

Process Fault Detection, Isolation, and Reconstruction by Principal Component Pursuit

Joshua D. Isom and Robert E. LaBarre

Abstract—A common approach to process monitoring based on principal component analysis (PCA) assumes that fault-free, noise-free data is sampled from a low-dimensional subspace. Although widely described and applied, process fault detection and isolation using PCA is not robust to outliers in the training data, is hard to properly tune, and is not capable of isolating multiple faults. A newly introduced method called principal component pursuit (PCP) optimally decomposes a data matrix as the sum of a low-rank matrix and a sparse matrix. When applied to the process monitoring problem, PCP simultaneously accomplishes the objectives of model building, fault detection, fault isolation, and process reconstruction with a single convex optimization problem, thereby overcoming the key shortcomings of PCA-based approaches for process monitoring. The use of PCP for process monitoring is described and illustrated using data from a manufacturing process.

I. INTRODUCTION

Data-driven approaches for process monitoring are widely applied because they are developed, deployed, and maintained at low cost. The principal component analysis (PCA) approach to process monitoring assumes that process data, in the absence of noise or faults, is sampled from a low-dimensional subspace. This paper compares existing PCA-based approaches for process monitoring to a technique based on a new matrix decomposition technique called principal component pursuit (PCP), which decomposes a matrix as the sum of a low-rank matrix and a sparse matrix. Like PCA-based approaches, the PCP approach to process monitoring assumes that process data, in the absence of noise or faults, is sampled from a low-dimensional subspace. However, the PCP approach remedies the key shortcomings of PCA-based approaches.

Section II provides a description of the PCA-based approach to process fault detection, isolation, and reconstruction. Section III describes the shortcomings of the PCA-based approach. Section IV develops a PCP-based approach for process fault detection, isolation, and reconstruction. Section V provides a comparison of the PCP-based approach to the PCA-approach, highlighting connections between the two methods and describing ways in which PCP solves longstanding problems. Section VI provides experimental results on the application of PCA- and PCP-based methods for process monitoring.

This work was supported by UTC Power, with technical assistance from Mithun Kamat and Ryan Boersma.

Joshua D. Isom and Robert E. LaBarre are with United Technologies Research Center, East Hartford, CT 06108, USA
isomjd@utrc.utc.com, labarrre@utrc.utc.com

II. PROCESS FAULT DETECTION, ISOLATION, AND RECONSTRUCTION USING PCA

Principal component analysis is widely used as the basis for data-driven fault detection and isolation of industrial processes [1]. Given a training data matrix X containing m rows of observations of n columns of process variables which have been autoscaled so that each process variable has zero mean and unit variance, a PCA model is constructed using the singular value decomposition $\frac{1}{\sqrt{m-1}}X = U\Sigma V^T$ where the matrix Σ contains the non-negative real singular values of decreasing magnitude along its main diagonal ($\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_{\min(m,n)} \geq 0$), and zero off-diagonal elements. The PCA loading vectors are the orthonormal column vectors in the matrix V . One selects the columns of the loading matrix P to correspond to the loading vectors associated with the first a singular values, and the projections of the observations in X into the lower-dimensional space are contained in the score matrix, $T = XP$. A low-rank reconstruction of the process data is $\hat{X} = XPP^T$. The PCA model-building process is equivalent to solving for a low-rank reconstruction \hat{X} of the data X by solving the following optimization problem:

$$\text{minimize } \|X - \hat{X}\|_2 \text{ subject to } \text{rank}(\hat{X}) \leq a. \quad (1)$$

The T^2 statistic for PCA-based fault detection is a scaled, squared distance from the mean, and is defined as $T^2 = \mathbf{x}^T P \Sigma_a^{-2} P^T \mathbf{x}$ where Σ_a contains the first a rows and columns of Σ . An alternative statistic for fault detection is the *squared prediction error* $Q = \mathbf{x}^T \tilde{P} \tilde{P}^T \mathbf{x}$ where $\tilde{P} = I - PP^T$ is the principal component residual space. Both fault detection statistics have the quadratic form $\mathbf{x}^T M \mathbf{x}$ with $M = P \Sigma_a^{-2} P^T$ for the T^2 statistic and $M = \tilde{P} \tilde{P}^T$ for the Q statistic.

For testing data, a fault is detected if the T^2 or Q statistic exceeds a threshold. Once a fault has been detected, fault isolation is performed. Although there a number of techniques for PCA-based fault isolation, the approach with the best theoretical grounding calculates the reconstruction-based contribution [2] of each process variable to the fault detection statistic.

The reconstructed vector in the direction \mathbf{e}_i of process variable i is

$$\mathbf{y}_i = \mathbf{x} - \mathbf{e}_i f_i \quad (2)$$

where f_i is the magnitude of the fault. The value of the fault detection statistic for the reconstructed data vector is $\mathbf{y}^T M \mathbf{y}$. The task of reconstruction is to find a value of the

fault magnitude f_i such that the fault detection statistic for the reconstructed process condition is minimized. The value of f_i meeting this objective is [2]

$$f_i = (\mathbf{e}_i^T M \mathbf{e}_i)^{-1} \mathbf{e}_i^T M \mathbf{x}. \quad (3)$$

The reconstruction-based contribution (RBC) of variable x_i to the fault detection statistic (Q or T^2) is $RBC_i = \|\mathbf{e}_i f_i\|_M^2$, and the isolated fault direction is that which has the maximal contribution. Once a fault has been detected and isolated, process reconstruction to remedy the effect of the fault can be accomplished with (2).

III. SHORTCOMINGS OF PCA-BASED APPROACHES

In the model building, fault detection, fault isolation, and reconstruction steps, PCA has a number of shortcomings which increase the complexity and limit the applicability of the approach. These limitations are discussed in the following subsections.

A. Model building

A first shortcoming of PCA-based approaches is that the calculated principal components are sensitive to the presence of outliers in the data set. It follows from Lemma 1 in the Appendix that a single anomalous row in a data matrix X can arbitrarily increase the singular values and rotate the principal components of an otherwise low-rank training data set. Thus it becomes critical to remove outliers from the training data set prior to the construction of the principal component model. A variety of techniques for robust outlier have been proposed [3], but the selection of one of these techniques and its application complicates the implementation of PCA for fault detection and isolation. More important, prior to the introduction of PCP, “none of the existing approaches yields a polynomial-time algorithm with strong performance guarantees under broad conditions” [4].

Another shortcoming of PCA-based methods is that one must decide on the number of principal components to be retained, as the fault detection and isolation results are sensitive to the reduction order. As noted in the literature, several techniques have been proposed but there appears to be no dominant technique [1].

B. Fault detection

A weakness of the PCA-based approach to fault detection is that the T^2 and Q statistics are known to be sensitive to different types of faults [1], the use of one without the other may mean that certain types of faults are not detected and isolated.

C. Fault isolation

The PCA-based approach to fault isolation with the best theoretical grounding consists of finding the magnitude of a fault in a specific fault direction that minimizes the value of the fault detection statistic:

$$\text{minimize}_{f \in R, \mathbf{e} \in E} (\mathbf{x} - f\mathbf{e})^T M (\mathbf{x} - f\mathbf{e}) \quad (4)$$

where \mathbf{e} is a nonnegative unitary vector corresponding to a fault direction, f is a real number representing the magnitude

of the fault, and E is the set of all fault directions considered. This approach is termed *reconstruction-based contribution* because $\mathbf{x} - f\mathbf{e}$ is the best reconstruction of the process data over a finite set of fault directions E if

$$(f, \mathbf{e}) = \arg \min_{f \in R, \mathbf{e} \in E} (\mathbf{x} - f\mathbf{e})^T M (\mathbf{x} - f\mathbf{e}). \quad (5)$$

Explicit solution of (4) is not viable for the unsupervised isolation of multiple faults (fault directions of cardinality greater than one). Even if the fault direction vectors are constrained such that each positive element is equal, the number of such fault directions for a given fault cardinality grows combinatorially with the number of process variables. Specifically, if there are n process variables and one limits E to unitary vectors of cardinality c with equal positive elements, the size of $|E|$ is $|E| = C(n, c) = \frac{n!}{c!(n-c)!}$. Furthermore, as will be described below, the PCA-based approach to fault detection and isolation becomes ill-posed in the presence of faults of large cardinality.

Because of the above limitations associated with PCA-based fault isolation, authors have suggested augmenting PCA-based fault isolation with expert knowledge of the process, such as that captured in a sign-directed graph [5] or in a structured residual model [6], in order to successfully isolate multiple faults.

D. Process reconstruction

Process reconstruction is the estimation of the values of process variables known to be affected by a fault in the scenario where the fault is corrected. Reconstructed process variables can be used for data reconciliation and for fault-tolerant control. The main limitation of PCA-based process reconstruction is a consequence of the main limitation PCA-based fault isolation: the unsupervised method is not capable of the isolation of multiple faults, and therefore, process reconstruction cannot remedy the effects of multiple faults.

IV. PROCESS FAULT MONITORING USING PRINCIPAL COMPONENT PURSUIT

This section describes principal component pursuit and develops techniques for fault detection, isolation, and reconstruction using PCP.

Given a data matrix X , principal component pursuit (PCP) is the solution of the convex optimization problem

$$\begin{aligned} &\text{minimize } \|Y\|_* + \lambda \|Z\|_1 \\ &\text{subject to } X = Y + Z \end{aligned} \quad (6)$$

with $\|A\|_*$, the nuclear norm, equal to the sum of the singular values of A and with $\|A\|_1$ equal to the sum of the absolute values of the elements of A [4]. Under certain conditions on a low-rank matrix Y_0 , a sparse matrix Z_0 , and a Lagrange multiplier λ , the optimization problem (6) recovers Y_0 and Z_0 exactly given the input $X = Y_0 + Z_0$ [4], [7]. The optimization problem is convex and linearly constrained, and a number of efficient algorithms are available [4].

Consider the process model: $X = Y + G + H$ where Y is a low-rank matrix corresponding to a fault-free process condition, G has nonzero entries corresponding to sensor

noise, and H has nonzero entries corresponding to sensor and process faults, and each matrix is $m \times n$. In what follows, $d_{\min} = \min(m, n)$ and $d_{\max} = \max(m, n)$.

For the purpose of process fault detection, isolation, and reconstruction, the following assumptions on Y , G , and H are made:

- (A) The matrix Y is μ -incoherent. If $Y = U\Sigma V^*$ is the reduced singular value decomposition, and r is the rank of Y , then Y is μ -incoherent if $\max_i \|U^* \mathbf{e}_i\|^2 \leq \mu r/m$, $\max_i \|V^* \mathbf{e}_i\|^2 \leq \mu r/n$, $\|UV^*\|_\infty \leq \sqrt{\mu r/mn}$ [8].
- (B) With probability $\rho < 1$, any entry g_{ij} of G is a random variable with a cumulative distribution function $F_j^g(\cdot)$ having a median of zero; with probability $1 - \rho$, $g_{ij} = 0$.
- (C) With probability $\pi < 1$, any entry h_{ij} of H is a random variable with a cumulative distribution function $F_j^h(\cdot)$ having a median of zero; with probability $1 - \pi$, $h_{ij} = 0$.

Let $\eta = 1 - (1 - \rho)(1 - \pi)$ be the probability that any particular entry of $Z = G + H$ is nonzero.

A qualitative statement of assumption (A) when μ is small is that the principal axes of Y are not closely aligned with the standard basis. This assumption is reasonable in situations where PCA would be applied because it is a key motivating assumption for the use of PCA for process monitoring [1].

One of the implications of assumptions (B) and (C) is that noise and faults of negative signs and positive signs occur with equal probability. Another implication is that sensor noise, sensor faults, and the fault-free process are independent random variables. Both implications are reasonable assumptions and would likely be default assumptions in the absence of specific contrary knowledge. Examples of probability distributions with a median of zero include a zero-mean Gaussian distribution, a zero-mean uniform distribution, and a binomial distribution over a negative and positive alternatives with parameter $\rho = 0.5$.

The following theorem establishes the optimality of PCP in recovering low-rank process data from noise and faults.

Theorem 1: If $X = Y_0 + G_0 + H_0$ and assumptions (A),(B), and (C) hold, then there exists a value of λ for which principle component pursuit (6) exactly recovers the matrices Y_0 and $Z_0 = G_0 + H_0$ with high probability if the rank r of Y_0 satisfies $r < \frac{C_1 d_{\min}}{\mu \log^2 d_{\max}}$ for some nonzero numerical constant C_1 .

Proof: Assumption (A) is the same as [8, Theorem 1]. If assumptions (B) and (C) are satisfied, then the elements of $Z = G + H$ take positive and negative signs with equal probability, and assumption (B) of [8, Theorem 1] is satisfied. The result follows from [8, Theorem 1], generalized to non-square matrices following [4]. ■

Given the optimality of PCP in recovering Y_0 and Z_0 from $X = Y_0 + Z_0$, and assumptions (B) and (C), the PCP-based approach for fault detection, isolation, and reconstruction is simple. One uses PCP to decompose the autoscaled process data X as $X = Y + Z$. A fault is detected and isolated in variable j in measurement i if $z_{ij} \neq 0$ and one rejects the null hypothesis that z_{ij} is sampled from the noise distribution $F_j^g(\cdot)$ in favor of the alternative hypothesis that h_{ij} , the entry of the fault matrix, is nonzero. The mechanism for separation

of noise from faults in Z takes the form of a threshold on the absolute value $|z_{ij}|$ if the noise and fault distributions are both Gaussian and the fault distribution has a larger variance. If a fault is detected and isolated, an estimate of the fault magnitude is z_{ij} and a reconstructed estimate of the fault-free, noise-free variable is y_{ij} .

V. ADVANTAGES OF PCP RELATIVE TO PCA

This section describes the advantages of a PCP-based approach to fault detection, isolation, and reconstruction relative to a PCA-based approach. Relationships between PCA process monitoring and an online version of PCP process monitoring are described.

A. Model building

The singular value decomposition used to construct a PCA model is not robust to outliers in the training data set. In contrast, the PCP is provably robust at recovering a low-rank reconstruction in the presence of outliers in the data set, as stated in Theorem 1. This is a key advantage of PCP relative to PCA.

Another advantage of PCP approach relative to PCA is that PCP does not require the selection of a model order. To apply PCP to a data matrix X and achieve an optimal low-rank reconstruction Y_0 , the rank of Y_0 need not be known or assumed *a priori*, it must only be low relative to the dimensions of X .

The only tuning parameter for PCP model building is the constant λ in (6). The optimal value for this parameter is known to be $O\left(\frac{1}{\sqrt{n}}\right)$ [4], and tighter bounds on the optimal value of λ are available if η is known [8].

B. Fault detection

The fault detection and isolation decision for PCP amounts to determining whether a deviation from zero can be explained well by the properties of noise of a particular sensor, whereas fault detection for the PCA-based method is based on a multivariate statistic. The process for establishing fault detection thresholds for PCP is a univariate problem where physical intuition is more readily applied.

A second advantage of PCP relative to PCA for fault detection, described below, is that PCP is sensitive to both faults that would elevate the PCA T^2 statistic and to faults that would elevate the PCA Q statistic.

Consider now the relationship between fault detection results produced by the PCA method and the PCP method. These relationships are illustrated with the following scenario. A training data set X is available. PCP has been used to perform the decomposition $X = Y + Z$, and PCA has been used to develop a singular values and loading vectors. A new measurement \mathbf{x} is made, and PCP is applied to find the optimal decomposition $\begin{bmatrix} X \\ \mathbf{x} \end{bmatrix} = Y' + Z'$.

At one extreme, all of \mathbf{x} can be assigned to low-rank Y' and none to sparse Z' . The singular value decomposition of Y , $Y = U\Sigma V^T$, where the matrix Σ contains the non-negative real singular values of Y along its main diagonal

($\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_{\min(m,n)} \geq 0$), and zero off-diagonal elements, will be used to analyze the nuclear norm of Y' .

If the matrix Y is augmented with the testing vector \mathbf{x} such that the augmented matrix Y' is $Y' = \begin{bmatrix} Y \\ \mathbf{x} \end{bmatrix}$, it is shown in Lemma 1 in the Appendix that the singular values of the matrix Y' correspond to the singular values of the matrix

$$\begin{bmatrix} \Sigma & 0 \\ \mathbf{t} & \sqrt{Q} \end{bmatrix} \quad (7)$$

where $\mathbf{t} = [t_1 \dots t_n]$ is a row vector of scores defined by $\mathbf{t} = V^T \mathbf{x}$ and Q is the PCA fault detection statistic, $\sqrt{Q} = \|(I - VV^T)\mathbf{x}\|$.

The singular values τ_j of the matrix (7) (which are the singular values of Y') satisfy the secular equation

$$0 = Q - \tau_j^2 - \sum_{i=1}^{n-1} \frac{t_i^4}{\sigma_i^2 - \tau_j^2}. \quad (8)$$

Consider two limiting cases on the nuclear norm of Y' , when a new observation is appended to Y . The first limiting case is when \mathbf{x} is orthogonal to a basis for Y . In this case, all the scores t_i are zero and Q is nonzero and it can be verified from (8) that the increase in the nuclear norm $\|Y'\|_*$ is simply $\sqrt{Q + \sum \sigma_i}$.

Another limiting case is when \mathbf{x} is in the direction of a principal component of Y . In this case, Q is zero and there is a single nonzero score t_i . In this case, it can be verified from (8) that the nuclear norm is given by $\|Y'\|_* = \sqrt{\sigma_i^2 + t_i^2}$.

In the general case, it is shown in Lemma 2 in the appendix that bounds on $\|Y'\|_*$ are

$$\sqrt{Q + \sum_i t_i^2 + \sigma_i^2} \leq \|Y'\|_* \leq \sqrt{(n+1)(Q + \sum_i t_i^2 + \sigma_i^2)}. \quad (9)$$

It is evident that $\|Y'\|_*$ is more closely related to the square root of the PCA-based statistics Q and T^2 than to the unmodified statistics. The following asymptotic relationships exist:

$$\|Y'\|_* - \|Y\|_* \rightarrow 0 \text{ as } Q \rightarrow 0, T^2 \rightarrow 0 \quad (10)$$

$$\|Y'\|_* - \|Y\|_* \rightarrow \sqrt{Q} \text{ as } \sum_i t_i^2 \rightarrow 0$$

$$\|Y'\|_* - \|Y\|_* \rightarrow \sigma_i \sqrt{T^2} \text{ as } t_i \rightarrow \infty, \frac{Q + \sum_{j \neq i} t_j^2}{t_i} \rightarrow 0.$$

C. Fault isolation

This section describes the relationship between PCA-based fault isolation and PCP-based fault isolation and explains why, unlike PCA, PCP is optimal for isolating multiple faults.

The PCP approach for fault isolation can be developed as a refinement of the PCA-based approach to fault isolation. A generalization of the PCA-approach to fault isolation (4) which allows for arbitrary fault directions is

$$\begin{aligned} & \text{minimize}_{\mathbf{y}, \mathbf{z}} \quad \mathbf{y}^T M \mathbf{y} \\ & \text{subject to } \mathbf{x} = \mathbf{y} + \mathbf{z}, \end{aligned} \quad (11)$$

where $M = P\Sigma_a^{-2}P^T$ (for the T^2 statistic) or $M = \tilde{P}\tilde{P}^T$ (for the Q statistic), \mathbf{y} is the reconstructed process data, and \mathbf{z} is an arbitrary fault vector. This problem is ill-posed, because it has a trivial solution $\mathbf{y} = \mathbf{0}$. A well-posed PCA-based version of the problem of isolation of multiple faults incorporates a constraint on the cardinality of the fault \mathbf{x} is

$$\begin{aligned} & \text{minimize}_{\mathbf{y}, \mathbf{z}} \quad \mathbf{y}^T M \mathbf{y} \\ & \text{subject to } \mathbf{x} = \mathbf{y} + \mathbf{z}, \|\mathbf{z}\|_0 \leq c \end{aligned} \quad (12)$$

where \mathbf{y} is the reconstructed process data and \mathbf{z} is a fault vector, and $\|\mathbf{z}\|_0$ is the cardinality of the fault direction. The standard PCA approach to fault isolation described in Section III solves this problem with $c = 1$.

The 1-norm is often used as a convex heuristic for the 0-norm. Making this substitution of norm and putting the cardinality constraint in Lagrangian form, one has

$$\begin{aligned} & \text{minimize}_{\mathbf{y}, \mathbf{z}} \quad \mathbf{y}^T M \mathbf{y} + \lambda \|\mathbf{z}\|_1 \\ & \text{subject to } \mathbf{x} = \mathbf{y} + \mathbf{z}. \end{aligned} \quad (13)$$

This formulation of PCA-based fault-detection, isolation, and reconstruction is a tractable convex optimization problem which allows for the isolation faults of cardinality greater than one. If one accepts the optimality of PCP, then it seems unlikely that the formulation (13) has optimal properties, due to the fact that, as shown in the previous section, $\mathbf{y}^T M \mathbf{y}$ is more closely related to the square of the nuclear norm than to the nuclear norm of the data matrix.

If the PCA-based fault detection term $\mathbf{y}^T M \mathbf{y}$ is replaced with the nuclear norm, one arrives at an online formulation of Principal Component Pursuit:

$$\begin{aligned} & \text{minimize}_{\mathbf{y}, \mathbf{z}} \quad \left\| \begin{bmatrix} Y \\ \mathbf{x} \end{bmatrix} \right\|_* + \lambda \|\mathbf{z}\|_1 \\ & \text{subject to } \mathbf{x} = \mathbf{y} + \mathbf{z}, \end{aligned} \quad (14)$$

where Y is a set of low-rank, fault-free data produced from the training problem (6). To compute this online version, one need not store the entirety of Y , but only its singular values Σ and loading vectors V , as shown in Lemma 1 in the Appendix.

D. Process reconstruction

The advantage of PCP relative to PCA for process reconstruction is a result of the fact that PCP is capable of isolating multiple faults, and is therefore capable of reconstructing the effects of multiple faults. Unlike PCA, PCP provably recovers low-rank, fault-free process data in the presence of multiple faults.

VI. EXPERIMENT AND RESULTS

This section describes the application of PCA- and PCP-based techniques for process fault detection, isolation, and reconstruction to process data collected during the quality assurance testing of fuel cell power plants.

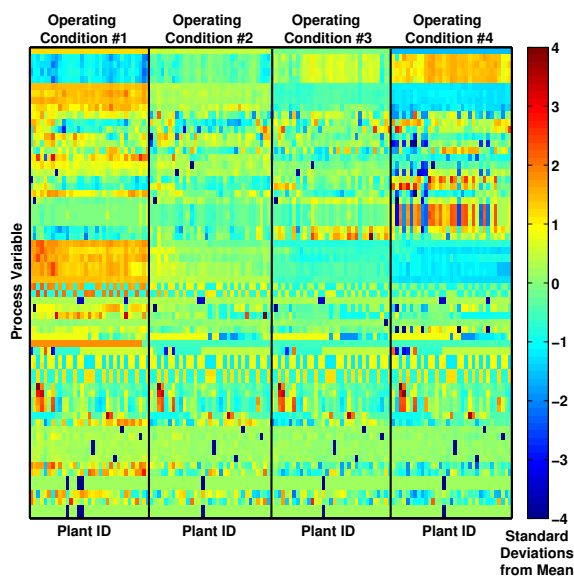


Fig. 1. Autoscaled process data from fuel cell power plants.

A. Data Description

Data was collected from a manufacturing process for fuel cell power plants. One of the final steps of the manufacture is an operational check of each power plant in which the power plant is operated at steady-state at four operating points as a quality check. Each of the 33 power plants in the data set has 66 measured variables. The data matrix X consists of 132 columns (four operating points for each of the 33 power plants) and 66 rows, one for each measurement variable. In general, process faults were not corrected until after each power plant had been operated at each of the four operating points.

B. Experimental Set-Up

Prior to application of the process monitoring techniques, the data was autoscaled so that each measured variable had zero mean and unit variance over the entire data set. The autoscaled process data is depicted in Figure 1.

Five different algorithms for fault detection, isolation, and reconstruction were tested: 1) the PCA-based technique with the T^2 statistic, applied without removing outliers from the original data set (“PCA T^2 ”); 2) the PCA-based technique with the Q statistic, applied without removing outliers from the original data set (“PCA Q ”); 3) principal component pursuit, as described in Section III (“PCP”); 4) a robust PCA-based technique with the T^2 statistic, applied after removing outliers from the data using PCP (“rPCA T^2 ”); and 5) a robust PCA-based technique with the Q statistic, applied after removing outliers from the data using PCP (“rPCA Q ”). Fault isolation and reconstruction for the PCA-based methods was performed using the reconstruction-based contribution method, with fault directions of cardinality one.

Principal component pursuit was performed using an augmented Lagrangian method [9], with parameter $\lambda = \frac{1}{\sqrt{n}}$.




Fig. 2. Number of faults consistently isolated (isolated in at least three of the four operating conditions) as a function of threshold for the process monitoring methods. The threshold for the PCA-based methods is expressed as a multiplier of the median value of the fault detection statistic over the training set. The threshold for separation of faults from noise for the PCP method is expressed in terms of number of standard deviations of the process variable in the training set.

Twenty principal components were retained for the PCA-based techniques.

C. Experimental Results

The fault-free process data Y and the fault and noise data Z reconstructed by PCP are shown in Figure 3.

The number of faults detected and isolated is a function of the threshold applied to the fault detection statistic (for the PCA based methods) and the threshold for separating noise from faults (for the PCP-based method). Because faults were not typically corrected during the course of the quality assurance test, a fault was said to be *consistently* isolated if the particular fault was isolated for a particular power plant in at least three of the four operating conditions. The number of faults consistently isolated for each of the process monitoring methods, as a function of the threshold, is shown in Figure 2; PCP is capable of consistently isolating more faults than the other methods. An engineer familiar with the quality-control process indicated that the majority of the faults consistently isolated by PCP corresponded to known sensor faults.

VII. CONCLUSIONS

Like PCA-based approaches for process monitoring, the PCP-based approach to process monitoring is based on the premise that process data in the absence of faults and noise is sampled from a low-dimensional subspace. The PCP-based approach remedies the key shortcomings of PCA-based methods and allows faults and noise to be isolated from otherwise low rank process data.

REFERENCES

- [1] L. Chiang, E. Russell, and R. Braatz, *Fault detection and diagnosis in industrial systems*. Springer Verlag, 2001.
- [2] C. Alcalá and S. Qin, “Reconstruction-based contribution for process monitoring,” *Automatica*, vol. 45, no. 7, pp. 1593–1600, 2009.
- [3] L. Chiang, R. Pell, and M. Seasholtz, “Exploring process data with the use of robust outlier detection algorithms,” *Journal of Process Control*, vol. 13, no. 5, pp. 437–449, 2003.

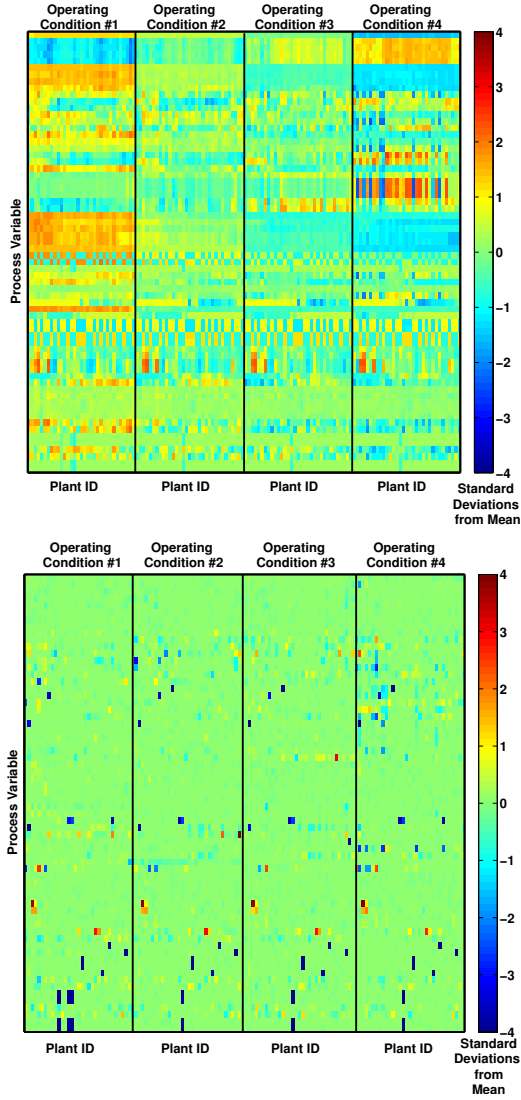


Fig. 3. Reconstructed process data (top) and fault and noise data (bottom) obtained using PCP.

- [4] E. Candes, X. Li, Y. Ma, and J. Wright, "Robust principal component analysis?," *Arxiv preprint arXiv:0912.3599*, 2009.
- [5] H. Vedam and V. Venkatasubramanian, "PCA-SDG based process monitoring and fault diagnosis," *Control Engineering Practice*, vol. 7, no. 7, pp. 903–917, 1999.
- [6] J. Gertler, W. Li, Y. Huang, and T. McAvoy, "Isolation enhanced principal component analysis," *AICHE Journal*, vol. 45, no. 2, pp. 323–334, 1999.
- [7] V. Chandrasekaran, S. Sanghavi, P. Parrilo, and A. Willsky, "Rank-sparsity incoherence for matrix decomposition," *Arxiv preprint arXiv:0906.2220*, 2009.
- [8] A. Ganesh, J. Wright, X. Li, E. J. Candes, and Y. Ma, "Dense error correction for low-rank matrices via principal component pursuit," in *Proceedings of IEEE International Symposium on Information Theory*, June 2010.
- [9] Z. Lin, M. Chen, L. Wu, and Y. Ma, "The augmented lagrange multiplier method for exact recovery of corrupted low-rank matrices," *Arxiv preprint arXiv:1009.5055*, 2010.
- [10] L. N. Trefethen and D. Bau, *Numerical linear algebra*. SIAM, 1997.
- [11] G. Strang, *Linear algebra and its applications*. Harcourt Brace Jovanovich, 3rd ed., 1988.
- [12] A. Polianin and A. Manzhirov, *Handbook of mathematics for engineers and scientists*. Chapman & Hall/CRC, 2007.

APPENDIX

Lemma 1: The singular values of a matrix $\begin{bmatrix} Y \\ \mathbf{x} \end{bmatrix}$ are the singular values of $\begin{bmatrix} \Sigma & 0 \\ \mathbf{t} & \sqrt{Q} \end{bmatrix}$, where $Y = U\Sigma V^T$ is the SVD of Y , $\mathbf{t} = V^T \mathbf{x}$ and $\sqrt{Q} = \|(I - VV^T)\mathbf{x}\|$.

Proof: Any m by n matrix A can be factored with the singular value decomposition $A = U\Sigma V^T$, where U is unitary, V is unitary, and Σ is diagonal [10].

Note that

$$\begin{bmatrix} Y \\ \mathbf{x} \end{bmatrix} = \begin{bmatrix} U\Sigma V^T \\ \hat{\mathbf{x}} + \tilde{\mathbf{x}} \end{bmatrix} = \begin{bmatrix} U & 0 \\ 0 & I \end{bmatrix} \begin{bmatrix} \Sigma & 0 \\ \mathbf{t} & \sqrt{Q} \end{bmatrix} \begin{bmatrix} V \\ \mathbf{r} \end{bmatrix}^T \quad (15)$$

where $\mathbf{t} = V^T \mathbf{x}$, $(I - VV^T)\mathbf{x} = \sqrt{Q}\mathbf{r}$, \mathbf{r} is a unitary vector orthogonal to the columns of V , and \sqrt{Q} , a scalar, is the length of $(I - VV^T)\mathbf{x}$.

Let

$$U'\Sigma'V'^T \stackrel{\text{SVD}}{\leftarrow} \begin{bmatrix} \Sigma & 0 \\ \mathbf{t} & \sqrt{Q} \end{bmatrix}. \quad (16)$$

Then

$$\begin{bmatrix} U & 0 \\ 0 & I \end{bmatrix} U'\Sigma' \left(\begin{bmatrix} V \\ \mathbf{r} \end{bmatrix} V' \right)^T \quad (17)$$

is a singular value decomposition of $\begin{bmatrix} Y \\ \mathbf{x} \end{bmatrix}$ because $\begin{bmatrix} U & 0 \\ 0 & I \end{bmatrix} U'$ is a unitary matrix, Σ' is diagonal, and $\begin{bmatrix} V \\ \mathbf{r} \end{bmatrix} V'$ is a unitary matrix. ■

Lemma 2: Let $Y' = \begin{bmatrix} Y \\ \mathbf{x} \end{bmatrix}$, with $Y = U\Sigma V^T$ the singular value decomposition of Y . Then

$$\sqrt{Q + \sum_i t_i^2 + \sigma_i^2} \leq \|Y'\|_* \leq \sqrt{(n+1)(Q + \sum_i t_i^2 + \sigma_i^2)} \quad (18)$$

where the t_i are the elements of the row vector $\mathbf{t} = V^T \mathbf{x}$ and the σ_i are the diagonal elements of Σ .

Proof: The nuclear norm is the sum of the singular values of a matrix. By the previous lemma, the singular values of $\begin{bmatrix} Y \\ \mathbf{x} \end{bmatrix}$ are the singular values of $\begin{bmatrix} \Sigma & 0 \\ \mathbf{t} & \sqrt{Q} \end{bmatrix}$. The singular values of this matrix can be found by taking the square root of each eigenvalue of the matrix $\begin{bmatrix} \Sigma & 0 \\ \mathbf{t} & \sqrt{Q} \end{bmatrix} \begin{bmatrix} \Sigma & 0 \\ \mathbf{t} & \sqrt{Q} \end{bmatrix}^T$. By the trace theorem [11], $\sum_{j=1}^{n+1} \tau_j^2 = Q + \sum_i \sigma_i^2 + t_i^2$. We desire to establish bounds on the sum of the singular values $\sum_{j=1}^{n+1} \tau_j$ in terms of a known sum of the eigenvalues $\sum_{j=1}^{n+1} \tau_j^2$. The lower bound follows from the generalized triangle inequality $\sqrt{\sum u_j} \leq \sum \sqrt{u_j}$ [12].

The upper bound follows from the fact that the square root is a concave function. For a concave function $f(u)$ and nonnegative mixing weights $q_i : \sum q_i = 1$,

$$f(q_1 u_1 + \dots + q_n u_n) \geq q_1 f(u_1) + \dots + q_n f(u_n). \quad (19)$$

The upper bound follows if one sets $f(u) = \sqrt{u}$ and $q_1 = q_2 = \dots = q_{n+1} = \frac{1}{n+1}$. ■