma}}(t)^{-1}\boldsymbol{r}(t+i)$ of all multivariate residuals is determined, and observations for which the corresponding residuals have a Mahalanobis distance larger than a certain constant d_n are trimmed from the sample. To accomodate for the correlation structure between the variables, the last step applies a multivariate least squares regression to the remaining observations in the local time window for estimation of the current signal level.

In short time windows this procedure is even more efficient than the MCD-based regression techniques. Also, it is feasible for the intensive care data and fast to compute.

6. Conclusions

For univariate signal extraction, the LMS filter offers the highest robustness against many large outliers and is able to track level shifts and trend changes well. The RM filter slightly smoothes such changes. Nevertheless, the repeated median is considered the best choice for signal extraction as it not only offers considerable robustness against outliers but is also stable w.r.t. moderate variations in the data. Additionally, computation of the RM filter is with the $O(n)$ update algorithm from Bernholt and Fried (2003) much faster than the fastest exact $O(n^2)$ LMS algorithm (Bernholt, 2005).

The advantages of the univariate techniques, particularly the high robustness and fast computation, directly transfer to the described multivariate signal extraction procedure which additionally takes into account possible dependence between the variables.

REFERENCES (RÉFERENCES)

Agulló, J., Croux, Ch., Van Aelst, S. (2007). The Multivariate Least Trimmed Squares Estimator. *J. Multivariate Anal.*, to appear.

Bernholt, Th. (2005). Computing the Least Median of Squares Estimator in Time $O(n^d)$. *Proc. ICCSA 2005, LNCS* **3480**, 697-706.

Bernholt, Th., Fried, R. (2003). Computing the Update of the Repeated Median Regression Line in Linear Time. *Inf. Process. Lett.* **88** (1), 111-117.

Davies, P.L. (1993). Aspects of Robust Linear Regression. *Ann. Stat.* **21** (4), 1843-1899.

Davies, P.L., Fried, R., Gather, U. (2004). Robust Signal Extraction for On-line Monitoring Data. *J. Stat. Plann. Inference* **122** (1-2), 65-78.

Edelsbrunner, H., Souvaine, D.L. (1990). Computing Least Median of Squares Regression Lines and Guided Topological Sweep. *JASA* **85**, 115-119.

Edgeworth, F.Y. (1887). *Metretike or the Method of Measuring Probabilities and Utilities*, London, Temple Company.

Gather, U., Schettlinger, K., Fried, R. (2006). Online Signal Extraction by Robust Linear Regression. *Comput. Statistics* **21** (1), 33-52.

Hampel, F.R. (1975). Beyond Location Parameters: Robust Concepts and Methods. *Bulletin of the International Statistical Institute* **46**, 375-382.

Hettmansperger, T.P. and Sheather, S.J. (1992). A Cautionary Note on the Methods of Least Median Squares. *Am. Stat.* **46** (2), 79-83.

Hössjer, O. (1995). Exact Computation of the Least Trimmed Squares Estimate in Simple Linear Regression. *Comput. Stat. Data An.* **22** (3), 1478-1501.

Koivunen, V. (1996). Nonlinear Filterung of Multivariate Images under Robust Error Criterion. *IEEE T. Image Process.* **5**, 1054-1060.

Lanius, V. (2005). *Statistische Extraktion relevanter Information aus multivariaten Online-Monitoring-Daten der Intesivmedizin.* Dissertation, University of Dortmund, Germany.

Maronna, R.A., Zamar, R.H. (2002). Robust Estimates of Location and Dispersion for High-dimensional Datasets. *Technometrics* **44**, 307-317.

Rousseeuw, P.J. (1983). Multivariate Estimation with High Breakdown Point. In W. Grossmann, G. Pflug, I. Vincze, W. Wertz (eds.) *Proceedings of the 4th Pannonian Symposium on Mathematical Statistics and Probability* **B**, D. Reidel Publishing Company, Dordrecht (The Netherlands).

Rousseeuw, P.J., Hubert, M. (1999). Regression Depth. *J. Am. Stat. Assoc.* **94**, 388-402.

Rousseeuw, P.J., Van Aelst, S. Van Driessen, K., Agulló, J. (2004). Robust Multivariate Regression. *Technometrics* **46**, 293-305.

Sheather, S.J., McKean, J.W. and Hettmansperger, T.P. (1997). Finite Sample Stability Properties of the Least Median of Squares Estimator. *J. Stat. Comput. Sim.* **58**, 371-383.

Siegel, A.F. (1982). Robust Regression Using Repeated Medians, *Biometrika* **69**, 242-244.

Stromberg, A.J. (1993). Computing the Exact Least Median of Squares Estimate and Stability Diagnostics in Multiple Linear Regression. *SIAM J. Sci. Comput.* **14** (6), 1289-1299.

Tukey, J.W. (1977). *Exploratory Data Analysis*, Reading, MA, Addison-Wesley.

Yohai, V.J., Zamar, R.H. (1993). A Minimax-Bias Property of the Least α-Quantile Estimates. *Ann. Stat.* **21**(4), 1824-1842.

ABSTRACT (RÉSUMÉ)

We present procedures for online signal extraction from intensive care data. These filtering methods use high-breakdown linear regression methods in moving time windows. In particular, we concentrate on the performance of Least Median of Squares *(LMS) and* Repeated Median *(RM) regression in this context. Comparing those to other robust regression techniques, we discuss their merits for preserving clinically relevant patterns such as trends, abrupt shifts and extremes and for the removal of irrelevant spikes or outliers.*