Computers and Chemical Engineering 48 (2013) 187–199

Contents lists available at SciVerse ScienceDirect

Computers and Chemical Engineering

journal homepage: www.elsevier.com/locate/compchemeng

Qualitative Representation of Trends (QRT): Extended method for identification of consecutive inflection points

Kris Villez[a,b,*], Christian Rosén[c], François Anctil[d], Carl Duchesne[e], Peter A. Vanrolleghem[a,d]

[a] BIOMATH, Faculty of Bioscience Engineering, Ghent University, Coupure Links 653, 9000 Ghent, Belgium
[b] Laboratory for Intelligent Process Systems, School of Chem. Eng., Purdue University, 480 Stadium Mall Drive, West Lafayette, IN 47907, USA
[c] Veolia Water – Solutions & Technology, AnoxKaldnes AB, Klosterängsvägen 11A, SE-226 47 Lund, Sweden
[d] Département de génie civil et de génie des eaux, Université Laval, 1065 av. de la Médecine, Québec, QC G1V 0A6, Canada
[e] Département de génie chimique, Université Laval, 1065 av. de la Médecine, Québec, QC G1V 0A6, Canada

ARTICLE INFO

Article history:
Received 6 September 2011
Received in revised form 3 May 2012
Accepted 14 August 2012
Available online 8 September 2012

Keywords:
Qualitative Trend Analysis (QTA)
Qualitative Representation of Trends (QRT)
Data mining
Qualitative analysis

ABSTRACT

In this paper a methodology for extraction of qualitative descriptions of data series is extended for proper inflection point recognition. The original method allows the extraction of qualitative information such as the sign of the first and second derivatives of the assumed signals underlying univariate noisy time series. However, it is not able to handle consecutive inflection points, i.e. inflection points which follow each other in time with no minimum or maximum in between them. To improve this method for proper assessment of inflection points, the original method is compared with three modifications of the original method. One of the alternatives based on repeated application of Witkin's stability criterion delivers better results for both identification of the qualitative descriptions and the locations in time of extrema and inflection points. Furthermore, the same modified method is shown to deliver the best fault diagnosis performance in a benchmark batch fermentation study. Furthermore, the same modified method is shown to deliver the best fault diagnosis performance in a benchmark batch fermentation study.

1. Introduction

Qualitative interpretation of time series has received fair attention in literature over the past 15 years. The use of qualitative information is often warranted and supported by the observation that an operator's reasoning or knowledge is to a large extent based on qualitative features in data series rather than their quantitative properties. Given that operators spend a large proportion of their time to the monitoring of trends in process measurements (Yamanaka & Nishiya, 1997), an automated assessment of process trends can facilitate operators in performing the task of process supervision. Part of the potential for such a tool lies in the context of fault detection and identification (Venkatasubramanian, Rengaswamy, & Kavuri, 2003; Villez, Keser, & Rieger, 2009; Villez, Rosén, Anctil, Duchesne, & Vanrolleghem, 2008). Indeed, information about trends may be addressed to the operator only in case of abnormalities. As such, an operator does not need to check normal data on a regular basis and can focus on true problems. This reduces the time spent on evaluation of normal data and will likely result in a faster analysis and reaction to abnormal situations. More advanced use of such tools may include the design of automated diagnosis and/or control systems that are based on qualitative assessment of process trends. Furthermore, methods focused on qualitative features are well-suited for data mining of process data (Villez et al., 2007).

In the context of automated control of biological wastewater treatment plants for nutrient removal, oxidation reduction potential (ORP) signals and the points within time series thereof where transitions from accelerating to decelerating behavior and vice versa occur (inflection points) have been addressed as key indicators (Andreottola, Foladori, & Ragazzi, 2001; Fuerhacker et al., 2001; Kim, Chen, Kishida, & Sudo, 2004; Li, Peng, Peng, & Wang, 2004; Plisson-Saune, Capdeville, Mauret, Deguin, & Baptiste, 1996; Ra, Lo, & Mavinic, 1999; Vanrolleghem & Coen, 1995; Wareham, Mavinic, & Hall, 1994; Yu, Liaw, Chang, & Cheng, 1998). As the inflection points in ORP signals can be used as indicators for the end of nitrification and denitrification processes in waste water treatment systems (Chang & Hao, 1996), an accurate assessment of inflection points with minimal delay is desired for such applications. In other biotechnological applications qualitative descriptions are used for model structure identification (Shaich, Becker, & King, 2001; Vanrolleghem & Van Daele, 1994).

Methods available for the assessment of Qualitative Representation of Trends can be divided into two main classes, based on the

* Corresponding author at: Eawag, Überlandstrasse 133, P. O. Box 611, 8600 Dübendorf, Switzerland.
E-mail address: kris.villez@gmail.com (K. Villez).

0098-1354/$ – see front matter © 2012 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.compchemeng.2012.08.010

principles of their core methodology. The first group is based on the training of a clustering (unsupervised) or classifier (supervised learning) model. Wang and Li (1999) use PCA-based clustering while Rengaswamy and Venkatasubramanian (1995) use neural networks. An essential characteristic and potential weakness of these methods is the necessity for training of the models on historical data prior to application. This may lead to errors related to extrapolation to new data for which the models are not trained. Methods of the second class do not require such explicit training and can be applied directly to any time series. Most of these are based on the fitting of polynomial functions in contiguous windows. A prior assessment of one or more parameters (e.g. noise level) for acceptable fits is however essential. This requirement limits the use of the techniques in case of changing noise characteristics and may lead to errors due to extrapolation as well. The methods of Charbonnier, Garcia-Beltan, Cadet, and Gentil (2005), Dash, Maurya, Venkatasubramanian, and Rengaswamy (2004), Flehmig, Watzdorf, and Marquardt (1998) share this characteristic. Within this second class, an exception to this is only found with the method explained by Bakshi and Stephanopoulos (1994) which does not require an assessment of noise characteristics. It is the latter Qualitative Representations of Trends (QRT) method and modifications thereof which are the subject of this study.

Qualitative analysis of time series has been judged promising to solve many problems, including data mining, fault detection and identification (FDI) and reaction end-point detection. However, effective and critical comparison of the available techniques is very limited. For this reason, this study includes a first benchmarking study in qualitative analysis. The original QRT method and the proposed modification are compared in terms of accuracy, computational demand and fault detection and identification (FDI) performance.

It is safe to say that the selected method, called Qualitative Representation of Trends (QRT), is fairly complicated. It serves therefore to evaluate what is favorable about this method. First of all, neither model training or assessment of statistical properties of the signal (e.g. noise level) are necessary prior to application of the method. This advantage results from the use of a feature selection rule, called Witkin's stability criterion (Witkin, 1983), which is aimed at the discrimination between noisy and acceptable qualitative features in a time series. This heuristic rule is based on the intuitive notion that true or meaningful qualitative features are more difficult to filter from a time series than noisy ones. Its application results in implicit control of the complexity of the resulting qualitative representation. Because of the heuristic, it is a very generic method which requires little tuning or adjustment once implemented. Second, the QRT method is based on the cubic spline wavelet decomposition which means that the method implicitly assumes continuity of the first and second derivatives of the assumed signal underlying the time series. This was a fair assumption to make in the context of the original study (Villez, 2007), as the focus lay on the analysis of pH and oxidation–reduction potential time series which were not characterized by visible discontinuities. Third, an initial comparison with the Qualitative Trend Analysis method, proposed by Dash et al. (2004), indicated some problems with the QTA method that were considered prohibitive for application for the particular case study. For example, a preliminary comparison (Villez, 2007) between the QTA and QRT method indicated that the QTA method as implemented does not enforce continuity of the first and second derivatives. In addition, the same noise-free signal resulted in a large variability of qualitative descriptions when repeated for different noise sequences. The original QRT method did not exhibit such large variability. A third observation was that the QTA method obtains qualitative descriptions through control of the parametric complexity of the fitted polynomials by choosing between zero-, first- and second-order approximations. This does not automatically imply control of the complexity of the qualitative descriptions as was demonstrated in Villez (2007).

To be fair, some arguments can be brought against the QRT method just as well. A first is that the method is fairly complex, both conceptually and in terms of implementation. At the same time, little existing software exists that can easily be re-used for this purpose and the literature offers little guidance on proper implementation. A second is that the original QRT method does not allow for consecutive inflection points (between two consecutive extrema) to be identified correctly and a third is that this method cannot tackle discontinuities in an effective manner, due to the cubic spline wavelet decomposition. It was decided to tackle the first and second remark based on investigative research while the third remark was not considered relevant to the intended application. As such, the modifications of the original QRT method are reported with the aim of tackling the second remark on QRT in particular. Next to that, gaps in the original method are completed. The implementation in Matlab/Octave is available publicly, by which the first criticism of QRT is also alleviated.

In what follows, materials and methods are explained first (Section 2), including the benchmark simulations set up to properly test the developed modifications of the QRT method and the four alternative QRT-based methods. Results are presented in the following section after which Section 4 follows. At the end of this paper, the major conclusions of this work are provided. In supplementary materials, the original and developed QRT methods are benchmarked for fault detection and identification (FDI) of the Penicillin fermentation model of Birol, Undey, and Çinar (2002) and Monroy, Villez, Graells, and Venkatasubramanian (2011).

2. Materials and methods

In the following sections, the data set as used for the benchmarking study reported in this article is explained first. This benchmarking study is focused on the identification of qualitative representations. After that, the deployed methods for qualitative time series analysis are explained. To this end, the paradigm of using primitives as templates for qualitative behavior is presented after which the original method by Bakshi and Stephanopoulos (1994) for Qualitative Representation of Trends (QRT) is explained in detail. Following that, modifications to the original method are presented which result in three alternative methods. In a last section, the performance measures used to evaluate the different methods are given. A second benchmarking study is provided in Supplementary Materials. This study evaluate the same methods for identification of qualitative representations as well as for fault detection and identification (FDI). This additional study further supports the conclusions made on the basis of the first study. It also leads to fair further criticism against using the QRT method for Fault Detection and Identification problems in noisy circumstances.

2.1. Benchmark simulations

The prime purpose of this paper is to extend the original QRT method so that the identification of consecutive inflection points becomes possible. To effectively evaluate whether a particular proposed alternative method fares better or worse at this task, a benchmark time series data set is produced for which the true qualitative descriptions are available as a reference. It is for this reason that simulated data, rather than real data, are used. Furthermore, the QRT method heavily relies on cubic spline wavelet, which inherently assumes that the obtained signal is approximated well by a signal which is continuous up to the second derivative. As such, the simulations deliver noise-free signals which are continuous up to the second derivative. These noise-free simulations are used to determine the true qualitative representation by simple differentiation.

To meet this requirement, a simple fed-batch fermentation process is simulated. The biokinetic model makes use of conventional Monod and inhibition switching functions of which the parameters can be set up so as to to produce consecutive inflection points in the simulated data. The fed-batch process is simulated for a fixed length of 480 min. At the start, the reactor contains 100 l of an aquatic solution containing two reagents (A and B). Both reagents react with a third reagent (C) which is added in solution at the beginning of the fed-batch operation with a fixed pump flow rate for a given time. The detailed model is given in Appendix A. All parameters and variables are listed with their (initial) values in Appendix B. A noise-free value of the reactor concentration of reagent C is taken every 5 min, leading to 5760 measurements in each simulation.

Three parameters are varied in this simulation study. First, the time length during which this pump is active is varied between 5 and 20 min in steps of 5 min. Depending on this time length, the reagent B is completely converted or not within the simulated batch time. Secondly, additional variation in the simulated results is obtained through applying five different values for the inhibition constant in the model (K_I) between 10^{-5} to 10^{-1} mol l^{-1}, each separated with a factor 10. This affects the smoothness of the simulated profiles. For low values of the inhibition constant one obtains a smoother profile which makes the inflection points harder to observe, i.e. they are not as crisp. Finally, the measurement noise standard deviation was set to values between 10^{-5} and 10^{-1} mol l^{-1} each separated with a factor 10. The values for the first two parameters were deliberately set so to obtain a set of 20 distinct target qualitative representations, including simulated time series with three consecutive inflection points present in the time series. The different noise levels permit to evaluate the robustness of the applied techniques against noise. The three parameters are varied according to a full factorial design, thereby leading to 100 ($4 \times 5 \times 5$) simulations for a given noise sequence. Each of these was repeated for 100 different noise sequences. The Signal to Noise Ratio (SNR, ratio of the variance of the noise-free signal to the variance of the noise sequences) computed over the whole data set ranges from 3.0210^8 (highest noise level) to 3.02 (lowest noise level), with factors of 100 apart. The ratio of the noise standard deviation to the concentration range (overall minimum to maximum value) ranges from 1.4×10^{-3}% to 14% (with factors of 10 apart). The true qualitative representations are obtained by analysis of the derivatives of the noise-free signal, $x_C(k)$. The actual qualitative analysis by means of the QRT methods takes the noisy time series, $y(k)$, as input.

2.2. Primitives

Common to most qualitative analysis methods for time series is that qualitative behavior is described on the basis of the sign of the first and second derivatives as inferred by the method of choice. Seven unique qualitative behaviors result, each of which is referred to as a primitive and to which a unique alphabetic character is assigned. Fig. 1 shows those seven primitives as well as the characters as used in this work. In addition, it shows two additional characters used in this work when only the sign of the first derivative is known and/or relevant. For upward trends, U (Upper) is used, for downward L (Lower) is used. The steady F primitive can always be used without ambiguity.

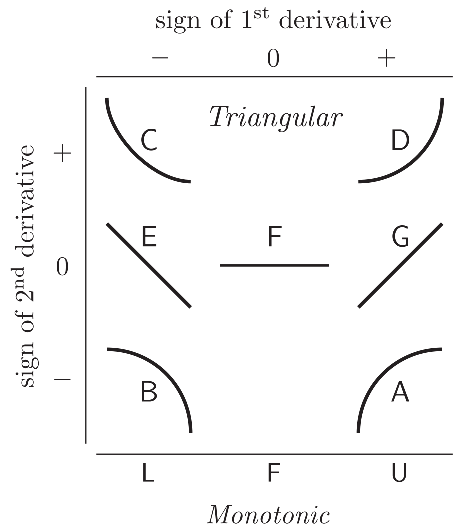

Fig. 1. Primitives for qualitative analysis.

2.3. Original method

The original method by Bakshi and Stephanopoulos (1994) proceeds in four major steps, referred to as (1) wavelet decomposition, (2) trending of details, (3) construction of the Wavelet Interval Tree and (4) feature selection. Each of these steps is explained separately in the following paragraphs. The first paragraphs of Section 3 report detailed results for a single time series which can be inspected for visualization of the following abstract explanation.

2.3.1. Step 1 – Wavelet decomposition

In essence, wavelet theory is apopular approach to analyze signals simultaneously in the frequency domain and the time domain. Several works and tutorials can be found today as a result. Daubechies (1992), Meyer (1993), Strang and Nguyen (1996) are suggested for extensive, theoretical study and (Torrence & Compo, 1998) for a more intuitive introduction to the subject.

In the case of the QRT method, the discrete wavelet transform (DWT) is used. This transform consists of passing a given signal through a dyadic filter bank consisting of corresponding pairs of high-pass and low-pass Finite Impulse Response (FIR) filters. Fig. 2 shows a scheme of this filtering process. In what follows, a given level in the filter bank is referred to as the wavelet scale (p). Practically speaking, one obtains the low-pass filtered signals by convolution with the so called *scaling function* (ϕ_p). The resulting low-pass signals are called the *approximations* ($a_p(t)$). In parallel, one also convolutes each low-pass signal with the so called *wavelet function* (ψ_p) which results in high-pass signals called *details* ($d_p(t)$). The maximum number of levels in the wavelet analysis, P, is set to $log_2(N)$ where N is the number of points in the studied series.

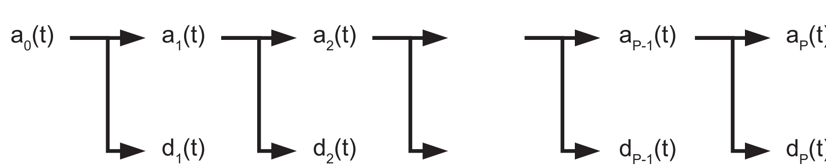

Fig. 2. Scheme of the wavelet filter bank.

Indeed, assessing the behavior at frequencies lower than $2/N$ has little meaning when the series length is only N.

In the QRT method one specifically uses the cubic spline wavelet. This wavelet has the particular property that a zero-crossing in the detail d_p is indicative of an extremum (maximum or minimum) in the approximation from which it is derived (a_{p-1}) (Mallat, 1991). Also, an extremum (maximum or minimum) in the detail signal corresponds to an inflection point in the approximation at the previous level. This property of the cubic spline wavelet is only shared with the Gaussian kernel, as used in (Witkin, 1983). However, the Gaussian kernel has infinite support while the cubic spline wavelet has compact support. As a result, the cubic spline wavelet is the only practically viable wavelet for the QRT method.

2.3.2. Step 2 – Trending

The trending step is by far the easiest to implement. In this step, each detail is checked for zero-crossings and extrema. These correspond to extrema and inflection points, respectively, in the approximations. Importantly, the original method requires to remove inflection points when they are not separated by an extremum in between them. Each time this occurs, the inflection point corresponding to the maximum absolute value for the detail signal is selected while the rest are discarded. This implies that the original method cannot identify consecutive inflection points by design.

Following the identification of the essential points in an approximation, one can obtain the qualitative description of that approximation. It is noted here that it is practically impossible to obtain the primitives E, F and G in a noisy scenario. Indeed, positive identification of the E, F and G primitives requires the detail signal to be constant for a period of time. Moreover, the F primitive implies a constant zero value for a period of time. As such, any practical result for noisy time series is expected to contain only A, B, C and D primitives. In Section 4, available options to enable identification of linear trend episodes are discussed. Each primitive in the sequence together with its start and end time is called a *triangular episode*. The contiguous set of triangular episodes defines the *triangular representation* of the respective approximation. Next to the triangular episode, a *monotonic episode* is defined as a contiguous time frame within which the sign of the first derivative is the same. The qualitative description with monotonic episodes is obtained by aggregating contiguous triangular episodes with the same sign for the first derivative. The primitives U, F and L as discussed before are used for monotonic episodes. In analogy to the triangular representation, the *monotonic representation* of a signal is defined as the contiguous set of monotonic episodes.

2.3.3. Step 3 – Wavelet Interval Tree

Construction of the Wavelet Interval Tree is the step in which structural relations between the qualitative representations at each wavelet scale are obtained. The Interval Tree concept was originally proposed in Witkin (1983) in the context of scale-space filtering, a technique which inspired the original QRT method. The underlying idea is that qualitative features disappear gradually during the filtering process. It is indeed so that the cubic spline wavelet filter only removes qualitative features such as extrema and inflection points. This means that a tree can be constructed which links features at a given wavelet scale with a feature at the subsequent wavelet scale. The original work does not specify the algorithm execute this task. This link-finding problem is posed as a dynamic program (DP) aimed at finding the feasible connection of essential points at neighboring wavelet scales with minimum misalignment as the cost function. This misalignment is computed as the sum of absolute shifts in time necessary to align the essential points at the coarser scale with those in the more detailed wavelet scale. Feasibility of an alignment requires that only essential points of the same type (see above) are aligned, that chronological order is preserved and that all essential points at the coarser scale are aligned with an essential point at the finer scale. The algorithm to solve this DP is adapted from the Viterbi algorithm for maximum likelihood state estimation of Hidden Markov Models (HMMs, Forney, 1973).

Following the alignment of essential points, one defines a branch–leaf relationship between the episode at the coarser scale (*branch*) and each individual episode (*leaf*) in the finer scale which is found between the essential points aligned with the essential points of the parent episode. Following alignment and identification of the branch–leaf relationships, one obtains one or more tree structures. There are as many trees as there are qualitative features at the coarsest wavelet scale. The whole set of such trees is referred to as the Wavelet Interval Tree. The construction of the WIT is executed on the basis of the triangular representations at each wavelet scale. Upon finalization of this Triangular Wavelet Interval Tree (TWIT), one can easily obtain a monotonic equivalent (Monotonic Wavelet Interval Tree, MWIT) by aggregation of the individual triangular episodes into monotonic ones as described above.

It has been indicated by Bakshi and Stephanopoulos (1994) and confirmed during implementation that the cubic spline wavelet filter is not perfect in the sense that spurious qualitative features may be generated at coarser wavelet scales which are not present at the finer scale. This occurs seldom but can prevent sound execution of step 4 in the QRT method. Therefore, qualitative features at a coarser scale are removed when they are not aligned with a feature at the finer scale. Care has been taken to do this so that the modified qualitative representation remains internally consistent (e.g. B should be followed by C following implied continuity constraints of the first and second derivatives). This removal is done before the branching step. For each pair of subsequent wavelet scales, the executed steps are thus (1) alignment, (2) spurious feature removal and (3) branching.

2.3.4. Step 4 – feature selection

In the final step of the QRT method, the multiple qualitative representations at different wavelet scales are reduced to one, final representation. In order to arrive at the final representation, one starts with the most coarse Monotonic representation available, the one at the last wavelet scale. For each episode in this representation one tracks over how many scales the same monotonic episode spans without branching into more than one episode. At the point where the considered episode branches into more than one branch, one counts the number of scales over which each of those connected episodes spans without branching themselves.

Now, one decides whether the simple episode is kept or whether it should be replaced with the more complex set of branched episodes according to *Witkin's stability criterion* (Witkin, 1983). This criterion states that if the number of scales over which the simple episode spans is smaller than the average number of scales over which the branched episodes span, then one should replace the simple representation with the more complex one. In the other case (larger), one does not replace the given episode. This heuristic was originally developed in Witkin (1983) and is based on the observation that the number of scales over which a qualitative feature spans and does not branch out into more features is indicative of its relevance. One continues this process of conditional replacement for all episodes until no more episode can be replaced by a more complex presentation based on this criterion. The procedure then ends and one has obtained the final monotonic qualitative representation. To obtain the triangular qualitative representation, one simply selects the triangular episodes which correspond to the episodes in the final monotonic episodes.

2.4. Modifications to the original method

In what follows, the proposed modifications to the original QRT method are explained. A modification which permits the identification of consecutive inflection points is discussed first. After that, three different ways to select triangular episodes for the final triangular representation are proposed.

2.4.1. Modification of step 2: allowing for consecutive inflection points

One proposed modification is related to the existence of consecutive inflection points. In the original QRT method, consecutive inflection points are removed in step 2 (Trending) by retaining one inflection point only between each pair of extrema that follow each other. However, the identification of consecutive inflection points may be relevant at least for some fed-batch processes in general and for systems with buffered states, such as redox or acid–base reaction systems. To make this possible, the removal of inflection points in the trending step is simply skipped in all proposed alternatives.

2.4.2. Modification of step 4: selecting triangular episodes

Three alternatives are proposed for the selection of triangular episodes in step 4. The selection of the monotonic episodes remains the same.

Alternative 1: Top The first manner of selecting the triangular episodes works in the same way as in the original method. This means that, upon selection of the monotonic episodes for the final representation, the triangular episodes present at the most detailed wavelet scale of those over which the selected monotonic episodes span are selected. Step 4 in the procedure thus remains the same as in the original method. Note that, due to the modification in step 2 actual results can differ. In particular, this alternative method is expected to deliver the most complex triangular representations of all tested methods.

Alternative 2: Bottom In this alternative method, one selects the triangular episodes at the coarsest scale at which the selected monotonic episodes exist. In this way, a lower number of triangular episodes is expected to be retained compared to the 'Top' alternative since cubic spline wavelet filtering removes qualitative features such as inflection points.

Alternative 3: Witkin The third alternative is based on repeated application of Witkin's stability criterion. To this end, one regards the portion of the Triangular WIT corresponding to a monotonic episode in the final monotonic representation. This is done by selecting those triangular episodes which correspond to the considered monotonic episode over all scales the latter spans. Now Witkin's stability criterion is applied to this portion of the Triangular Wavelet Interval Tree (TWIT) in exactly the same way it was done for the whole Monotonic WIT. After doing so for all monotonic episodes in the final monotonic representation, one has obtained the final triangular representation. This 'Witkin' alternative is expected to deliver triangular representations with complexity in between the 'Top' and 'Bottom' alternatives.

2.5. Software

The original method and the proposed alternatives were implemented as a self-consistent toolbox for Matlab and Octave. It has been tested succesfully in several versions of Matlab (R14, R2009b, R2010b, R2011a) and Octave (v3.4.2). This software is provided as supplementary material to this article and made public under the GPLv3 license.

2.6. Performance evaluations

In the following paragraphs, the performance measures evaluated in this study to compare the different methods are explained. As pointed out by (Dash et al., 2004), several characteristics are desirable for qualitative analysis of time series, including accuracy and speed which are considered the most important ones. To evaluate accuracy, two measures are computed while for speed one suffices.

Fraction of correct qualitative sequences. The first and simplest way to assess the quality of the proposed methods is to compute the fraction of qualitative representations for which the sequence of primitives is an exact match to the true representation. This fraction is computed for each of the proposed methods (4) and for each applied noise level (5), and averaging over all fed-batch simulations (20) and noise sequence repetitions (100).

Misalignment. Even if the obtained and targeted representation match qualitatively, it may be expected that the identified extrema and inflection points are not found at the same location in time. To evaluate such temporal mismatch, the misalignment is computed as the sum of absolute distances in time that essential points (extrema, inflection points) have to be shifted to match the target representation exactly. This is only computed for the cases where the sequence of identified primitives matches the target representation perfectly.

Speed. Next to the measures of accuracy described above, the computational time needed for each method is also evaluated. This is done with the *tic* and *toc* commands in Matlab. In addition, the time needed for steps 1 and 2 combined, step 3 and step 4 are computed. For the purpose of these evaluations, all computations were executed with Matlab R2010B on a desktop computer equipped with a Pentium 4 3 GHz CPU, 1.49 GB of RAM and MS Windows operating system.

3. Results

Results are presented in two parts. The first part shows results for one exemplary time series taken from the benchmark simulations. By means of these results, the differences among the applied methods are demonstrated in detail. In the second part, the performance results are shown and interpreted for the whole set of benchmark simulations.

3.1. Detailed example

Results are shown for the benchmark simulation with parameters set as follows: $T_{pump} = 20$ min, $K_I = 10^{-2}$ mol l^{-1}, and $\sigma = 0.01$ mol l^{-1} ($SNR = 302$). The most important benefit of these results is the visual interpretation to the method, thereby rendering the method less abstract. In addition, the particular example allows to demonstrate the original and alternative methods and is informative with respect to a more general analysis further in this text. The noise-free signal as well as the noisy time series are shown in Fig. 3.

3.1.1. Using the original method

Fig. 4 shows the results of the first step in QRT, *wavelet decomposition*. The left side of the figure shows the original signal at the top and the approximations below from the first wavelet scale to the last. The right side shows the corresponding details.

The second step, *trending*, is demonstrated in Fig. 5 for wavelet scale 9. At this scale, three consecutive extrema appear in the detail signal, corresponding to three consecutive inflection points in the approximation at scale 9. If these consecutive inflections were to be kept, then the resulting triangular representation would be DABCBC (Fig. 5b). However, with the original method only one

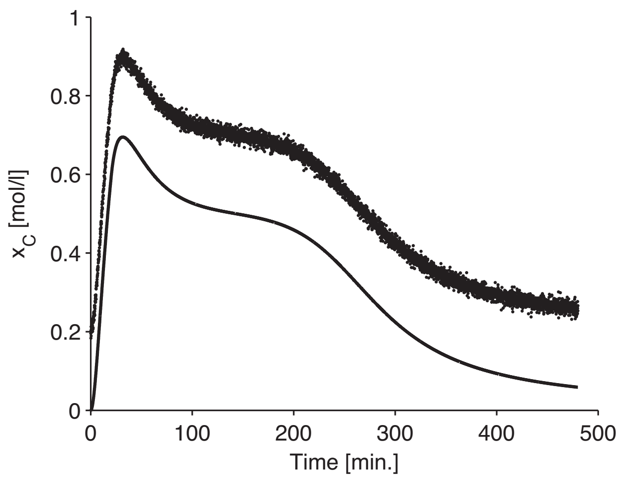

Fig. 3. Noise-free signal (line) and noisy time series (dots). The noisy time series is off-set for visualization.

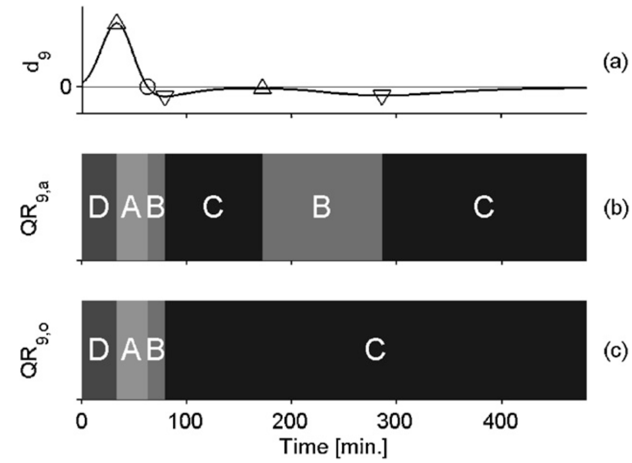

Fig. 5. Trending (step 2) at wavelet scale 9: detail signal with identified maxima (△), minima (▽) and zero-crossing (◯) (a), qualitative representation prior to inflection point selection (b) and after inflection point selection (c).

of the three consecutive inflection points is retained, namely the one at 79.2′ in the simulation, thus leading to the triangular representation DABC at wavelet scale 9 (Fig. 5c). In what follows, the scale number will correspond to the wavelet scale at which the corresponding detail signal was found.

Fig. 6 shows the result of the alignment process as part of the third step, *construction of the Wavelet Interval Tree*. All the obtained triangular representations are shown with lines indicating the alignment across scales. One remarkable observation is that the second inflection points at scale 8 and 9 are connected with each other while being quite afar in time. This is attributed to the fact that the removed inflection points at scale 8 and 9 are not the corresponding ones. As will be shown further, the alternative methods do not exhibit this effect of inflection point removal. Fig. 7 shows the final result of the alignment procedure. In this figure, all essential points are now aligned with their corresponding location at the first, most detailed, wavelet scale. Also, episodes which are not split across several scales are grouped and marked with a single character where space allows. For instance, the B episode at scale 12 branches into a BCDAB sequence at scale 7.

Fig. 8 shows the Monotonic Wavelet Interval Tree (MWIT) which follows from aggregation of the triangular episodes into monotonic ones. DA sequences are replaced with a U episode and BC sequences with a L episode. Note that the latter characters are not shown to permit better visualization of the branches in the MWIT. A node is shown for each episode at the highest scale it is present. Lines indicate the branches in the wavelet tree. For instance, the L episode (32.1–480′) at scale 12 remains until scale 8 after which it is split into three episodes (LUL).

The last step is determining which episodes are relevant. One starts with the UL representation at scale 12 and verifies whether any of the episodes in this representation should be replaced with a finer representation. The first episode U is present over scale 12 to 6 (7 scales). At scale 5, this episode branches into two new ones, which span over 5 and 1 scale respectively. Thus, the average number of scales over which they span without further split is 3 whereas the parent episode spans over 7 scales. Following Witkin's stability

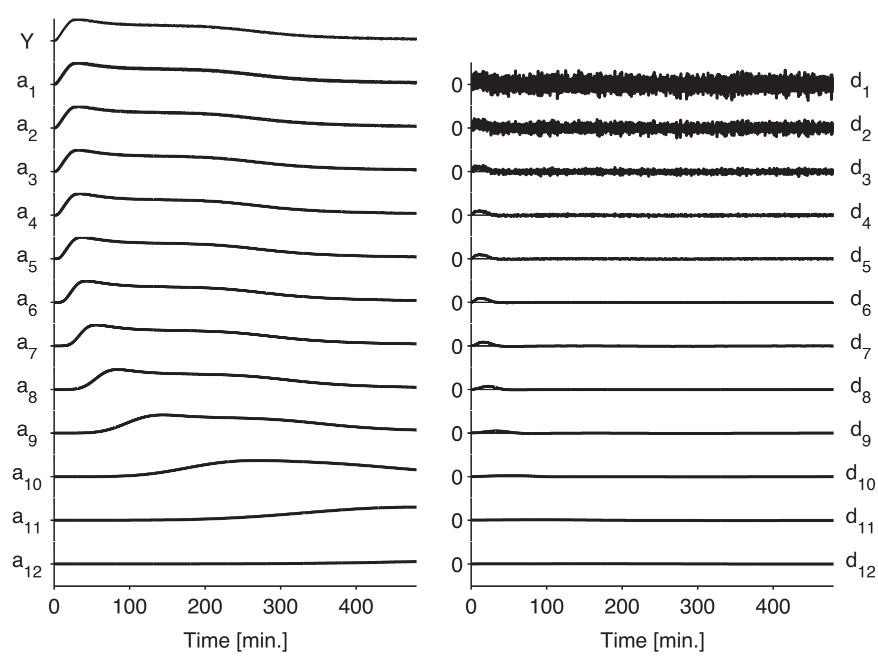

Fig. 4. Wavelet decomposition (step 1). Left panels show the original time series and approximations. The right panels show the details.

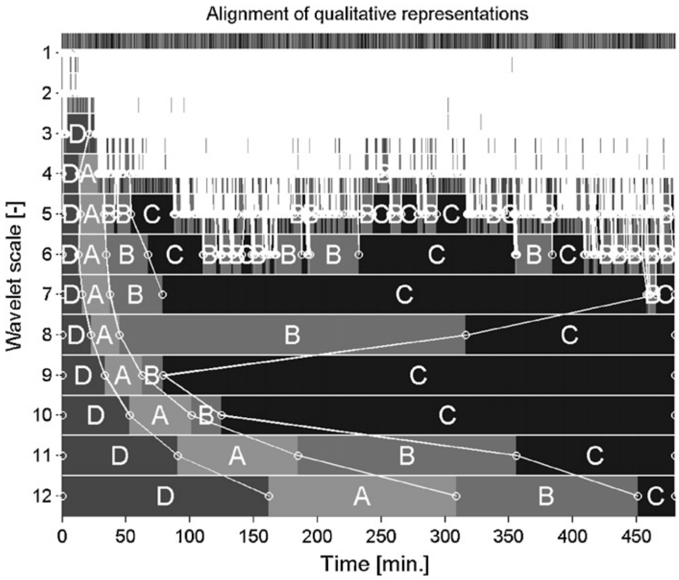

Fig. 6. Construction of the Wavelet Interval Tree (WIT, step 3) – alignment of essential points. White lines connect essential points that are aligned with each other. Alphabetic characters indicate identified primitives where space allows.

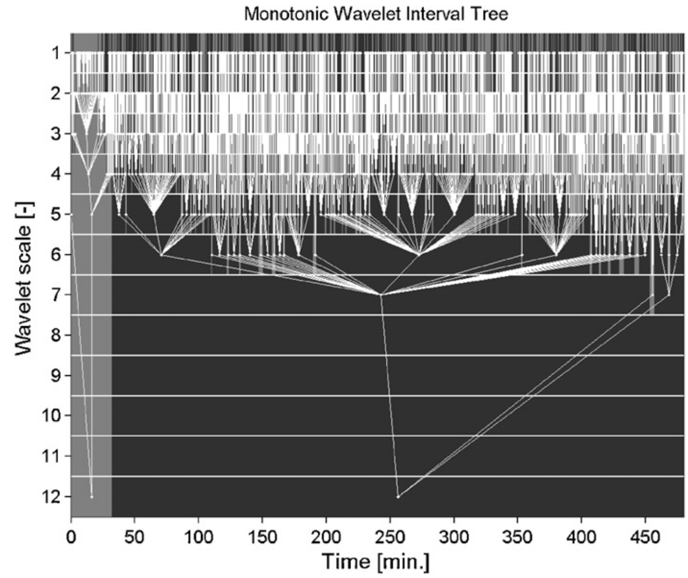

Fig. 8. Construction of the Wavelet Interval Tree (WIT, step 3) – Monotonic Wavelet Interval Tree (MWIT). White lines visualize the interval tree explicitly.

criterion, one keeps the coarse representation, not the finer one. Based on similar reasoning, the second episode at scale 12 is not split either. The final monotonic representation is thus UL. The maximum between the two episodes is located at 32.1 min. This corresponds very well with the true representation which is also UL with a maximum at 31.6 min.

The triangular representation now follows easily. The final monotonic representation contains a U which spans until scale 6. By inspecting Fig. 7, one finds the corresponding DA sequence at scale 6. The identified inflection point (maximum slope) is found at 20.9′. For the second L episode in the monotonic episode one finds a BC sequence at scale 8 with the inflection point at 463.3′. As such, the final triangular representation obtained with the original method is thus DABC with 20.9, 32.1 and 463.3 min as the locations in time of the corresponding essential points. These values are compared later with the true values and the values obtained with the alternative methods.

3.1.2. Using the newly proposed methods

Now the results for the newly proposed methods are presented based on the same example. For all proposed alternative methods, consecutive inflection points are kept in the trending step (step 2). This means, for example, that the representation of Fig. 5b is used, rather than the one in Fig. 5c. Naturally, the triangular representations and resulting alignments in Fig. 9 are different than the ones obtained earlier (Fig. 6). In particular, one can see that the second inflection point at scale 9 is now connected to a closer inflection point at scale 8 which was removed with the original method. As a result, the alignment result appears more regular in this case. The Triangular Wavelet Interval Tree (TWIT), common to all three proposed alternative methods, is shown in Fig. 10. Note that the

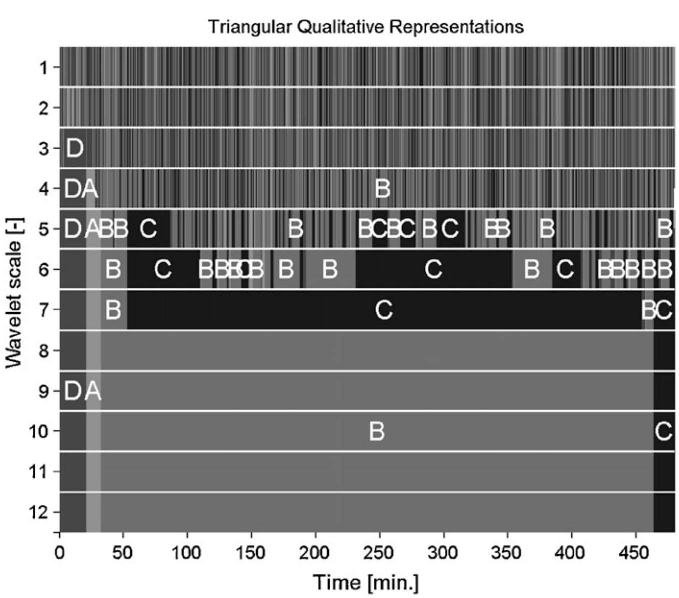

Fig. 7. Construction of the Wavelet Interval Tree (WIT, step 3) – Triangular Wavelet Interval Tree (TWIT). Alphabetic characters indicate identified primitives where space allows.

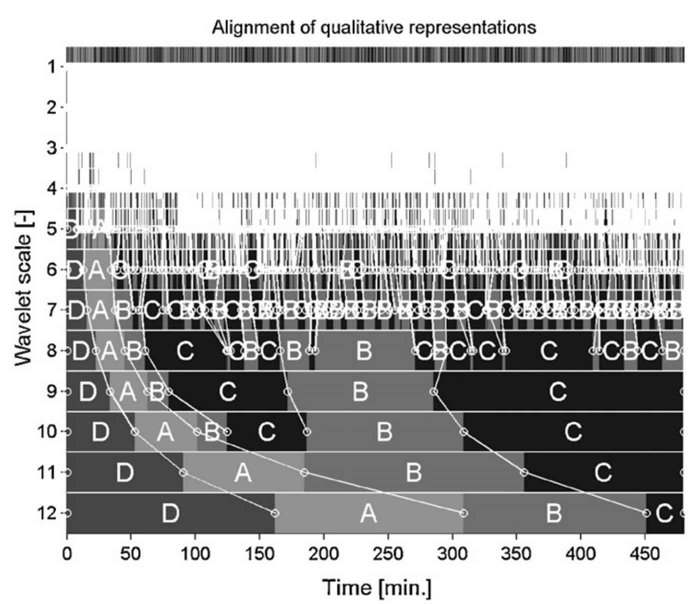

Fig. 9. Construction of the Wavelet Interval Tree (WIT, step 3) – alignment of essential points for the alternative methods. Alphabetic characters indicate identified primitives where space allows.

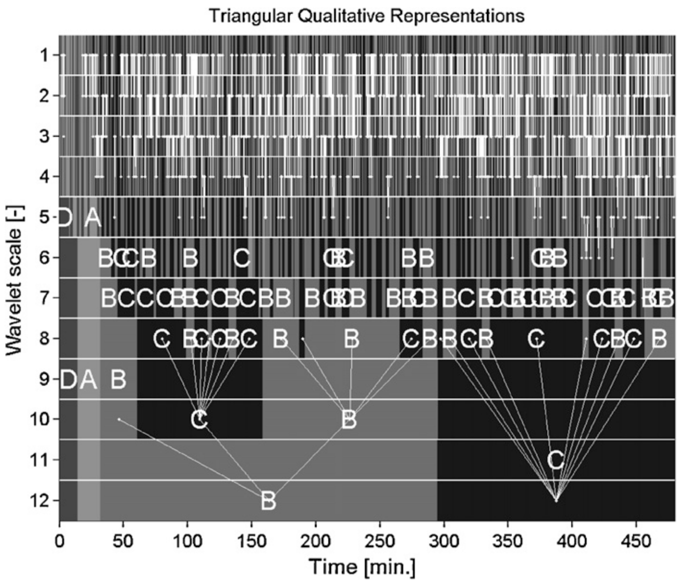

Fig. 10. Construction of the Wavelet Interval Tree (WIT, step 3) – Triangular Wavelet Interval Tree (TWIT) for the alternative methods. White lines visualize the interval tree explicitly for the monotonic episodes included in the final representation. Alphabetic characters indicate identified primitives where space allows.

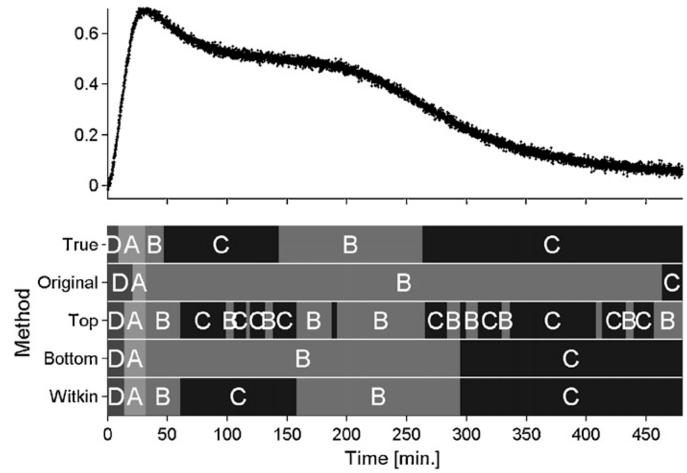

Fig. 11. Final result with all evaluated methods: original time series (top) and triangular representations (bottom). Alphabetic characters indicate identified primitives where space allows.

Monotonic Wavelet Interval Tree remains the same in this case (Fig. 8). As a result, the final monotonic representation is exactly the same as with the original method.

To obtain the final triangular representation with the first alternative (top), one selects the triangular episodes in the same way as the original method, namely at the most detailed scale the monotonic episodes are present. This is at scale 6 for the first episode, at scale 8 for the second episode. Following inspection of Fig. 10, the resulting triangular representation that results is thus DABCBCBCBCBCBCBCBCBCBCBCB. The short hand notation $\mathsf{DA(BC)_{11}B}$ is proposed for such long representations. Quite clearly, this result is too complex when compared to the true representation DABCBC. This is likely the main reason why consecutive inflection points are removed in the original method.

The second alternative (bottom) consists of selecting the triangular episodes at the most coarse scale that the monotonic episodes are present. This is at scale 12 for both monotonic episodes. As such, by means of Fig. 10, one finds that the resulting triangular representation is DABC. Clearly, this sequence is too simple compared to the true sequence. Note however that it is the same as the one obtained with the original method.

The third and last alternative (Witkin) evaluates the stability criterion over the parts of the triangular wavelet tree that correspond to the already selected monotonic episodes. The three corresponding parts of the Triangular WIT can be observed in Fig. 10. Application of Witkin's stability criterion for the leftmost part is very simple as the sequence remains to be DA over all considered scales. For the second part, the situation is different. Here, a B episode is found at scale 12 which branches into a BCB sequence at scale 10. The first of these four episodes spans three scales; the other two span two scales. This means that the average span of the branched episodes is 2.33 which is larger than the span of the coarse, parent episode (span = 2). As such, the application of the stability criterion leads to acceptance of the proposed split. Based on the same criterion, none of these episodes is split further. The same holds for the C episode found at scale 12. This episode spans over 4 scales and branches into 10 episodes ($\mathsf{(CB)_5}$) at scale 8 spanning an average of 1 scale. This means that the given episodes are not split further. With this, the execution of the 'Witkin' alternative method is completed and the DABCBC triangular representation results, thus giving the true qualitative sequence.

Fig. 11 shows the results obtained with all methods as well as the true representation. The true locations of the essential points are at 9.2, 31.5, 47.1, 143 and 263.5 min. It is clear that only the third alternative (Witkin) delivers the correct sequence of primitives. In this case, the identified location of the essential points are 14, 32.1, 60.7, 158.3 and 194.83 min. As such, the identified locations are relatively accurate. The results in the next section allow to elaborate further on this. Another interesting observation is that the original method and the second alternative (Bottom) deliver the same qualitative sequence but not the same locations for the essential points. More specifically, it appears that the locations for the identified inflection points are more accurate for the second alternative (Bottom). This is due to the inconsistent alignment resulting from the removal of inflection points at neighboring wavelet scales which do not correspond to each other.

3.2. Benchmarking results

In the overview of the benchmarking study, results related to the quality and speed of computation are shown for the original as well as all proposed alternative methods. In addition, results related to the accuracy of the monotonic representations are presented. Finally, an assessment of computational requirement is presented as well.

3.2.1. Fraction of correct qualitative sequences

The fraction of correctly identified qualitative sequences is shown as function of the applied method for inflection point identification and noise level in the top panel of Fig. 12. One can see that with increasing noise, the performance of the method is reduced irrespective of the chosen method. For instance, for the monotonic representations, the correct fraction reduces from 95% at the lowest noise level ($\sigma = 10^{-5}$ mol l^{-1}, $SNR = 3.0210^{8}$) to 10% for the highest level ($\sigma = 10^{-1}$ mol l^{-1}, $SNR = 3.02$). Next, one can see that the performance is reduced further when the obtained triangular representations are considered, irrespective of the method. This is expected as a correct triangular representation requires a correct monotonic representation. The subsequent reduction in performance is dramatic for all proposed methods, except for alternative 3 (Witkin) where the overall performance is 80% at the lowest noise level and drops to 10% at the highest level. As such, this alternative can be considered best on the basis of the fraction of correctly

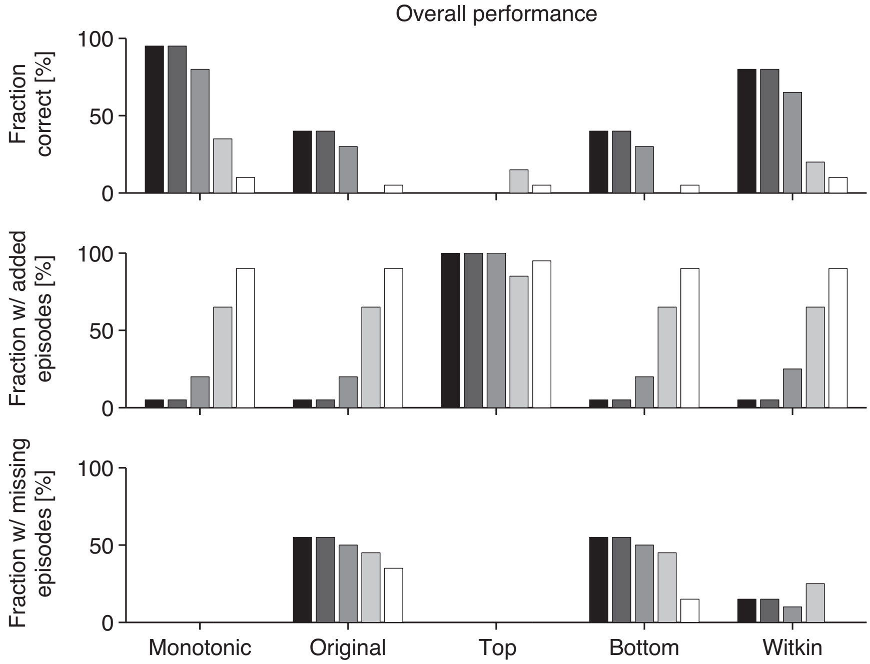

Fig. 12. Accuracy of identified qualitative sequences as function of method and noise level (increasing noise level from left to right in each group). Fraction of correct qualitative sequences (top), fraction of representations with unwanted episodes (middle) and fraction of representations with missing episodes (bottom).

identified qualitative sequences. The other methods fare worse with alternative 1 (top) being the worst (15% correct identification at best). The original method and alternative 2 (bottom) deliver practically the same result with results ranging between 5% and 40%.

The middle and bottom panel serve to indicate where the major causes of misidentification lie. The middle panel shows the fraction of results for which the obtained result includes episodes which are not present in the true representation. The bottom panel shows the fraction of results for which the obtained results miss one or more episodes present in the true representation. Note that the fractions do not necessarily sum to 100% over the three panels as it is possible that a qualitative representation is simultaneously missing some episodes in the true representation as well as including episodes not in the true representation. However, the simultaneous missing of episodes and inclusion of unwanted episodes does not occur frequently so that general trends can still be observed easily. In the middle panel one can see that for all methods except for alternative 1 (top) the fraction of representations which include unwanted episodes increases with increasing noise for all methods. For the alternative 1 (top) method, the fraction is high irrespective of the noise level. This is due to the applied heuristic, namely choosing the triangular episodes at the most detailed level included in the selected monotonic episodes. We note further that the results for the original method, alternative 2 (bottom) and alternative 3 (Witkin) are very similar. As such, the fraction of representations with unwanted episodes is not a driver to choose one of the alternatives over the original method.

In the bottom panel, results are more diverse. First, one can see that the fraction of too simple representations is zero for the monotonic representation. This suggests that the QRT method is biased towards complex representations. However, the true monotonic representation is UL for all cases, which cannot be simplified much further. A zero fraction is also obtained for alternative 1 (top). Given that (1) this method choses the most complex one available conditional to the monotonic representation and (2) the monotonic representation is never too simple, this should be expected. Once more, the original method and alternative 2 (bottom) deliver similar results. The fraction of results with missing features drops from 55% to 35% (original) and 15% (alternative 2). Finally, alternative 3 (Witkin) leads to the missing of episodes at the first four noise levels (10-25%) and none at the highest noise level.

3.2.2. Misalignment

In the above paragraphs only the sequence of primitives was considered for comparison. While some general trends are clearly observed, it serves to evaluate whether the identified locations of the essential points (extrema, inflection points) match the true locations well. To this end, the misalignment is computed for those results for which the qualitative sequence matches the true sequence exactly. This misalignment is averaged over all repetitions and simulations and shown in Fig. 13. It is observed easily

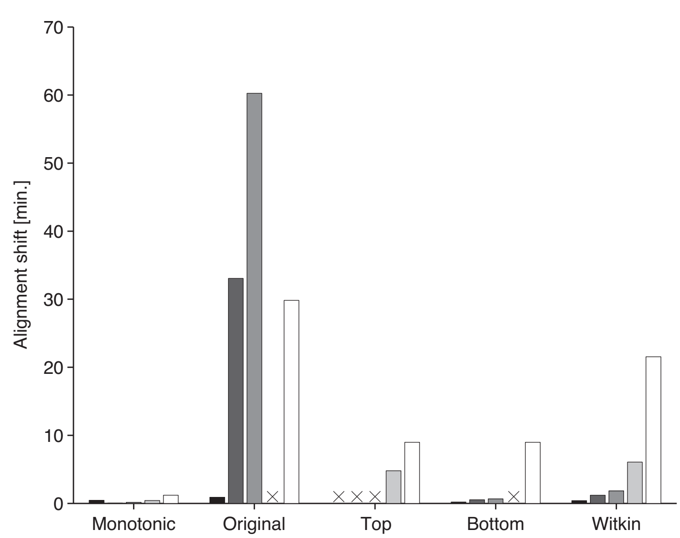

Fig. 13. Misalignment as function of method and noise level (increasing noise level from left to right in each group). Crosses (×) indicate where misalignment is not computed.

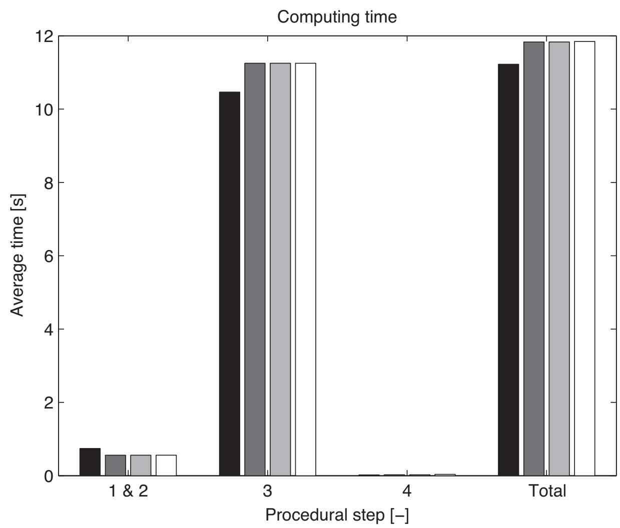

Fig. 14. Average total computational time for $\sigma = 10^{-3}$ mol l^{-1}; each group shows the time for the four methods (left to right: original and alternatives 1–3).

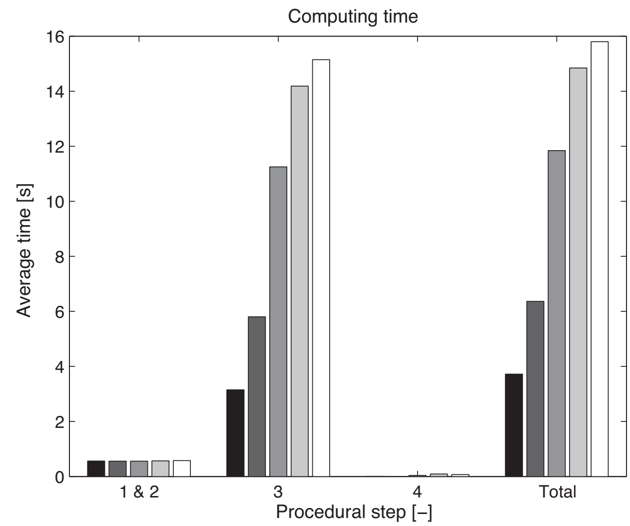

Fig. 15. Average step-wise computational time for the alternative method 3 (Witkin); each group shows the time for different noise levels increasing from left to right ($\sigma = 10^{-5}$ mol l^{-1} to $\sigma = 10^{-1}$ mol l^{-1} with factors of 10 apart).

that the shift for monotonic representations is relatively small, less than 1.5 min irrespective of the noise level. As such, the method does well at locating the (single) extremum (maximum) in the time series conditional to a correctly identified qualitative sequence. In contrast, the results for triangular representations with the original method are much worse. Indeed, the time shift required to match the true locations is larger for this method than for any of the alternative methods with values ranging from 0.9 s ($\sigma = 10^{-1}$ mol l^{-1}) to 60.3 s ($\sigma = 10^{-3}$ mol l^{-1}). Moreover, no clear trend with respect to noise level is seen which contrasts with the other results. Indeed, for the alternative methods, an increase in misalignment is seen with increasing noise levels. This is intuitively expected as increased noise can mask the true location of an extremum or inflection point. Alternative method 1 delivers a lower misalignment than the other methods. However, it is by far the worst method based on the previously discussed fraction of correct qualitative sequences. The misalignment values for alternative method 2 (bottom) are lower than for alternative 3 (Witkin). Of course, it should be noted in this case that the Witkin method delivered better results on the basis of fraction of correct qualitative sequences.

Among these observations, the high misalignment values obtained for the original method are the most surprising. Detailed inspection of the results indicated that this is largely attributed to the removal of inflection points in step 2 (trending) of the QRT method which makes consistent alignment of inflection points in step 3 (Wavelet Interval Tree) impossible for a large fraction of the simulated time series. This effect was already demonstrated by means of the detailed results for one time series discussed earlier. Both alternative methods 2 and 3 deliver acceptable results for the location of inflection points. This suggests that even when the assumption that only one single inflection point can exist between extrema, the original method is not the best to find its location.

3.2.3. Computational time

In addition to the quality of the produced results, computational time is a crucial element in method evaluation given that most applications relate to on-line monitoring and diagnosis. Fig. 14 shows the total computational time for the four methods at the middle noise level ($\sigma = 10^{-3}$ mol l^{-1}). It is clear that the total computational time does not change much with respect to the chosen method. It is a little lower for the original method (11.23) than for the other methods (alternative 1 and 2: 11.83; alternative 3: 11.84). The obtained values suggest that on-line application is feasible for many processes with relevant dynamics in time scales of minutes and higher for the given sampling rate. One should expect a higher computational demand with higher sampling rates. However, this is not investigated further in this work. The time needed for alternative 3 is just slightly higher than alternative 1 and 2 because of the repeated application of Witkin's heuristic which is not part of alternative 1 and 2. Step 3 (wavelet tree construction) clearly is the most demanding step. This is due to the need for dynamic programming to solve the sequence alignment problem. This alignment is faster for the original method because a lesser number of features have to be aligned following removal of consecutive inflection points. It is this difference which makes the total computational time a little lower for the original method. The time needed for step 2 in the original method is slightly higher because of the time needed for removal of consecutive inflection points. However, this difference is not large enough to offset the reduced demand for step 3.

Fig. 15 shows the average computational time for complete analysis as well as for each constituting step for alternative method 3 (Witkin). The total computational time increases from 3.72 s at the lowest noise level to 15.8 s at the highest. It is easily seen, that this effect is largely due to a similar trend in the computational time for step 3 (Wavelet Interval Tree). This, in turn, can be explained by the fact that with increasing noise levels more features are present in the qualitative representations at several scales which makes the dynamic program for alignment (step 3) larger in size and the resulting Wavelet Interval Tree more complex. As such, this result is not very surprising.

4. Discussion

4.1. General aspects and limitations of the methodology

With this work, the original QRT method as well as three alternatives based on the original method were evaluated on two benchmark data sets. The first study is based on a simulation of time series which match the desired qualitative representation exactly and was presented in this article. The second study is based on a realistic benchmark fermentation model developed specifically for research in advanced monitoring and diagnosis. This study is reported in Supplementary Materials. For this study, both the accuracy of the QRT methods as well as the accuracy of QRT-based fault diagnosis were investigated.

In short, the QRT method consists of a procedure to process data towards a qualitative representation. All currently available methods share this characteristic, i.e. they are defined as a procedure. One disadvantage is that such procedures are not related to formal definitions of optimality. In contrast, generative approaches have been shown to lead to formal definitions for optimality in the context of latent variable modeling (Tipping & Bishop, 1999). A generative approach to qualitative analysis does not exist today and is therefore considered for further research.

The original publication reporting on QRT (Bakshi & Stephanopoulos, 1994) left some practical aspects of the methods open to the reader. In particular, the original work does not explicitly indicate how one should align essential points from scale to scale to construct the Wavelet Interval Tree. A dynamic programming algorithm was settled for to execute this task.

For noisy data, the QRT method cannot identify E, F and G primitives, which have zero second order derivatives. For this reason, the benchmarking study was focused on the identification of episodes with non-zero first and second derivatives. As such, the results are particularly relevant to settings where one is not interested in the identification of time windows with zero value for the first and second order derivative. Other methods are able to do so (Dash et al., 2004). Alternatively, one may be able to combine the QRT method as is with another one particularly oriented at the identification of such episodes. This was done as an ad hoc solution in Villez et al. (2009) by combination with the method of Cao and Rhinehart (1995). The method by Flehmig et al. (1998), which allows identifying time episodes with a zero value for a (single) particular derivative, can also be used for this purpose. In addition, the wavelet framework lends itself to shrinkage which can be used to force the detail signals towards zero (Donoho & Johnstone, 1994). For an objective comparison with available methods, a more general benchmark study should naturally include episodes with zero first and second order derivatives in the true representations (e.g. E, F, G).

It is also noted that the original method is not particularly designed with on-line applications in mind. Indeed, except for the first two steps (wavelet filtering and trending), it is necessary to reprocess all provided information in steps 3 and 4 (Wavelet Interval Tree and heuristic selection) as new information becomes available. While the computational requirements permit such an approach for processes with relevant dynamics in the minute scale and higher, other methods may still do better. As a last remark, it is concluded that the existence of a generally valid, optimal method for qualitative analysis has not been provided as of yet.

4.2. Method performance and selection

Based on the presented results, it is clear that the QRT method works well in terms of first order trends (monotonic representation) at very low noise levels. However, the performance drops dramatically with increasing noise level. While this is expected, it is clear from the results that the method tends to be biased towards more complex representations. This stands in contrast to statistical methods for model selection where simplicity is preferred over complexity when in doubt or left without sufficient evidence, a concept known as Occam's razor. This is an unfortunate result as this means that the QRT method is not very robust to noise. This holds for both Monotonic and Triangular representations and for both benchmark studies. These observations are also true with respect to Fault Detection and Identification as demonstrated by the second benchmark study in supplementary materials.

Potential improvements of the method may result from tuning certain aspects of the method. For instance, one may consider to analyze the qualitative representations at certain wavelet scales only, by ignoring some of the more detailed scales or coarsest scales. Furthermore, one may replace Witkin's stability criterion with another heuristic better suited to a particular application. Naturally, such tuning may render the method less attractive.

With respect to the triangular representation, the third alternative method (Witkin) is shown to be the best in terms of finding the proper qualitative sequence and in terms of fault identification accuracy. In addition, this method permitted the best identification of the qualitative representation as well since the misalignment remained fairly low at low noise levels. It is not surprising that one of the alternative methods comes out as the best since the original method cannot identify consecutive inflection points by design. The question then remains whether one should still choose the original method if the identification of a single inflection point between extrema is sufficient. The answer is clearly no since this is shown to lead to severe misalignment of the identified inflection point. This result is due to the fact that the removal of inflection points in series of consecutive inflection points does not guarantee that the removed inflection points at different wavelet scales correspond to each other. As a better alternative to such ends, one can apply the first alternative method (top) with an additional heuristic as follows. If the representation obtained via this method contains consecutive inflection points, select the inflection point with maximum slope (value for detail signal) at the highest scale at which the corresponding monotonic episode exists. As a final remark, it is noted that the differences in computational requirements are minimal, thus not being a driving factor in the choice for this method.

4.3. Future work

With this work, two benchmarking studies were executed to compare the original QRT method for qualitative time series analysis with some proposed alternative methods based on the original method. The obtained benchmarking results allows to select a best method for qualitative analysis when (1) zero values for the second order derivative are of no concern and (2) consecutive inflection points may exist. Nevertheless, the benchmark simulations were purposely adjusted to the known limitations of the QRT method. To this day, comparison with other methods for qualitative analysis is limited to a preliminary exercise reported in Villez (2007).

It is the opinion of the authors that objective benchmarking is crucial to advancement in the particular area of qualitative analysis. To make this possible, two important remarks can be made already. First, not all qualitative methods have been set up for the exact same purpose (e.g. steady-state identification only vs. full qualitative representation, univariate vs. multivariate time series) and execution (e.g. on-/off-line). Naturally, this should be reflected in such a benchmarking study. Second, the benchmarking study should compare results to the true qualitative representations while being realistic for real-world applications, which is challenging. A possible pathway is to provide reference representation based on expert opinion. This can be motivated easily since initial efforts in qualitative analysis were proposed on the basis of the idea that human experts rely on a qualitative interpretation of data rather than a quantitative one. It is only natural then to compare the qualitative representations of computer-based methods with (human) expert-based analysis.

5. Conclusions

In this paper, an existing method for Qualitative Representation of Trends is modified to allow identification of consecutive inflection points and better location in time. Such identification is relevant for certain processes yet could not be achieved with the original method. Three different methods are obtained by means of changes at two points in the QRT procedure. By means of both a detailed example and two benchmarking studies, it is shown that the alternative method which includes a repeated application of Witkin's stability criterion allows proper identification of consecutive inflection points provided the signal-to-noise ratio is low. Furthermore, it is the best for the intended context of application and only requires a small increase in computational burden.

The benchmarking studies as executed as part of this study are the first of their kind. The observed results indicate that the selected QRT method is not as good as one may expect on previous, more practically oriented studies, including studies written by the authors of this paper. This shows that benchmarking is essential to determine whether the qualitative analysis problem is inherently difficult or whether better tools are available to this end. Such extensive benchmarking is considered for future research.

Acknowledgements

This work was originally supported by the Institute for Encouragement of Innovation by means of Science and Technology in Flanders (IWT) and has benefited from additional support by the ICIS Distinctive Signature at Idaho National Laboratory (INL). Peter Vanrolleghem holds the Canada Research Chair in Water Quality Modelling. The first author wishes to thank Prof. V. Venkatasubramanian for the provided opportunity to continue this research.

Appendix A. Fed-batch process model

The bioconversion model consists of two reactions written as follows:

$$\begin{aligned} A + C &\rightarrow B \\ B + C &\rightarrow D \end{aligned} \tag{A.1}$$

The complete model simulates the flow rate (q), reactor volume (v) and four chemical species concentration (x_A, x_B, x_C, x_D) based on the following set of six Ordinary Differential Equations (ODEs):

$$\begin{aligned} \dot{Q} &= \frac{1}{\tau} \cdot (Q_{SP} - Q) \\ \dot{V} &= Q \\ \dot{x}_A &= \frac{\dot{v}}{v} \cdot x_A - r_1 \\ \dot{x}_B &= \frac{\dot{v}}{v} \cdot x_B + r_1 - r_2 \\ \dot{x}_C &= \frac{\dot{v}}{v} \cdot x_C - r_1 - r_2 + \frac{\dot{v}}{v} \cdot x_{C,in} \\ \dot{x}_D &= \frac{\dot{v}}{v} \cdot x_D - r_2 \end{aligned} \tag{A.2}$$

In this model, the actual flow rate is simulated to be the first-order response to changes in the inflow rate setpoint (q_{SP}), in turn evaluated as:

$$Q_{SP} = \begin{cases} Q_{max}, & t \leq T_{pump} \\ 0 & \text{otherwise} \end{cases} \tag{A.3}$$

The conversion rates (r_1, r_2) in the ODE model are computed as follows:

$$\begin{aligned} r_1 &= q \cdot \frac{x_A}{K_A + x_A} \cdot \frac{x_C}{K_C + x_C} \\ r_2 &= k \cdot \frac{K_I}{K_I + x_A} \cdot \frac{x_C}{K_C + x_C} \cdot x_B \end{aligned} \tag{A.4}$$

Noisy time series are generated by adding white noise to the sampled concentration of species C:

$$y(k) = x_C(k) + e(k), \quad e(k) \sim N(0, \sigma) \tag{A.5}$$

Appendix B. Model parameters and variables

Parameter	Description	Value	Unit
σ	Measurement standard deviation	10^{-5}–10^{-1}	mol l^{-1}
K_A	Affinity constant	1	mol l^{-1}
K_C	Affinity constant	1	mol l^{-1}
K_I	Inhibition constant	10^{-5}–10^{-1}	mol l^{-1}
T_{pump}	Duration of feeding phase	10–25	min
k	Specific reaction rate	0.1	min^{-1}
q	Maximum conversion rate	0.1	mol l min^{-1}

Variable	Description	Initial value	Unit
Q_{max}	Maximal flow rate	5	l min^{-1}
V	Reactor volume	100	l
x_A	Reactor concentration	1	mol l^{-1}
x_B	Reactor concentration	0	mol l^{-1}
x_C	Reactor concentration	0	mol l^{-1}
x_D	Reactor concentration	0	mol l^{-1}
$x_{C,in}$	Influent concentration	2	mol l^{-1}

Appendix C. Supplementary Data

A supplementary benchmark study and Matlab/Octave code associated with this article can be found, in the online version, at http://dx.doi.org/10.1016/j.compchemeng.2012.08.010.

References

Andreottola, G., Foladori, P., & Ragazzi, M. (2001). On-line control of a SBR system for nitrogen removal from industrial wastewater. *Water Science and Technology*, *43*(3), 93–100.

Bakshi, B. R., & Stephanopoulos, G. (1994). Representation of process trends – Part III. Multiscale extraction of trends from process data. *Computers and Chemical Engineering*, *18*, 267–302.

Birol, G., Undey, C., & Çinar, A. (2002). A modular simulation package for fed-batch fermentation: Penicillin production. *Computers and Chemical Engineering*, *26*, 1553–1565.

Cao, S., & Rhinehart, R. R. (1995). An efficient method for on-line identification of steady state. *Journal of Process Control*, *5*, 363–374.

Chang, H. C., & Hao, O. J. (1996). Sequencing batch reactor system for nutrient removal: ORP and pH profiles. *Journal of Chemical Technology and Biotechnology*, *67*, 27–38.

Charbonnier, S., Garcia-Beltan, C., Cadet, C., & Gentil, S. (2005). Trends extraction and analysis for complex system monitoring and decision support. *Engineering Applications of Artificial Intelligence*, *18*, 21–36.

Dash, S., Maurya, M. R., Venkatasubramanian, V., & Rengaswamy, R. (2004). A novel interval-halving framework for automated identification of process trends. *AIChE Journal*, *50*, 149–162.

Daubechies, I. (1992). *Ten lectures on wavelets*. Philadelphia: SIAM Press.

Donoho, D. L., & Johnstone, I. M. (1994). Ideal spatial adaptation by wavelet shrinkage. *Biometrika*, *81*, 425–455.

Flehmig, F., Watzdorf, R., & Marquardt, W. (1998). Identification of trends in process measurements using the wavelet transform. *Computers and Chemical Engineering*, *22*, S491–S496.

Forney, D. G. J. (1973). The Viterbi algorithm. *Proceedings of the IEEE*, *61*, 268–327.

Fuerhacker, M., Bauer, H., Ellinger, R., Sree, U., Schmid, H., Zibuschka, F., et al. (2001). Approach for a novel control strategy for simultaneous nitrification/denitrification in activated sludge reactors. *Water Research*, *34*, 2499–2506.

Kim, J.-H., Chen, M., Kishida, N., & Sudo, R. (2004). Integrated real-time control strategy for nitrogen removal in swine wastewater treatment using sequencing batch reactors. *Water Research*, *38*, 3340–3348.

Li, Y., Peng, C., Peng, Y., & Wang, P. (2004). Nitrogen removal from pharmaceutical manufacturing wastewater via nitrite and the process optimization with on-line control. *Water Science and Technology*, *50*(6), 25–30.

Mallat, S. G. (1991). Zero crossings of a wavelet transform. *IEEE Transactions on Inform. Theory*, *37*, 1019–1033.

Meyer, Y. (1993). *Wavelets and operators*. UK: Cambridge University Press.

Monroy, I., Villez, K., Graells, M., & Venkatasubramanian, V. (2011). Fault diagnosis of a benchmark fermentation process: A comparative study of feature extraction and classification techniques. *Bioprocess and Biosystems Engineering*, *35*, 689–704. http://dx.doi.org/10.1007/s00449-011-0649-1

Plisson-Saune, S., Capdeville, B., Mauret, M., Deguin, A., & Baptiste, P. (1996). Real-time control of nitrogen removal using three ORP bending-points: Signification, control strategy and results. *Water Science and Technology*, *33*(1), 275–280.

Ra, C., Lo, K., & Mavinic, D. (1999). Control of a swine manure treatment process using a specific feature of oxidation reduction potential. *Bioresource Technology*, *70*, 117–127.

Rengaswamy, R., & Venkatasubramanian, V. (1995). A syntactic pattern-recognition approach for process monitoring and fault diagnosis. *Engineering Applications of Artificial Intelligence*, *8*, 35–51.
Shaich, D., Becker, R., & King, R. (2001). Qualitative modelling for automatic identification of mathematic models of chemical reaction systems. *Control Engineering Practice*, *9*, 1373–1381.
Strang, G., & Nguyen, T. (1996). *Wavelets and filter banks*. UK: Wellesley-Cambridge Press.
Tipping, E., & Bishop, C. (1999). Probabilistic principal component analysis. *Journal of the Royal Statistical Society: Series B, Statistical Methodology*, *61*, 611–622.
Torrence, C., & Compo, G. P. (1998). A practical guide to wavelet analysis. *Bulletin of the American Meteorological Society*, *79*, 61–78.
Vanrolleghem, P., & Coen, F. (1995). Optimal design of in-sensor-experiments for on-line modelling of nitrogen removal processes. *Water Science and Technology*, *31*(2), 149–160.
Vanrolleghem, P., & Van Daele, M. (1994). Optimal experimental design for structure characterization of biodegradation models: On-line implementation in a respirographic biosensor. *Water Science and Technology*, *30*(4), 243–253.
Venkatasubramanian, V., Rengaswamy, R., & Kavuri, S. N. (2003). A review of process fault detection and diagnosis – Part II: Qualitative models and search strategies. *Computers and Chemical Engineering*, *27*, 313–326.
Villez, K. (2007). Multivariate and qualitative data analysis for monitoring, diagnosis and control of sequencing batch reactors for wastewater treatment. Ph.D. Thesis. Ghent University, Ghent, Belgium.
Villez, K., Keser, B., & Rieger, L. (2009). Qualitative representation of trends (QRT) as a tool for automated data-driven diagnostics for on-line sensors. In *Proceedings of the 7th IFAC symposium on fault detection, supervision and safety of technical processes (SafeProcess2009)* Barcelona, Spain, 30 June–July 3, 2009, appeared on CD-ROM.,
Villez, K., Pelletier, G., Rosén, C., Anctil, F., Duchesne, C., & Vanrolleghem, P. A. (2007). Comparison of two wavelet-based tools for data mining of urban water networks time series. *Water Science and Technology*, *56*(6), 57–64.
Villez, K., Rosén, C., Anctil, F., Duchesne, C., & Vanrolleghem, P. A. (2008). Qualitative representation of trends: An alternative approach to process diagnosis and control. *Water Science and Technology*, *57*(10), 1525–1532.
Wang, X. Z., & Li, R. F. (1999). Combining conceptual clustering and principal component analysis for state space based process monitoring. *Industrial and Engineering Chemistry Research*, *38*, 4345–4358.
Wareham, D. G., Mavinic, D. S., & Hall, K. J. (1994). Sludge digestion using ORP-regulated aerobic-anoxic cycles. *Water Research*, *28*, 373–384.
Witkin, A. P. (1983). Scale-space filtering. In *Proc. 8th Int. Joint Conf. Art. Intell., Karlsruhe, Germany* (pp. 1019–1022).
Yamanaka, F., & Nishiya, T. (1997). Application of the intelligent alarm system for the plant operation. *Computers and Chemical Engineering*, *21*, S625–S630.
Yu, R.-F., Liaw, S.-L., Chang, C.-N., & Cheng, W.-Y. (1998). Applying real-time control to enhance the performance of nitrogen removal in the continuous-flow SBR system. *Water Science and Technology*, *38*(3), 271–280.