

A Service of

ZBW

Leibniz-Informationszentrum Wirtschaft Leibniz Information Centre for Economics

Schettlinger, Karen; Fried, Roland; Gather, Ursula

Working Paper Robust Filters for Intensive Care Monitoring: Beyond the Running Median

Technical Report, No. 2006,23

Provided in Cooperation with:

Collaborative Research Center 'Reduction of Complexity in Multivariate Data Structures' (SFB 475), University of Dortmund

Suggested Citation: Schettlinger, Karen; Fried, Roland; Gather, Ursula (2006) : Robust Filters for Intensive Care Monitoring: Beyond the Running Median, Technical Report, No. 2006,23, Universität Dortmund, Sonderforschungsbereich 475 - Komplexitätsreduktion in Multivariaten Datenstrukturen, Dortmund

This Version is available at: https://hdl.handle.net/10419/22666

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



WWW.ECONSTOR.EU

Robust Filters for Intensive Care Monitoring – Beyond the Running Median

Karen Schettlinger

Roland Fried

Ursula Gather

Fachbereich Statistik Universität Dortmund 44221 Dortmund

Abstract

Current alarm systems on intensive care units create a very high rate of false positive alarms because most of them simply compare the physiological measurements to fixed thresholds. An improvement can be expected when the actual measurements are replaced by smoothed estimates of the underlying signal. However, classical filtering procedures are not appropriate for signal extraction as standard assumptions, like stationarity, do no hold here: the measured time series often show long periods without change, but also upward or downward trends, sudden shifts and numerous large measurement artefacts. Alternative approaches are needed to extract the relevant information from the data, i.e. the underlying signal of the monitored variables and the relevant patterns of change, like abrupt shifts and trends. This article reviews recent research on filter based online signal extraction methods which are designed for application in intensive care.

1 Introduction

Monitoring systems in intensive care need to be credible tools for judging the state of the critically ill. Apart from the actual measurements, relevant patterns of change, like abrupt shifts or monotonic trends, contain essential information about a patient's condition. Therefore, methods are required for the reliable extraction of this information from the data. At the same time the methods have to be able to deal with many artefacts and irrelevant minor fluctuations.

Most alarm systems currently used for the haemodynamic monitoring in intensive care are essentially based on thresholds: violations of the upper or lower control limit activate an alarm – sometimes after a certain offset time. For example, the monitoring system we study here triggers an alarm for the systolic arterial blood pressure if the measurements exceed the upper control limit or fall below the lower control limit for at least four seconds. This offset time is one possibility to make the system robust against single measurement artefacts. However, experience with real data sets suggests that in practice such artefacts also occur in 'patches' of several consecutive values. Thus, even this proceeding does not completely avoid false alarms (see Fig. 1).



Figure 1 Current alarm systems trigger false alarms, e.g. because of measurement artefacts.

Pre-processing the input data for an alarm system by robust online filtering can be expected to yield considerably less false alarms as it does not only remove single but also short patches of artefacts. We review some robust versions from the broad variety of filtering methods, exhibiting certain characteristics which are desirable in the online monitoring context.

Measurements $(y_t)_{t \in \mathbb{Z}}$ from a physiological variable, recorded in short time intervals up to once per second, can be represented by a simple 'signal plus noise' model

$$y_t = \mu_t + u_t$$
 for $t \in \mathbb{Z}$.

Here, μ_t symbolises the underlying 'true' biosignal, which is assumed to vary smoothly with a few sudden changes, and u_t defines the noise component. This component can contain large aberrant values, e.g. due to measurement artefacts, which are called 'outliers' in the statistical literature.

A simple approach for extracting the signal μ_i 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 (see **Fig. 2**).

A running median, as suggested by Tukey [1], 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 (see also Fig. 2). In the following, we review filtering techniques for signal extraction which are robust against outliers but additionally capable of tracing trends, trend changes and level shifts. We compare these methods by applications to intensive care data, discuss their performance in situations which are of particular interest in the online monitoring context, and point out further demands for future research.



Figure 2 Moving average and running median applied to a time series of pulse measurements from intensive care.

2 Signal Extraction Methods

2.1 Simple Robust Regression Filters

In view of the weakness of the running medians in trend periods, Davies, Fried and Gather [2] achieve a better adaptation to temporal trends by assuming the signal to be locally linear instead of locally constant. This means, within a time window centred at time point t the following model is applied:

$$y_{t+i} = \mu_t + \beta_t i + \varepsilon_{t,i}$$
 for $i = -m,...,m$,

where μ_t again represents the underlying signal, and β_t is the slope in the window centre, while $\varepsilon_{t,i}$ describes the noise component.

Standard methods for the estimation of μ_t and β_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 srongly affected. Denoting the residuals in a window by

$$r_{t+i} = y_{t+i} - (\widetilde{\mu}_t + \beta_t i)$$
 for $i = -m, \dots, m$,

Davies, Fried and Gather [2] survey the following techniques for estimating μ_t and β_t :

- L₁ Regression

$$(\widetilde{\mu}_{t}^{L_{1}},\widetilde{\beta}_{t}^{L_{1}}) = \arg\min_{\widetilde{\mu}_{t},\widetilde{\beta}_{t}} \left\{ \sum_{i=-m}^{m} |r_{t+i}| \right\}$$

- Least Median of Squares (LMS) Regression [3]

$$(\widetilde{\mu}_{t}^{LMS},\widetilde{\beta}_{t}^{LMS}) = \arg\min_{\widetilde{\mu}_{t},\widetilde{\beta}_{t}} \{ \underset{i=-m,...,m}{med} (r_{t+i}^{2}) \}$$

- Repeated Median (RM) Regression [4]

$$\widetilde{\beta}_{t}^{RM} = \underset{i=-m,...,m}{med} \left\{ \underset{j=-m,...,m; j \neq i}{med} \frac{y_{t+i} - y_{t+j}}{i-j} \right\}$$
$$\widetilde{\mu}_{t}^{RM} = \underset{i=-m,...,m}{med} \{ y_{t+i} - \widetilde{\beta}_{t}^{RM} i \}$$

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 because it does not only offer considerable robustness against outliers, but it is also stable w.r.t. moderate variations in the data. Additionally, computation of the RM filter is much faster: In [5] an algorithm for the RM regression line is presented which only needs linear time for an update. 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 a lot of computation time as the estimates do not have to be calculated for each window from scratch.

Estimation of the parameters μ_t and β_t in the centre of the time window means a delay of *m* time units for the filter output. Taking into account the urgency of reliable output on intensive care units, only minimal delays are acceptable. Thus, signal extraction as described above is rather suitable for retrospective analyses when applying a large window width. For signal extraction without time delay, Gather, Schettlinger and Fried [6] examine the online estimates

$$\widetilde{\mu}_{t+m}^{online} = \widetilde{\mu}_t + \widetilde{\beta}_t m ,$$

estimating the signal value at the most recent time point. Since both RM and LMS regression show certain advantages in [2], these methods are considered again and compared to two further robust regression methods in the online situation: *least trimmed squares* (LTS) *regression* [7] and *deepest regression* (DR) [8].

It turns out that the differences in the outcomes between LMS and LTS regression are negligible, and also that there is little difference between the repeated median and deepest regression filters (see **Fig. 3**). In the online situation, LMS and LTS track shifts with a longer delay than their competitors and tend to overshoot shifts, while RM and DR show more stable results (see also Fig. 3). Considering the computational speed, again the repeated median is recommended for applications in intensive care.



Figure 3 Online signals extracted with four different regression filters.

2.2 Repeated Median Hybrid Filters

As pointed out above, a simple RM filter does not preserve sudden level shifts as such but 'smears' them somewhat [2]. Heinonen and Neuvo [9], [10] emphasise the advantages of linear median hybrid filters for preserving such signal edges. FIR median hybrid (FMH) filters are computationally even less demanding than running medians and preserve shifts similarly good or even better than these. An FMH filter is defined as the median of several linear subfilters:

$$FMH(y_t) = med\{\Phi_1, \Phi_2, \dots, \Phi_M\}.$$

For signal extraction from blood pressure measurements, Heinonen, Kalli, Turjanmaa and Neuvo [11] use a simple FMH filter with M = 3 subfilters, consisting of two one-sided moving averages and the central observation as central subfilter:

$$\Phi_1(y_t) = \frac{1}{m} \sum_{i=1}^m y_{t-i}, \ \Phi_2(y_t) = y_t, \ \Phi_3(y_t) = \frac{1}{m} \sum_{i=1}^m y_{t+i}.$$

Similar to running medians, such simple FMH filters assume the signal to be locally constant. Predictive FMH (PFMH) filters use one-sided weighted averages instead of ordinary half-window averages for tracking linear trends [10]. Combined FMH filters, finally, combine the structures for a local constant and for a local linear signal. However, these filters can only remove single isolated outliers and hence, they are not sufficiently robust for applications in intensive care.

Fried, Bernholt and Gather [12] construct hybrid filters based on RM regression to combine the robustness of the repeated median with the better shift preservation of FMH filters.

Several filters are investigated, using either the central observation y_t or the median of all observations in the window $\tilde{\mu}_t$ as central subfilter. Instead of one-sided means they use one sided medians

$$\widetilde{\mu}_{t}^{F} = med\{y_{t-m}, ..., y_{t-1}\} \text{ and } \widetilde{\mu}_{t}^{B} = med\{y_{t+1}, ..., y_{t+m}\},\$$

and instead of the one-sided weighted averages they apply one-sided RM filters

$$RM_{t}^{F} = med\{y_{t-m} + m\widetilde{\beta}_{t}^{F}, ..., y_{t-1} + \widetilde{\beta}_{t}^{F}\}$$

where $\tilde{\beta}_t^F$ is the RM slope estimate based on the observations $y_{t-m}, ..., y_{t-1}$, and RM_t^B is defined analogously for the other half of the window. Since these subfilters make predictions for the central value, the procedures are called 'predictive' – or 'combined' if both, median and RM subfilters, are used.

In general RM based filters are not affected by trends and attenuate Gaussian as well as spiky noise well. The smoothest signal estimations are obtained by the ordinary RM filter, but on the other hand it also smoothes out shifts and trend changes. In contrast, the predictive RM hybrid filter

$$PRMH(y_t) = med\{RM_t^F, y_t, RM_t^B\}$$

can preserve trend changes and level shifts almost exactly – even within trends – but it attenuates Gaussian noise less efficiently, and like the other RM hybrid filters it is more affected by many outliers. Also, RM hybrid filters are designed for delayed signal extraction and hence, for online signal extraction different subfilters had to be applied.

2.3 Nested Filters

An approach for combining the smoothness of the moving average with the robustness and shift preservation of the running median is given by modified trimmed means (MTM) [13]. The idea is to calculate the median of all observations in the window and then 'trim', i.e. discard, those observations which deviate more than a specified multiple of a robust scale estimate, e.g. the *median absolute deviation about the median* (MAD) [14]

$$\widetilde{\sigma}_{t}^{MAD} = \underset{i=-m,\dots,m}{med} \{ |r_{t+i}| \}.$$

The arithmetic mean of the remaining observations is then taken as signal estimate in the centre of the time window. These MTM estimates are both robust against outliers and efficient for Gaussian noise. Also, they can preserve large shifts in an otherwise constant level better than ordinary running medians [15].

Since the location-based MTM deteriorates in trend periods Gather and Fried [16] extend this idea to the trimmed repeated median (TRM): Within each time window a RM regression line is fitted and the MAD calculated from its residuals for estimating the local variability [17]. Observations deviating more than a multiple of the residual MAD from the fitted line are trimmed, and the final signal estimate is derived by a least squares fit to the remaining observations. This TRM filter is almost as robust as a variant applying another RM regression in the second step, but it is more efficient for Gaussian errors.

To further improve the preservation of shifts, Bernholt, Fried, Gather and Wegener [18] use a smaller window width in the first step for the initial RM fit. Because of the nested design of the windows for the first and the second regression step, the prefix 'double window' (DW) is added to the estimates which results in DWMRM and DWTRM (see **Fig. 4**).



Figure 4 DWTRM fit to a single time window of width n=31: In the second step only the observations within the trimming boundaries around the RM line are used to calculate the least squares fit.

Using this double window technique considerably improves the performance of the RM filters concerning the preservation of shifts. In general, shifts which are large relative to the observational noise are traced more accurately than smaller shifts.

If the application allows for a relatively large outer window width, the signal estimation can also be improved by using a short inner window for the initial RM slope estimation and a larger outer window for the level estimation. First experiences show that the DWTRM filter seems even more promising for delayed signal extraction – keeping in mind the demands for robustness and the allowable time delay. Yet, these methods have not been investigated carefully in full online analysis.

2.4 Weighted Repeated Median Filters

In analogy to the popular weighted median (WM) filters, Fried, Einbeck and Gather [19] construct weighted repeated median (WRM) filters. While the former are based on the idea that a constant level is more likely for close-by observations, the latter filters assume the signal slope to be more likely to be the same in short time lags. Suitable symmetric bell-shaped (in delayed/retrospective analysis) or monotonic (in full online analysis) weighting schemes allow to use longer time windows than ordinary running medians or RM filters which correspond to uniform weights.

Considering *n* observations $(x_i, y_i)^{*}$, i=1,...,n, where the x_i are not necessarily equidistant, and two sets of integer weights w_i and \tilde{w}_j , the weighted repeated median is defined by

$$\widetilde{\beta}^{WRM}(x) = \operatorname{med}_{j=1,\dots,n} \left\{ \widetilde{w}_{j} \circ \left(\operatorname{med}_{i \neq j} \widetilde{w}_{i} \circ \frac{y_{i} - y_{j}}{x_{i} - x_{j}} \right) \right\}$$
$$\widetilde{\mu}^{WRM}(x) = \operatorname{med}_{j=1,\dots,n} \left\{ w_{j} \circ \left(y_{j} - (x_{j} - x) \widetilde{\beta}^{WRM}(x) \right) \right\}.$$

Here, the operator \circ symbolises replication, i.e. $w_i \circ y_i$ means that y_i is replicated w_i times. This newly defined method is then compared to L₁ and weighted L₁ filters.

Among other things, the study determines the minimal window width which is necessary for the investigated methods to resist a certain number h of successive outliers, while taking these deviant values into account if their number is larger than h. The reason for this lies in the fact that, moving a time window through a series of measurements, at some point the time series contains h subsequent outliers or 'spikes' (which are still regarded as a sequence of artefacts) while in the subsequent time window the presence of h+1 successive outliers with the same size and sign may already indicate a shift [20]. In this way, window widths are determined which allow for tracking shifts lasting at least h+1 observations while eliminating a smaller number of outliers.

For the RM weighting improves the adjustment to nonlinear trends, allows for larger window widths, and increases the efficiency, while for the L_1 filter weighting can increase robustness and efficiency.

For online signal extraction, the WRM filter tracks shifts better than the L_1 filter, which has some difficulties in distinguishing relevant from irrelevant patterns. The weighting reduces the bias of the RM, implying that the WRM also outperforms the standard RM filter in tracing shifts. Also, the WRM filter shows generally the smoothest signal estimations in application to time series, while the L_1 filter overshoots shifts and is wiggly. In conclusion, a suitably designed weighted RM filter can be recommended for online signal extraction. In the retrospective situation the weighted L_1 filters provide even better results than the WRM filters. Particularly for moderate outliers, weighted L_1 filters show the least biased results and further they trace large shifts with a smaller time delay.

However, if it is possible that several outlier patches occur close to each other and thus intrude into the same time window, the standard RM filter may still be the best choice because of its maximal breakdown point.

2.5 Extended Robust Regression Filters

In contrast to LMS filters, RM filters are more vulnerable to large outliers while they accommodate small outliers well (see e.g. [16] and [17]). Also, large outliers are usually easier to detect than small ones. Therefore, it is worthwhile to add automatic rules for outlier detection and replacement to the repeated median to increase the robustness of the signal estimation [21]. Likewise we can apply automatic rules for level shift detection to the RM filters investigated in [2] and [6].

Similarly to the nested filters approach, an observation is regarded as outlier if the corresponding absolute deviation from the current regression line is larger than a specified multiple d of a robust scale estimation, i.e. if $|r_{t+i}| > d \cdot \tilde{\sigma}_t$. However, here only the next, incoming observation is screened for outlyingness before entering the actualised time window by extrapolating the previous regression line. Detected outliers are replaced and no longer considered in the following analysis. In this way they lose their influence on the estimations. For certain 'worst case' scenarios, replacing outliers by the simple extrapolation of the regression line, gives better results than other 'down-sizing' replacement strategies, for the price of reduced Gaussian efficiency.

For the scale estimation, several robust estimators are investigated and their respective advantages elaborated. These are, in addition to the MAD (see Section 2.3), Rousseeuw's and Croux' S_n and Q_n estimators [22] and the 'length of the shortest half' (LSH) [23], [24]. The Q_n and the LSH scale estimator give the best results in case of many large outliers of similar size, but the Q_n provides better efficiency, especially when identical measurements occur, e.g. due to rounding.



Figure 5 Comparison of the simple (delayed) RM filter with its extended version including outlier and shift detection.

For shift detection, a simple majority rule is added to the filtering procedure: Considering the most recent *m* observations in the time window, the number of observations with residuals larger than a certain bound and same sign is counted. If this number exceeds m/2, this indicates a level shift and the procedure moves to the next window not overlapping the current one.

This rule enables the regression filters to detect and thus preserve shifts, and hence it overcomes the biggest disadvantage of the RM filter (see **Fig. 5**). Also, the delay in following shifts decreases – ideally to a minimal delay of $\lfloor m/2 \rfloor + 1$ time units. In this context, regression based filters with additional shift detection rules seem preferable to other shift preserving procedures like LMS or FMH filters.

Further, some simple rules can be added to overcome problems in the infrequent case that too many observations are identified as outliers and thus replaced [21].

The rules for outlier treatment and shift detection can also be applied for online signal extraction. However, the minimal delay of shift detection cannot be reduced further because of the necessary differentiation of shifts and outlier patches.

3 Applications

The different approaches described in the previous sections, produced different filters which seem promising for application to online monitoring data from intensive care. Especially filters based on the repeated median show good results. However, the choice of the appropriate filter should depend on the characteristics of the underlying signal whenever known.

Summarising the outcomes described above, the following recommendations can be given: For retrospective signal extraction the predictive RM hybrid (PRMH) filter (Sec. 2.2) seems to be the best choice if the signal is assumed to contain many jumps and trend changes, while the simple RM filter (Sec. 2.1) yields better results if many outliers but no abrupt changes are expected. A compromise between these two methods is given by the DWTRM filter (Sec. 2.3), while the RM filter in combination with additional rules (Sec. 2.5) works better than the simple RM filter in the occurrence of many outlier patches and level shifts (see Fig. 5).



Figure 6 Comparison of a nested RM (DWTRM) and an RM hybrid filter (PRMH) with the simple RM filter for retrospective application.

For online signal extraction, the weighted version of the RM filter (Sec. 2.4) seems the best choice at the moment but another promising approach based on the RM is currently under research, adapting the window width at each time point. To provide a comparison of the specific benefits of the proposed filters, we present some applications to intensive care time series here.

The comparison of the simple RM filter with its extended version (Sec. 2.5) in Fig. 5 shows how much a shift detection rule can improve the RM filter. However, local extremes, i.e. sudden trend changes, cannot be traced as well with this extended RM filter. In that case, the application of the PRMH (Sec. 2.2) or the DWTRM filter (Sec. 2.3) is more recommendable (see **Fig. 6**).

In Fig. 6 we see that the predictive RM hybrid filter (PRMH) traces the sudden shifts and local extremes very accurately. However, the PRMH signal shows the largest variability, especially in relatively constant periods, e.g. from 14:30h to 14:40h here. The simple RM filter output is the smoothest but smears sudden shifts and 'cuts' local extremes. As pointed out before, the DWTRM signal is a compromise between the RM and the PRMH filter output: It is smoother than the PRMH signal but traces trend changes and shifts better than the RM filter.

Real-time application of these filters implies a time delay of half the window width used for the signal extraction. Therefore, filters have been examined for their online application without any time delay. As displayed in Fig. 3, simple regression filters (Sec. 2.1) are suitable for this purpose but even the online version of the RM filter still possesses some disadvantages – such as the slow reaction to level shifts. Weighted RM filters (WRM, Sec. 2.4), which are under current research, can possibly improve upon simple online RM filters.

Online signal extraction by such methods can be used for improving monitoring systems for haemodynamic variables: Basing an alarm system on the signal instead of the actual measurements would not trigger alarms at the occurrence of single measurement artefacts or irrelevant patches of outliers. Therefore, we claim that such an alarm system will have a lower false alarm rate than alarm systems based on raw measurements. **Fig. 7** shows that basing the alarm system on the extracted RM online signal avoids false alarms due to artefacts. However, the system would still react to the sudden change of heart rate around 15:00h.



Figure 7 A filter based alarm system does not trigger alarms in case of measurement artefacts.



Figure 8 Depending on the alarm settings, the overall alarm rate of current monitoring systems can be very high – even in the absence of outliers.



Figure 9 A filter based alarm system can reduce the overall alarm rate drastically.

Further, an alarm system based on the data signal can also reduce the overall alarm rate: **Fig. 8** shows half an hour of systolic arterial blood pressure measurements with the alarm settings and alarms triggered by the currently used alarm system.

In **Fig. 9** it is demonstrated that - although these data do not contain any measurement artefacts - an alarm system based on the filter output can reduce this high alarm rate considerably. This result does not imply a decrease of the *false alarm* rate. However, a decrease of the number of total alarms may also be considered an improvement.

One should note that the figures and applications given here are only exemplary. The general superiority of a filter based alarm system has yet to be shown – based on a sensible and careful definition of the term 'false alarm'. The full assessment of such a new alarm system is one aim of a currently running clinical study conducted at the Hospital of the University of Regensburg and supported by the collaborative research centre SFB 475 at the University of Dortmund.

4 Discussion

The methods recommended for univariate signal extraction here, are based on a simple linear regression approach. The ordinary repeated median regression filter improves on running medians in trend periods but lacks the property of preserving sudden shifts. Different approaches to overcome this problem have been proposed and work well for particular situations, but there is no 'universal' procedure without any deficiencies. Double window TRM filters are promising for delayed signal extraction while weighted RM filters are hopeful candidates for the online analysis.

Further investigations show that median based filters are also robust against the presence of autocorrelations. Compared to procedures based on least squares, robust location or regression filters, based on the median, trimmed means or the repeated median respectively, gain relative efficiency in the frequent case of positive correlations. In that case, they also outperform filters incorporating the autocorrelations explicitly into the analysis (Fried [25], Fried and Gather [26]). In the infrequent situation of strong negative autocorrelations a Prais-Winsten transformation of the data is worthwhile and improves the ordinary RM filter. With the linear time RM update algorithm developed by Bernholt and Fried [5] or its advancement to a linear storage algorithm [12], the RM filter is computationally feasible even for high frequency data. This update algorithm also means linear computation time for all of the recommended filters. Outlier and shift detection as described in Section 2.5 does not add further computation time when using e.g. an $O(\log n)$ MAD update algorithm [18] for the scale estimation. Thus, an RM filter with such extensions offers an acceptable choice for signal extraction.

Another approach to improve the computational speed is proposed by Fried and Gather [27]. Dividing the time window into n_2 disjoint segments, each of length n_1 , the level within each segment is estimated by an ordinary median or by repeated median regression. Then the RM, or another procedure, can be applied to this pre-processed output window of width n_2 . Hence, the computation time can be shrunk by a factor n_1 when using a linear time algorithm and by n_1^2 when using an algorithm needing quadratic time.

For retrospective analyses, computation times are not as crucial as for realtime applications but they are still important because of the possible magnitude of the data sets. Hence, the computability of the filters should always be taken into account when choosing the 'right' filter.

For the filtering procedures described above, the window width n = 2m + 1 has been assumed to be fixed throughout. The suitable choice of the width is no trivial task and depends on statistical as well as medical demands. Larger window widths generally imply a smoother filter output, but they also increase the bias at shifts.

To tackle the problem of accurate tracing of level shifts Gather and Fried [16] introduce a procedure with a variable, data-adaptive choice of the window width for delayed signal extraction with the RM filter. This adaptive procedure improves the tracing of trend changes and level shifts, while still achieving smooth signal estimations and high robustness. Therefore, it is of special interest for future research to convert this method to the situation of signal extraction without any time delay to make it applicable for online monitoring in intensive care.

Using extracted signals for judging the patient's state of health may improve the work on intensive care units notably. However, the superiority of filter based alarm systems is a task yet to be completed, and there is still need for research on fully adaptable signal extraction methods best suitable in the online monitoring context. Nevertheless, the presented signal extraction filters provide a solid background for the development of enhanced alarm systems with lower false alarm rates.

Acknowledgements

We gratefully acknowledge the support of the German Research Foundation DFG through the collaborative research centre SFB 475 "Reduction of Complexity for Multivariate Data Structures".

4 References

- Tukey, J.W.: Exploratory Data Analysis, Reading, MA, Addison-Wesley, 1977.
- [2] Davies, P.L.; Fried, R.; Gather, U.: Robust Signal Extraction for Online Monitoring Data, Journal of Statistical Planning and Inference 122 (2004), pp. 65-78.
- [3] Rousseeuw, P.J.: *Least Median of Squares Regression,* Journal of the American Statistical Association 79 (1984), No. 388, pp. 871-880.
- [4] Siegel, A.F.: *Robust Regression Using Repeated Medians,* Biometrika 69 (1982), pp. 242-244.
- [5] Bernholt, T.; Fried, R.: *Computing the Update of the Repeated Median Regression Line in Linear Time*, Information Processing Letters 88 (3) (2003), pp. 111-117.
- [6] Gather, U.; Schettlinger, K.; Fried, R.: *Online Signal Extraction by Robust Linear Regression*, to appear in: Computational Statistics, Issue 02/2006.
- [7] Rousseeuw, P.J.: Multivariate Estimation with High Breakdown Point, Proceedings of the 4th Pannonian Symposium on Mathematical Statistics and Probability Vol. B, Dordrecht: Reidel, 1983.
- [8] Rousseeuw, P.J.; Hubert, M.: *Regression Depth*, Journal of the American Statistical Association 94 (1999), No. 446, pp. 388-402.
- [9] Heinonen, P., Neuvo, Y.: FIR-Median Hybrid Filters, IEEE Transactions on Acoustics, Speech, and Signal Processing 35 (1987), pp. 832-838.
- [10] Heinonen, P., Neuvo, Y.: FIR-Median Hybrid Filters with Predictive FIR Substructures, IEEE Transactions on Acoustics, Speech, and Signal Processing 36 (1988), pp. 892-899.

- [11] Heinonen, P., Kalli, S., Turjanmaa, V., Neuvo, Y.: Generalized Median Filters for Biological Signal Processing, Proceedings of the Seventh European Conference on Circuit Theory and Design (1985), pp. 283-386.
- [12] Fried, R., Bernholt, T., Gather, U.: *Repeated Median and Hybrid Filters*, to appear in: Computational Statistics and Data Analysis, Special Issue on 'Statistical Signal Extraction and Filtering' (2006).
- [13] Lee, Y., Kassam, S.A.: Generalized Median Filtering and Related Nonlinear Filtering Techniques, IEEE Transactions on Acoustics, Speech and Signal Processing 33 (1985), pp. 672-683.
- [14] Huber, P.J.: Robust Statistics, Wiley, New York, 1981.
- [15] Himayat, N., Kassam, S.A.: Approximate Performance Analysis of Edge Preserving Filters, IEEE Transactions on Signal Processing 4 (1993), pp. 2764-2776.
- [16] Gather, U., Fried, R.: *Methods and Algorithms for Robust Filtering*, Proc. in Computational Statistics, J. Antoch (ed.), Physica, Heidelberg, 2004, pp. 159-170.
- [17] Gather, U., Fried, R.: *Robust Estimation of Scale for Local Linear Temporal Trends*, Tatra Mountains Mathematical Publications 26 (2003), pp. 87-101.
- [18] Bernholt, T., Fried, R., Gather, U., Wegener, I.: *Modified Repeated Median Filters*, conditionally accepted for Statistics and Computing (2004).
- [19] Fried, R., Einbeck, J., Gather, U.: *Weighted Repeated Median Smoothing and Filtering,* submitted to Journal of the American Statistical Association (2005).
- [20] Imhoff, M., Bauer, M., Gather, U., Fried, R.: Pattern Detection in Intensive Care Monitoring Time Series with Autoregressive Models: Influence of the Model Order, Biometrical Journal 44-6 (2002), pp. 746-761.
- [21] Fried, R.: *Robust Filtering of Time Series with Trends*, Journal of Nonparametric Statistics 16 (3) (2004), pp. 313-328.

- [22] Rousseeuw, P.J., Croux, C.: Alternatives to the Median Absolute Deviation, Journal of the American Statistical Association 88 (1993), pp. 1273-1283.
- [23] Grübel, R.: Length of the Shorth, Annals of Statistics 16 (1988), pp. 619-628.
- [24] Rousseeuw, P.J., Leroy, A.: *A Robust Scale Estimator Based on the Shortest Half*, Statistica Neerlandica 42 (1988), pp. 103-116.
- [25] Fried, R.: *Robust Location Estimation under Dependence,* to appear in: Journal of Statistical Computation and Simulation (2006).
- [26] Fried, R., Gather, U.: Robust Trend Estimation for AR(1) Disturbances, Austrian Journal of Statistics, Vol. 34 (2) (2005), pp. 139-151.
- [27] Fried, R., Gather, U.: Fast and Robust Filtering of Time Series with Trends, Proc. in Computational Statistics, W. Härdle and B. Rönz (eds.), Physica, Heidelberg, 2002, pp. 367-372.