KV_C042
Villez, K., Habermacher, J. (2015). Shape Constrained Splines with Discontinuities for Anomaly Detection in a Batch Process. 12th International Symposium on Process Systems Engineering and 25th European Symposium on Computer Aided Process Engineering (PSE2015/ESCAPE25), Copenhagen, Denmark, May 31 - June 4, 2015, 1805-1810.

Krist V. Gernaey, Jakob K. Huusom and Rafiqul Gani (Eds.), 12th International Symposium on Process Systems Engineering and 25th European Symposium on Computer Aided Process Engineering.
31 May - 4 June 2015, Copenhagen, Denmark.

Shape Constrained Splines with Discontinuities for Anomaly Detection in a Batch Process

Kris Villez[a] and Jonathan Habermacher[a]

[a]*Department of Process Engineering, Eawag, Swiss Federal Institute of Aquatic Science and Technology; Dübendorf, Switzerland*
kris.villez@eawag.ch

Abstract

A previously developed technique for qualitative trend analysis (QTA) based on shape constrained splines (SCS) has been favourably compared to pre-existing techniques for the purpose of batch process diagnosis. Thanks to the branch-and-bound algorithm, this approach leads to a deterministic and global solution for QTA. One limitation of this method is that local discontinuities in otherwise continuous derivatives of the fitted spline function are not permitted. Recent work however allows to relax the optimization problem further so that provable bounds can be computed for this more complicated case. In this contribution, the resulting shape constrained splines with discontinuities (SCSD) method is applied for anomaly detection in batch process data. Importantly, the QTA approach to anomaly detection proves worthwhile because *(i)* tuning of the SCSD method is limited to setting an upper control limit, *(ii)* the resulting sum of squared errors statistic shows almost no drift for SCSD in contrast to the similar Q-statistic computed by principal component analysis (PCA), and *(iii)* true positive rates by means of SCSD are over 80% while the PCA method delivers at most 70% for false positive rates up to 5% based on a data set consisting of 410 batches.

Keywords: Anomaly detection, Fault diagnosis, Batch process monitoring, Statistical process control, Qualitative trend analysis

1. Introduction

Qualitative trend analysis (QTA) consists of a rather wide set of techniques for segmentation of data series. Such segmentation is popularly used for fault diagnosis applications. For such applications, QTA offers the possibility of an intuitive fault diagnosis mechanism. Its advantages include *(i)* that coarse-grained information about the faults symptomatic effects on measured data is sufficient to characterize and recognize a fault, *(ii)* that the required coarse-grained information corresponds to intuitive notions by human plant operators, and *(iii)* that consequently sparse occurrence of faults can effectively be dealt with. Despite these advantages, available techniques are often based on heuristic rules or greedy optimization schemes tuned for a particular application only. The lack of theoretical optimality makes it challenging to transfer such techniques from their original application context as was observed in Villez et al. (2012). This motivated the conception of a globally optimal approach to QTA based on shape constrained splines (SCS, Villez et al., 2013). This SCS method is limited in a number of ways. One important limitation is that any derivative of a degree smaller than or equal to $o-2$, in which o is the order of the spline function (for cubic splines: $o=4$), must be continuous over the whole function domain. Indeed, knots with multiplicity are not permitted. This is critical for QTA purposes as this means that certain shape constraints are not permitted. For example, one cannot fit a spline function which is characterized by a segment which is decreasing (negative 1st derivative) and concave (negative 2nd derivative) followed by a segment which is increasing (positive 1st derivative) and concave (negative 2nd derivative). Practically, this is solved by adding a knot of multiplicity $o-1$ when the exact location where the segments meet, a.k.a. transition, is known. This is however not the case in general. Unfortunately, the SCS method does not permit optimizing transitions which imply a discontinuity in one or more derivatives which are otherwise continuous. This paper discusses results obtained with a modified version of the SCS method which accounts for shape-implied discontinuities. The method is therefore further referred to as the shape constrained splines with discontinuities (SCSD) method. Similar to the SCS method, it is based on a combination of second order cone programming (SOCP) and the branch-and-bound algorithm. This algorithm enables the identification of all transitions. The method is successfully demonstrated by means of oxidation-reduction potential (ORP, redox) measurements obtained from a laboratory scale sequencing batch process for wastewater treatment.

2. Methods

2.1. Case study

The data used in this study stems from an experimental side-stream reactor (SSR) setup operated at Eawag for a period of two years by the second author of this contribution. This SSR setup consists of two connected tanks in which the first tank is a sequencing batch reactor (SBR). The second tank is the so called SSR and is operated as a continuously stirred tank reactor for sludge digestion. The SBR is operated with a fixed recipe which lasts 6 hours. Each SBR cycle starts with pumping of sludge from the SSR to the SBR (7 min.) followed by addition of fresh wastewater (10 min.) under anoxic conditions. A reaction phase follows in which air is supplied to provide aerobic conditions in the SBR (285 min.). The SBR cycle ends with excess sludge withdrawal, settling, and decanting (58 min.). Each new cycle starts right after decanting.

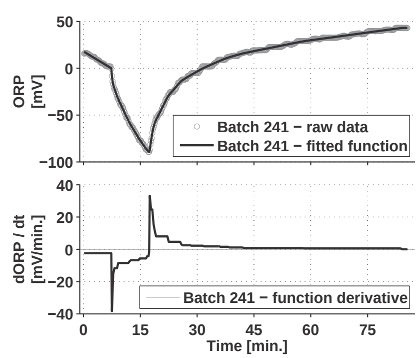

Figure 1: Exemplary time series corresponding to normal operation.

The SBR is equipped with an oxidation-reduction potential (ORP) sensor, amongst others. A typical profile of the first 513 data points (85.33 min.) is shown in Fig. 1 (top panel). The first phase corresponds to an almost linear decreasing trend, followed by a convex decreasing trend in the second phase. In the third phase, the trend roughly corresponds to a concave increasing trend. The ORP time series consisting of the first 513 data points of 410 consecutive cycles are studied in this work. The selected batch cycles represent a period of 3.5 months in which no changes were made to the SBR recipe, the installed sensors (e.g., no replacements), or the data acquisition and control system. Each of these time series were inspected visually by two operators which are familiar with the experimental setup and the use of ORP sensors, including the second author of this paper. These operators addressed the question *whether the observed time series can be explained by normal circumstances alone* by means of a simple yes or no. To this end, the time series were visually represented to these operators on separate occasions and in a randomized order. Following this initial classification of the ORP time series as normal or anomalous, batches for which the classification by the two operators was different were visualized again with both operators present so to find a consensus classification result. For 16 cycles out of 410, no consensus could be reached. Only the time series for which consensus was reached (394 batch cycles) are considered in the remainder of this study. The consensus classification is used as the ground truth or true class for each analysed time series.

2.2. Qualitative trend analysis (QTA): Definitions

The identification of so called *qualitative representations* (QRs) is the primary goal of QTA. Such QRs are defined as a sequence of *episodes*. Each of these episodes is completely characterized by means of *(i)* a start time, *(ii)* an end time, and *(iii)* a *primitive*. Such a primitive corresponds a unique combination for the signs of the measured variable and/or one or more of its derivatives with respect to time. A primitive is usually represented by means of a unique character, which can however be chosen arbitrarily. Fig. 2 shows the primitives and characters used in this study. Only the signs of the first and second derivatives are considered relevant. The episodes within a QR are contiguous, meaning that the end time of one episode is the start time of the next episode. The locations (in time) where two episodes meet are known as *transitions*. A sequence of primitives without specification of the transitions is known as a qualitative sequence (QS). Conform these definitions, the QS for a typical profile of ORP measurements as in Fig. 1 reads as EAC. The SCSD method explained in the next paragraphs allows to identify the values for the corresponding transitions.

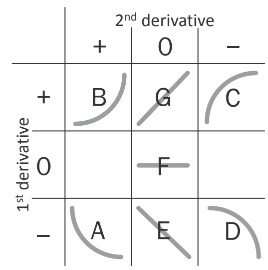

Figure 2: Primitives and associated characters.

2.3. Shape constrained splines (SCS)

In essence, the fitting of a shape constrained spline function can be written as the following mathematical problem:

$$\min_{\beta,\theta} J(\beta,\theta) = J(\beta,\theta,\mathbf{x},\mathbf{y}) \tag{1}$$

subject to:

$$\beta \in \Omega(\theta) \tag{2}$$

$$\theta \in \Theta \tag{3}$$

Table 1: Definitions of variables and parameters for SCSD.

Symbol	Definition
β	Function parameters
θ	Transitions
$\Omega(\theta)$	Feasible set for β
Θ	Feasible set for θ
$\mathbf{x}$	Argument vector
$\mathbf{y}$	Measurement vector
J	Objective function

with definitions given in Table 1. The spline function is defined by a number, m, of parameters represented by the vector β. In Villez et al. (2013), the parameters are the spline coefficients. Given a QR defining the shape constraints, the feasible set for these coefficients, $\Omega(\theta)$, is a convex subset of the real space ($\Omega(\theta) \subseteq \mathbb{R}$). This set depends on the values for the transitions, θ. In the context given by Villez et al. (2013) and this study, optimal values for θ are found by means of the branch-and-bound algorithm. Proofs for the applied bounds were given in Villez et al. (2013). The existing proofs are limited however in the sense that the bounding procedures only apply to spline functions without knot multiplicity and any associated discontinuities of the derivatives in such knots. This limitation is removed in the next paragraphs.

2.4. Shape constrained splines with discontinuities (SCSD)

To enable application of the branch-and-bound algorithm when discontinuities are implied by the imposed QS, the fitted function is now redefined by representing the spline function explicitly as a contiguous set of piece-wise polynomials. Once more, the problem is convex as long as θ

is given. To solve the non-linear search for θ the branch-and-bound algorithm can be applied again. The upper bound procedure is practically the same as in Villez et al. (2013). The lower bound procedure is fairly different however. In Villez et al. (2013) it was sufficient to implement only those constraints on β which are guaranteed to be included for any candidate solution for θ within a given solution set. In this case, one further needs to modify the objective function as well as set of implemented continuity constraints. The details of the lower and upper bounding procedures are given in Villez (2014). Apart from the modified bounding procedures, the branch-and-bound algorithm is applied exactly as in Villez et al. (2013). The branch-and-bound algorithm was solved until a resolution of 1 sampling interval (10 seconds) was reached for the location of the transitions. The fitted function is a natural cubic spline function ($o = 4$) with fixed spline knots placed at each sampling time. The resulting function is an interpolating spline function when no shape constraints are applied. As a result, any deviation of the SSE from zero is only due to a mismatch between the measured time series and the imposed shape (EAC).

2.5. *Shape constrained function fitting as a method for process monitoring*

To test the shape constrained splines methodology for process monitoring, the following strategy is applied. One optimizes the values for β and θ so to minimize the objective function value J. Importantly, there is no calibration phase for the SCSD method as long as the expected shape is known. In this study, this objective function is the sum of squared errors (SSE). In general, the optimized SSE, further referred to as SSE_{SCS}, will be high when the profile of the analysed data series does not correspond well to the imposed shape on the fitted function. As such, this SSE can be used as a lack-of-fit statistic to detect deviations from the expected shape. To this end, one defines a fixed upper limit, U_{SCS}, a priori and one classifies a data series as anomalous (normal) when the obtained SSE is higher (lower) than this limit. There are no theoretical distributions available for the SSE_{SCS} statistic. As such, an empirical evaluation is obtained by evaluating the false positive rate (FPR, fraction of normal cycles which are erroneously classified as anomalous) as well as the true positive rate (TPR, fraction of abnormal cycles which are correctly classified as anomalous) by varying the value for U_{SCS} between the lowest and highest value for SSE_{SCS} obtained. To this end, the ground truth classification as described above is used as a reference. This leads to the so called receiver operating characteristic (ROC, Fawcett, 2006).

2.6. *Principal component analysis (PCA)*

Principal component analysis (PCA) is by far the most commonly used method for process monitoring in a multivariate context. PCA is of use primarily when significantly large data sets are available and first principles modelling is impossible or cost-prohibitive. The successful application of classical PCA depends on a number of factors such as *(i)* whether relationships between measured variables are linear and time-invariant, and *(ii)* whether sufficient and representative data is available for PCA modelling of the normal process data variation. Despite these restrictive statistical requirements, PCA has been applied successfully to many processes. In this study, the batch process data is organized in a matrix with rows corresponding individual batches and columns to individual time instants in each of the analysed batches. A PCA model is then fitted to a subset of the data by *(i)* selecting a number of rows corresponding to normal process conditions, *(ii)* subtracting the column-wise mean, and *(iii)* computing a number of principal component loadings and associated scores by means of singular value decomposition. In this study, a simple scree plot is used as the device to select the number of principal components (PCs). Once the PCA model is identified, one projects the data not used for calibration and evaluation the so called reconstruction errors or residuals. Following this step, one computes the corresponding SSE, SSE_{PCA}, which is also known as the Squared Prediction Error (SPE, Kresta et al., 1991) and as the Q statistic (Jackson and Mudholkar, 1979). Theoretical and/or empirical confidence limits can be computed for this statistic (Jackson and Mudholkar, 1979). In this study, the ROC is computed as is done for the SCSD method to enable an effective comparison.

3. Results

3.1. Example

Fig. 1 displays the exemplary oxidation-reduction potential (ORP) time series, the optimized shape constrained spline function with EAC shape, and the 1^{st} derivative of this fitted function. Finding the optimal transitions required 34 branching steps in the branch-and-bound algorithm which is equivalent to solving 69 SOCPs. This took 1.2 minutes on a dedicated desktop machine (Intel® Core™ i7-4770 CPU, 3.40 GHz, 16.0 GB RAM). The optimized SSE_{SCS} is equal to 332.

3.2. Benchmarking

All batch time series were analysed with the SCSD method on a dedicated desktop machine (as above). Computing the optimal transitions took less than 20 minutes in all cases. The PCA model was identified by means of the first 100 normal batch cycles and included two principal components which capture 86.9% and 9.6% of the variability (total: 96.5%). The statistics obtained with SCSD and PCA are shown in Fig. 3 as a function of batch index. The SSE_{SCS} is relatively stable with a number of outlying values sparsely present during the covered time span. In contrast, the SSE_{PCA} shows more variation and is generally higher. This is believed to be a result of non-linear parametric changes of the ORP profiles which are hard to capture by a linear PCA model. The SCSD method is robust to these variations because of a greater parametric flexibility of the spline function compared to PCA. In the top panel, an upper detection limit, U_{SCS}, is set at 610.33 which is the lowest value for which no false alarms (no false positives) result. Similarly, in the bottom panel, the upper detection limit, U_{PCA}, is set at 37659. With SCSD 38 of the 46 abnormal time series have an SSE_{SCS} which is higher than the set limit (TPR: 83%). With PCA only 30 positive detections result (TPR: 65%). Closer inspection reveals that 29 cycles are identified as abnormal by both methods. The SCSD method correctly flags 8 batch cycles as anomalous which are not detected by PCA. One batch cycle not identified with SCSD is identified as anomalous by PCA.

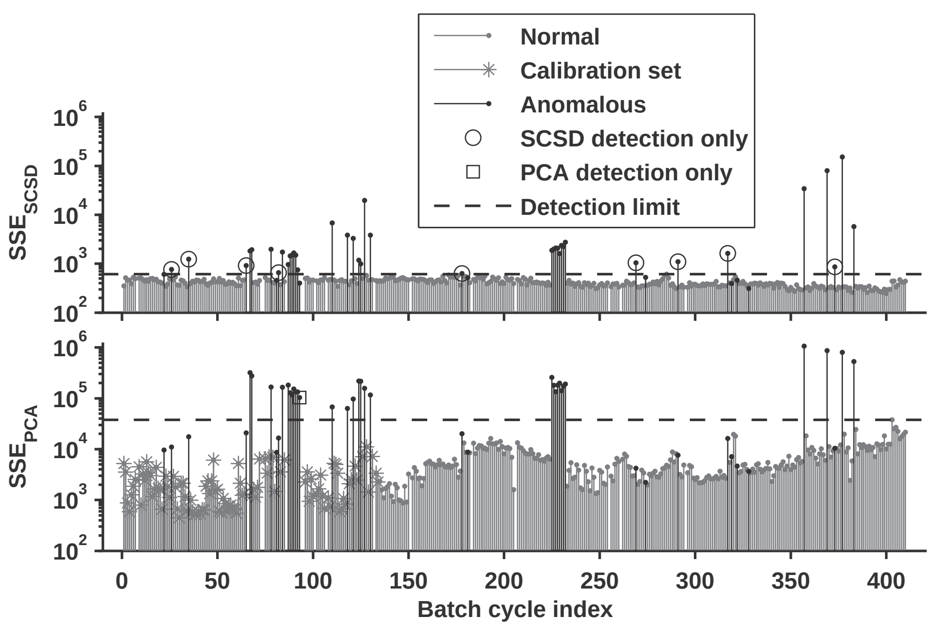

Figure 3: Sum of squared errors (SSE) statistics obtained with SCSD (top) and PCA (bottom).

Decreasing values of the upper limits for the SSR statistics increase both the FPR and the TPR. For example, at an FPR of 5% the TPR for SCSD is 85% while the TPR for PCA is 70%. The complete receiver operating characteristics (ROCs) are shown in Fig. 4. One can observe that the SCSD method performs better for FPRs up to 69%. PCA dominates only for FPRs ranging from 88% to 93%. In other sections of the curve, the methods are performing equally. In the present application as well as for most anomaly detection schemes, normal operating conditions are more frequent than anomalous conditions. It is thus reasonable to give larger weight to a low FPR than a high TPR. Selecting FPRs lower or equal than 5% is common. At such rates, the SCSD method clearly outperforms PCA-based detection.

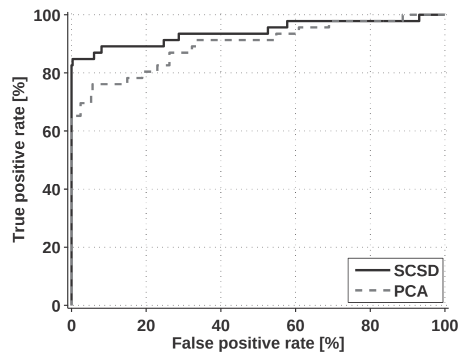

Figure 4: Receiver operating characteristics (ROC) showing the true positive rate (TPR) as a function of the false positive rate (FPR).

4. Conclusion

With this contribution, an extended version of the shape constrained spline (SCS) method for qualitative trend analysis (QTA) has been tested for batch process monitoring. The proposed extension allows to optimize the location of abrupt changes (discontinuities) in one or more derivatives of measured variable trends. Importantly, the identification of shape constrained splines with discontinuities (SCSD) is based on the optimization of the sum of squared errors which is has been successfully used as a lack-of-fit statistic for anomaly detection. A preliminary comparison with principal component analysis underscores the value of a qualitative approach to fault detection.

5. Acknowledgement

The authors want to thank Nicolas Derlon at Eawag for interpreting and classifying the studied data series and for his creative input and constructive feedback during execution and reporting of this study. All computations were executed by joint use of Matlab (The MathWorks Inc., 2012); the Functional Data Analysis software package accompanying Ramsay and Silverman (2002); and MOSEK (MOSEK ApS, 2012). The first author thanks O. Peth for everlasting inspiration.

References

Fawcett, T., 2006. An introduction to ROC analysis. Pattern recognition letters 27, 861–874.

Jackson, J., Mudholkar, G., 1979. Control procedures for residuals associated with principal component analysis. Technometrics 21, 341–349.

Kresta, J., MacGregor, J., Marlin, T., 1991. Multivariate statistical monitoring of process operating performance. Can. J. Chem. Eng. 69, 35–47.

MOSEK ApS, 2012. MOSEK Optimization Software for MATLAB, Version 6.0. Available from http://www.mosek.com/.

Ramsay, J. O., Silverman, B. W., 2002. Applied Functional Data Analysis: Methods and Case Studies. Springer-Verlag.

The MathWorks Inc., 2012. Matlab, version 8.0.0.783 (R2012b). Natick, Massachusetts.

Villez, K., 2014. Bounding Procedures for Shape Constrained Splines with Discontinuities. Eawag, Dübendorf, Switzerland, Technical Report TR-001-01-0.

Villez, K., Rengaswamy, R., Venkatasubramanian, V., 2013. Generalized shape constrained spline fitting for qualitative analysis of trends. Comp. Chem. Eng. 58, 116–134.

Villez, K., Rosén, C., Anctil, F., Duchesne, C., Vanrolleghem, P. A., 2012. Qualitative representation of trends (QRT): Extended method for identification of consecutive inflection points. Comp. Chem. Eng. 48, 187–199.