Journal of Water Process Engineering 38 (2020) 101556

Contents lists available at ScienceDirect

Journal of Water Process Engineering

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

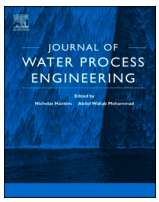

Case studies in real-time fault isolation in a decentralized wastewater treatment facility

Molly C. Klanderman [a], Kathryn B. Newhart [b], Tzahi Y. Cath [b], Amanda S. Hering [a,*]

[a] Department of Statistical Science, Baylor University, One Bear Place 97140, Waco, TX 76798, USA
[b] Department of Civil and Environmental Engineering, Colorado School of Mines, Golden, CO 80401, USA

ARTICLE INFO

Keywords:
Fault detection
Fault isolation
Lasso
PCA
Wastewater treatment

ABSTRACT

Decentralized wastewater treatment (WWT) can be an energy and resource efficient alternative to the traditional, centralized WWT paradigm for water-stressed communities. However, to operate economically, decentralized facilities do not typically have a WWT operator on-site full-time, so a real-time monitoring scheme is needed to quickly detect system faults and isolate the features associated with or affected by faults to ensure adequate treated water quality. Data collected from WWT facilities exhibit temporal dependence and experience natural fluctuations in the mean due to environmental and operator-controlled factors, violating the assumptions of many existing fault detection and isolation (FD&I) methods. To address this, we develop a complete data-driven FD&I method tuned to handle the unique features of WWT data that can be run in real-time and illustrate how it performs with data from a decentralized WWT facility in Golden, Colorado, USA. Enhanced visualization techniques are designed to assist operators in identifying features associated with the fault. We present three case studies with known faults and demonstrate how this method can aid operators in detecting and diagnosing the cause of a fault more quickly.

1. Introduction

In the traditional water management paradigm, municipal and industrial wastewater is collected and conveyed to centralized wastewater treatment (WWT) facilities for removal of contaminants before discharge to the environment (e.g., rivers, lakes). However, the demand for additional water resources in arid or water-stressed regions makes decentralized WWT systems an attractive option. In decentralized WWT systems, wastewater is treated close to the point of generation, which allows the treated water to be reused locally. Treatment can even be tailored to provide water of different qualities for various reuse applications such as toilet flushing or irrigation. Not only can this new paradigm provide an additional water source for a community's water portfolio, but it is also more resource and energy efficient because it does not require pumping wastewater over long distances for collection and distribution [1–3].

The major drawback of decentralized WWT is that these smaller facilities have higher operating expenses. It is often not feasible for decentralized facilities to have an operator on-site twenty-four hours a day to monitor system operations. However, without operator oversight, faults can quickly propagate through the system, which can have serious operational, financial, environmental, or public health implications. Thus, a fault monitoring platform is needed that would allow one operator to simultaneously monitor multiple facilities remotely. *Fault detection* (FD) is the detection of abnormal system behavior, or behavior that deviates from accepted, in control (IC) conditions. In a high-dimensional setting like WWT in which dozens of process variables are monitored, knowledge of the presence of out of control (OC) conditions is not completely informative; it is also important to identify the variables within the system that have shifted, known as *fault isolation* (FI). Then, operators are well-equipped to begin diagnosing the fault and correcting it.

Newhart et al. [4] summarize the existing monitoring paradigms used in the majority of WWT facilities and emphasize the need for more advanced data-driven monitoring strategies tailored for WWT. The most common process control approach in WWT is to identify an acceptable for each monitored variable and when the value exceeds a threshold, the variable is identified as OC, and an alarm is triggered. However, static univariate thresholds tend to detect faults slowly because they do not account for the relationship among variables within the system [4]. Typically, after an alarm is triggered by one variable in the system, operators attempt to identify other variables affected by the fault in

* Corresponding author.
E-mail address: mandy_hering@baylor.edu (A.S. Hering).

https://doi.org/10.1016/j.jwpe.2020.101556
Received 18 May 2020; Received in revised form 15 July 2020; Accepted 21 July 2020

order to diagnose the cause. This course of action includes visually interpreting time series plots of the individual variables. Even for highly-skilled operators who are familiar with the system, it is difficult to visually identify the affected variables among many; especially for faults that increase in magnitude and severity slowly over time. Therefore, an advanced real-time monitoring strategy is required to detect faults in a WWT system *and* identify the components associated with the fault.

A wealth of statistical literature exists for FD and FI, collectively referred to as *multivariate statistical process monitoring* (MSPM). Review papers by Woodall and Montgomery [5] and Reis and Gins [6] provide a good overview of the current methods and problems in the field. Control charts such as the multivariate exponentially weighted moving average (MEWMA) chart [7], Hotelling's T^2 chart [8], and multivariate cumulative sum (MCUSUM) chart [9] remain popular FD methods among practitioners, and many variations and improvements have been proposed (e.g., [10–15]). Common FI methods include contribution plots, variable thinning methods, and post-signal diagnostic variable selection methods. Contribution plots visualize the contribution of each variable to the fault, but the contribution of one variable can propagate to other variables, known as the "smearing effect," which can lead to misdiagnosis [16]. An alternative to contribution plots that reduces the smearing effect is reconstruction-based contribution (RBC) [17–19], but RBC can still result in misdiagnosis in complex fault settings in which multiple variables are affected. Variable thinning methods focus on thinning or reducing the number of variables *prior* to monitoring in order to have more power in detecting the fault [20–25]. Alternatively, some methods monitor all variables, and after a fault is detected, variable selection approaches such as forward selection and lasso estimation are used to identify the shifted variables [26–29].

Most fault detection and isolation (FD&I) methods require an assumption that the data are independent and identically distributed; many also require a distributional assumption such as multivariate normality. Data collected from decentralized WWT facilities violate these assumptions, making most MSPM methods inappropriate for direct application. Influent is more highly variable in quality and quantity than that of centralized facilities, resulting in a natural fluctuation in variables over the course of the day, week, and season, which causes *nonstationarity* in the mean. Furthermore, changes in operational settings such as flow rates or whether a valve is open or closed can impact monitored variables, and these changes are not indicative of a fault. Because the data are collected at a high frequency (e.g., seconds to minutes), the data have a strong positive *autocorrelation* present. Finally, the data often do not follow a normal distribution and may exhibit some skewness. Because of the nonstationarity, autocorrelation, and non-normality present in WWT facility data, it can be especially challenging to detect faults. Ignoring these characteristics can result in a high false alarm rate, which can be costly in terms of time and money for a decentralized WWT facility. Using data to then isolate the shifted variables can be even more challenging.

FD in WWT settings with real data has been addressed in a few papers (e.g., [30–35]), but it is much more common to use artificial simulated data from models such as the Benchmark Simulation Model (e.g., [36–40]). When real data are used, often only a small subset of the monitored variables (less than 10) are considered, which artificially reduces the complexity and difficulty of monitoring in this setting [31,32,34,35,41]. Kazor et al. [30] and Odom et al. [33] developed advanced real-time FD methods that were able to effectively monitor more than 25 variables. However, FI has never been applied to WWT processes outside of Klanderman et al. [42], who proposed a retrospective FI method using fused lasso to identify shifted variables in a period of historical data, so this is the first time wherein real-time FI has been tested in a WWT setting.

In this paper, we present a fully integrated, real-time FD&I method that accounts for nonstationarity, autocorrelation, and non-normality, all of which are present in WWT data. First, we account for nonstationarity by removing any trends due to environmental or operational factors, which is a common approach among practitioners [42–46]. Then, we apply an FD method based on principal component analysis (PCA) proposed by Kazor et al. [30] to the detrended data, which can account for temporal dependence, non-normality, and nonstationarity caused by long-term or seasonal trends not captured in the detrending process. For FI, we apply the post-signal diagnostic variable selection method proposed by Ebrahimi et al. [28], which uses adaptive lasso to recover the shifted variables from the principal components (PCs).

Using three case studies from a demonstration-scale decentralized WWT facility, we illustrate the method's performance through an enhanced visualization tool that we design to assist operators in the fault diagnosis step. We compare the results of the method with the operators' post-hoc analysis of the fault to see how quickly the method detects the fault, which variables are identified as shifted, and if the analysis can provide any additional insight into the fault. The major contributions of our work are (1) the integration of FD and FI for a fully functional FD&I methodology; (2) the development of an enhanced visualization tool to assist with diagnosis; and (3) the illustration of the method with real and complex cases. The approach that we develop is tailored to the unique features of decentralized WWT data and is the first time that an FD&I approach has been applied to such a complex system, but the methodology is flexible enough to be applied in other fields of application. Section 2 introduces the real-time FD&I method; Section 3 presents three case studies with known faults from a WWT facility in Golden, Colorado, USA; and Section 4 discusses the capabilities and limitations of the method when considering applying it to similar systems.

2. Methodology

The method is comprised of three primary components: training and detrending, detection, and isolation. A flow chart summarizing the method is shown in Fig. 1, which is described in the remainder of this section.

2.1. Training and detrending

Prior to monitoring, we must identify the *response variables*, or the variables in the system that we wish to monitor, and *predictor variables*, which can explain some of the variability in the response variables. We denote the p response variables and q predictor variables observed at time t as column vectors x_t and w_t, respectively. The first τ observations, which are known to be IC, are used as an initial training window. Numerous methods known as *Phase I* methods exist to identify an appropriate IC data set for training, estimate parameters, and set thresholds [47–49], and these are used in real-time monitoring schemes known as *Phase II* methods. However, Phase I methods make many of the same assumptions about the data as Phase II methods do and are therefore not appropriate for WWT data. Instead, to identify an IC period, we rely on the expertise of operators and other subject matter experts. Occasional outliers are common due to electrical surges in one or more sensors, so to account for any such extreme values, we linearly interpolate any observations in which at least one of the values exceeds 10 standard deviations of the median value in the training set. In the three case studies, zero observations were imputed for case study #1, two were imputed for case study #2, and one was imputed for case study #3.

Prior to monitoring, we detrend the response variables using the predictor variables to account for any variability in the system from known sources such as system state, control settings, or other external factors that are not indicative of a fault. Multi-state methods such as Odom et al. [33] have also been proposed, but a detrending approach provides greater flexibility in the sources of variability that can be accounted for in the model. For the i-th response variable at time t,

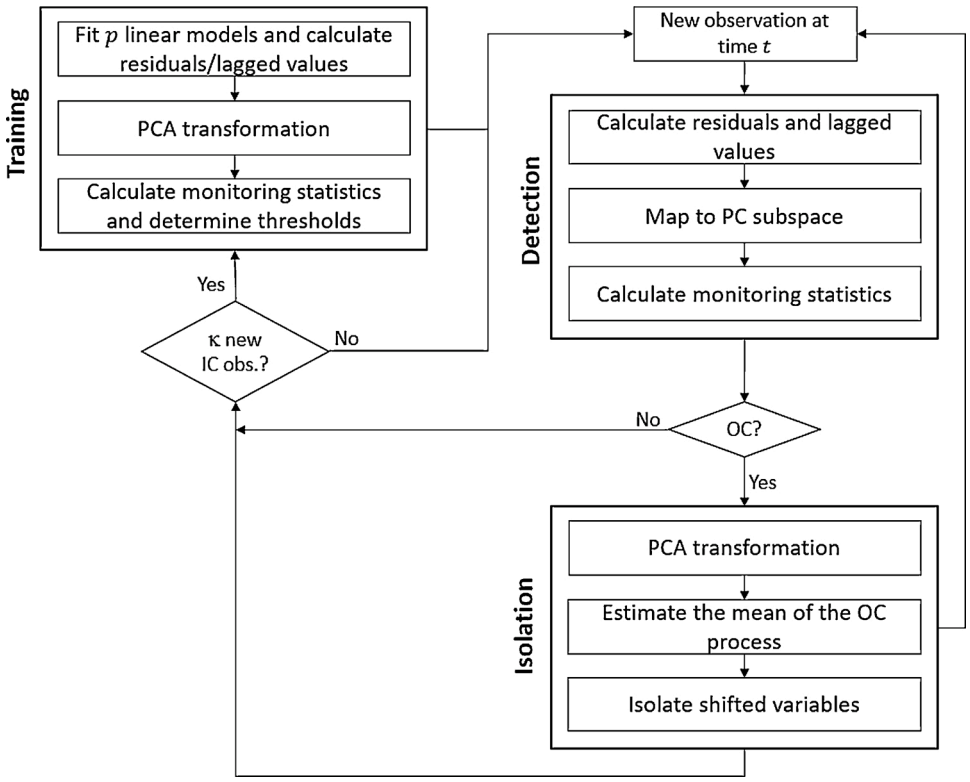

Fig. 1. Flow chart of the FD&I method used in the three case studies.

which we denote $x_t^{(i)}$, we fit a linear model

$$\mathbb{E}\left[x_t^{(i)}\right] = \beta_0^{(i)} + \mathrm{w}_t^T \boldsymbol{\beta}_1^{(i)}, \quad t = 1, 2, \ldots, \tau, \tag{1}$$

using the training window of observations, where $\beta_0^{(i)}$ and $\boldsymbol{\beta}_1^{(i)}$ are the intercept and regression coefficients, respectively, of the linear model for the i-th response variable. Using the fitted model, we can find the fitted value $\widehat{x}_t^{(i)} = \widehat{\beta}_0^{(i)} + \mathrm{w}_t^T \widehat{\boldsymbol{\beta}}_1^{(i)}$ for the i-th response variable at time t, where $\widehat{\beta}_0^{(i)}$ and $\widehat{\boldsymbol{\beta}}_1^{(i)}$ are the estimated intercept and regression coefficients, respectively. The residuals for the i-th variable, $r_t^{(i)} = x_t^{(i)} - \widehat{x}_t^{(i)}$, contain information about the behavior of the process that is not accounted for by other factors. We can collect the residuals for all variables at time t into a single vector $\mathrm{r}_t = \mathrm{vec}(r_t^{(1)}, r_t^{(2)}, \ldots, r_t^{(p)})$, which we monitor for faults using a method known as adaptive dynamic PCA (AD-PCA) described in Kazor et al. [30].

2.2. Detection

PCA is a common technique used in MSPM to reduce the dimension of the data while retaining as much information as possible. Two modifications, adaptive and dynamic, have been proposed to account for nonstationarity and autocorrelation, respectively. The adaptive extension of PCA regularly updates the training window to account for nonstationarity in the data over time [50,51]. In this case, the adjustment allows us to account for nonstationarity due to long-term or seasonal trends that remain after the effects of the predictor variables have been removed. To use the dynamic extension of PCA, we augment r_t with a *lagged* version of the residuals, which we denote r_t^+ [52]. For each variable, we use the lag with the highest partial autocorrelation function from the training window, so r_t^+ is of length $2p$. When applied to nonlinearly related variables, Kazor et al. [30] demonstrate through a simulation study of autocorrelated and nonstationary data that AD-PCA outperforms PCA in terms of the false alarm rate under IC conditions and detection speed under OC conditions.

Applying PCA to the correlation matrix of the residuals from the training window, $\mathrm{r}_1^+, \mathrm{r}_2^+, \ldots, \mathrm{r}_\tau^+$, we obtain a $2p \times k$ linear projection matrix P_k^+ with associated eigenvalues $\boldsymbol{\Lambda}_k^+ = \mathrm{diag}(\lambda_1^+, \ldots, \lambda_k^+)$, where k is the number of PCs necessary to capture 80% of the variability. Therefore, an observation projected into the PC subspace is $\mathrm{y}_t^+ = \left(\mathrm{P}_k^+\right)^T \mathrm{r}_t^+$. For each observation in the training period, we calculate two monitoring statistics: Hotelling's T^2 and squared prediction error (SPE), which we denote as Q. Hotelling's T^2 is calculated as

$$T^2 = \left(\mathrm{y}_t^+\right)^T \left(\boldsymbol{\Lambda}_k^+\right)^{-1} \mathrm{y}_t^+ = \left(\mathrm{r}_t^+\right)^T \mathrm{P}_k^+ \left(\boldsymbol{\Lambda}_k^+\right)^{-1} \left(\mathrm{P}_k^+\right)^T \mathrm{r}_t^+,$$

which measures the variability in the PC subspace. The squared prediction error (SPE) measures the variability in the residual subspace and is calculated as

$$Q = \left\| \left[\mathrm{I} - \mathrm{P}_k^+ \left(\mathrm{P}_k^+\right)^T \right] \mathrm{r}_t^+ \right\|^2,$$

where I is the $2p \times 2p$ identity matrix. Then, we can estimate the IC distribution and determine thresholds of T^2 and Q based on a prespecified false alarm rate, α. The use of parametric thresholds is popular among practitioners [5], but they can be highly affected by deviations from normality. Therefore, if the underlying process distribution is not known sufficiently to assume normality, nonparametric adjustments are recommended [53]. We identify the nonparametric thresholds using the $(1-\alpha)\%$ quantile of the kernel density estimate (KDE) of T^2 and Q from the training period.

Given a new observation x_t for $t = \tau + 1, \ldots$, we calculate the residuals based on the IC fitted model from Eq. (1), incorporate lagged values to produce r_t^+, map r_t^+ into the lower-dimensional subspace y_t^+, and compute T^2 and Q. If either statistic exceeds the corresponding threshold, we flag the observation as OC and perform FI. Otherwise, we continue to monitor new observations. Using the adaptive version of PCA, we retain the observations classified as IC by AD-PCA and update

the training window when we obtain κ new IC observations, removing the oldest κ observations from the training window. Each time we update the training window, we (1) re-estimate p fitted models to calculate the residuals for each variable; (2) re-estimate P_k^+ and Λ_k^+; and (3) recalculate the nonparametric thresholds for the monitoring statistics T^2 and Q.

2.3. Isolation

When an observation is flagged as OC, we want to identify which of the variables' means have shifted. The residuals from the training window are of mean zero, so identifying the shifted variables is equivalent to identifying the non-zero components of the mean of r_t, which we denote $\boldsymbol{\mu}_t$. The method proposed by Ebrahimi et al. [28] recovers $\boldsymbol{\mu}_t$ from the PCs. However, in order to isolate the non-zero components of the mean of r_t rather than r_t^+, we re-estimate the PCs as a function of r_t. Given the *entire* linear projection matrix P_p with associated eigenvalues $\boldsymbol{\Lambda}_p = \mathrm{diag}(\lambda_1, \ldots, \lambda_p)$, the PC score is

$$\mathrm{y}_t = \mathrm{P}_p^T \mathrm{r}_t = \mathrm{P}_p^T \boldsymbol{\mu}_t + \boldsymbol{\epsilon},$$

where $\boldsymbol{\epsilon}$ is the error term with mean zero and covariance $\boldsymbol{\Lambda}_p$.

To avoid any issues of non-constant variance across the PCs, we transform y_t and P_p as $\mathrm{y}_t^* = \boldsymbol{\Lambda}_p^{-1/2}\mathrm{y}_t$ and $\mathrm{P}_p^* = \boldsymbol{\Lambda}_p^{-1/2}\mathrm{P}_p$. Then, an estimate of the mean of the OC process is found using adaptive lasso, which enforces sparsity in $\widehat{\boldsymbol{\mu}}_t$ by putting a penalty on the components of $\widehat{\boldsymbol{\mu}}_t$ and driving small components to zero. An estimate of $\widehat{\boldsymbol{\mu}}_t$ is given by

$$\widehat{\boldsymbol{\mu}}_t = \underset{\boldsymbol{\mu}_t}{argmin} \|\mathrm{y}_t^* - (\mathrm{P}_p^*)^T \boldsymbol{\mu}_t\|_2^2 + \gamma \sum_{j=1}^{p} w_j |\mu_j|,$$

where γ is a regularization parameter, and $\mathrm{w} = \mathrm{vec}(w_1, \ldots, w_p) = \frac{1}{\widehat{\boldsymbol{\mu}}_t^{\mathrm{OLS}}}$. We solve this using the R package glmnet [54], where the regularization parameter γ is chosen using cross-validation. All variables with non-zero components of $\widehat{\boldsymbol{\mu}}_t$ are classified as *shifted* at time t; all variables with zero components of $\widehat{\boldsymbol{\mu}}_t$ are classified as *unshifted* at time t. We refer to this FD&I method as adaptive dynamic sparse PCA, or ADS-PCA, because we are enforcing sparsity in the components of $\widehat{\boldsymbol{\mu}}_t$.

3. Case studies

We demonstrate the performance of ADS-PCA by applying it to three known faults from Mines Park Water Reclamation Test Site (Mines Park), a demonstration-scale decentralized WWT facility in Golden, Colorado, USA. The facility treats approximately 7000 gallons of municipal wastewater daily from a university student apartment complex on the campus of Colorado School of Mines (Mines). This facility has been used as a case study to develop and study MSPM methods [30,31,33,35,42].

3.1. Mines Park system description

The system at Mines Park is a sequencing-batch membrane bioreactor (SB-MBR), which uses a combination of physical, chemical, and biological processes to treat municipal wastewater. There are two major subsystems, the two bioreactors (BRs) and two membrane bioreactors (MBRs), which are connected by the return activated sludge (RAS) and the MBR overflow. A diagram of the two subsystems is shown in Fig. 2.

Wastewater is diverted from the municipal sewer system to a holding tank. Then, once an hour, a batch of raw wastewater is delivered to one

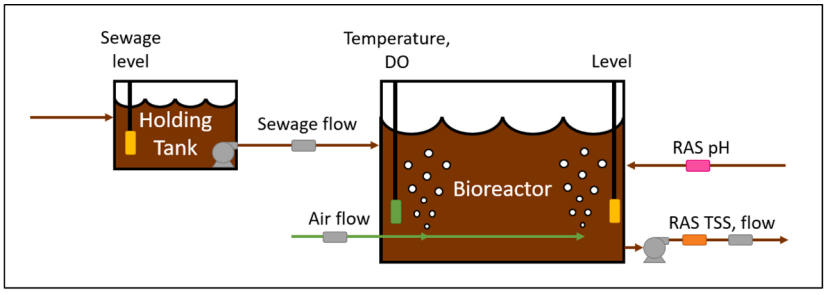

(a) BR subsystem

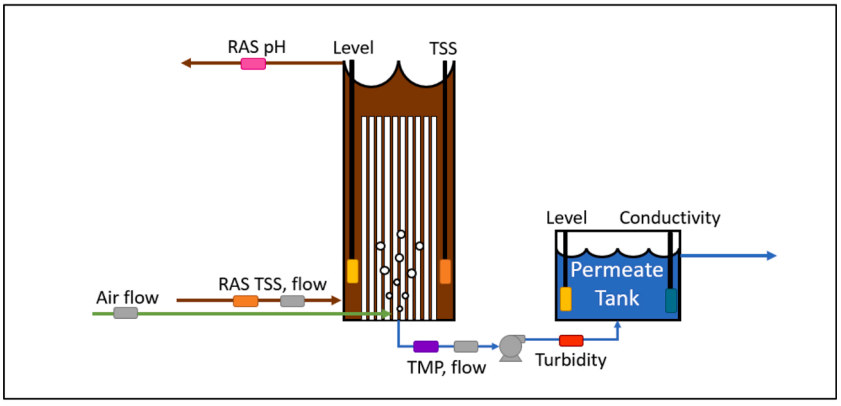

(b) MBR subsystem

Fig. 2. Process diagram of the two subsystems in the SB-MBR system at Mines Park. The sensors used to collect data on the response variables in the system are denoted throughout the diagram using colored rectangles. Note, TMP represents transmembrane pressure, and TSS represents total suspended solids.

of two BRs using a submerged grinder pump and a 2 mm fine screen to remove large, inert solids. The carbon and nutrient-rich wastewater is digested by the dense microbial mass in the BRs, known as *activated sludge*. To activate particular microorganisms to degrade specific contaminants, oxygen is provided via industrial air blowers. These air blowers are timed strategically to efficiently transform organic and inorganic nitrogen into nitrogen gas and phosphorus into microbial biomass, thus removing the nutrients from the water. When the blowers are on, dissolved oxygen (DO) sensors measure the concentration of aqueous oxygen, and a proportional-integral controller increases or decreases the air blower speed to maintain a DO concentration setpoint.

After an hour of alternating aeration conditions (i.e., aerobic, anaerobic), a RAS pump circulates the activated sludge between one of the BRs and the two MBRs operating in parallel. The activated sludge flows from the bottom of the fully-mixed MBR tanks upward and overflows to a trough, where the activated sludge is returned to the original BR. The MBR achieves solid–liquid separation of the activated sludge using hollow-fiber membranes with a nominal pore size of 0.03 μm, rejecting bacteria and some viruses. Aeration is provided to the MBRs to mix the activated sludge and to scour the solids that may accumulate on the membrane surface. A pump pulls water through the membrane, known as *permeate*, to a holding tank where it will overflow to a second, larger basin for subsequent reuse. The two BRs alternate between filling/aeration and recirculating to provide a continuous flow of activated sludge to the MBRs and, consequently, a continuous flow of permeate.

The faults observed at Mines Park derive from two primary sources: operational or environmental. Faults due to operational failures, such as a clog in a pipe or equipment malfunction, can propagate quickly throughout the system, and evidence of these problems are associated with changes in the behavior of variables such as flow rates and equipment speed. Operational faults require immediate corrective action and can result in system downtime in which either no water can pass through the system or water passes through under-treated. However, operational faults are usually short-lived and can be addressed with sufficient labor hours, tools, and equipment. Faults caused by environmental changes can have more long-lasting and severe consequences, but are slower to impact system variables because influent water is heavily diluted by water already in the system. Regardless, the microorganisms responsible for treating the wastewater are sensitive to the environmental conditions of the BRs, so even slight changes in water quality can threaten their survival. Diagnosing an environmental fault is difficult because regulated water quality measurements are only taken a few times a week, and once a change has been identified in these measurements, the biological system has already been compromised and requires supplemental chemical addition or drastic operational changes to recover.

3.2. Data description and detrending

More than 30 variables are recorded every minute in the Mines Park SB-MBR using sensors placed throughout the system, which are denoted by colored rectangles in Fig. 2. First, we identify the set of response variables that we wish to monitor and the corresponding predictor variables needed to detrend the response variables. Information on the system's state and control settings known as *control variables* are recorded every minute, and these can explain some variability in the behavior of the response variables. For example, if a blower's air flow rate increases, it is expected that the DO concentration in the BR will increase as well. Often, there is a delay in the effect of a control variable on a response variable, so we account for this delay by using lagged versions of the control variables in the fitted regression models given in Eq. (1). Using the partial cross correlation function (PCCF) between the control and response variable, we select the lag with the largest PCCF value. Based on exploratory data analysis, we also identified an hourly, two-hour, and daily trend in the observations, so we include *cyclic variables* in the form of sine/cosine pairs with a corresponding period to account for these temporal frequencies. Together, the control variables, lagged control variables, and cyclic variables make up the set of predictor variables used for detrending.

The operation of the two subsystems (BR and MBR) at a given time are semi-independent; over time, they become dependent if a fault goes undetected and propagates throughout the entire system, but there is a delay in the impact to the other subsystem. Therefore, we monitor the response variables from the two subsystems separately and include the RAS-related variables in both data sets. We monitor 11 variables in the BR subsystem and 16 variables in the MBR subsystem. For some response variables, we found that including control variables from both subsystems substantially increased the predictive power of the model, so we include the control variables from both subsystems to model all response variables. A list of all response, control, and cyclic variables is included in the supplementary materials.

For ADS-PCA, we use a training window of $\tau = 7200$ observations, or five days of observations monitored every minute. We update the model every $\kappa = 1440$ observations, which is equivalent to one day of observations assuming every observation is classified as IC. We must also specify the false alarm rate α, which determines the thresholds for the monitoring statistics. A larger α will result in a smaller threshold, so faults will be detected sooner, but there will also be more false alarms. We set the false alarm rate to be $\alpha = 0.005$, which means that when the process is IC, we can expect a false alarm every 200 observations on average, which translates to one false alarm every 3 h and 20 min.

For each response variable, we fit a linear model using adaptive lasso [55], which performs parameter estimation and variable selection simultaneously to remove some of the unimportant predictor variables from the model. Lasso regression methods are sensitive to differences in scale among variables, so prior to model fitting, we standardize the predictor and response variables from the training period. An example of the fitted models using a heat map to show the estimated regression coefficients is included in the supplementary materials.

Fig. 3 shows an example of the detrending process for sewage_level, one of the shifted variables, from case study #3. A time series plot of sewage_level is shown in Fig. 3(a) with the observed values in black and fitted values in red, and the residuals are shown in black in Fig. 3(b). The green vertical lines in both plots indicate when the model is updated, where the first five days of observations are the initial training window. A fluctuation in the mean of sewage_level is evident over time, even throughout the training period. The fitted model is able to capture some of the nonstationarity with an R^2 of 30% during the initial training window, reducing the variability in the residuals. On January 29, there is a dramatic decrease in sewage_level, which is not captured by the fitted model, and is therefore preserved for detection by ADS-PCA.

For these case studies, we present an enhanced visualization technique called a *checkerboard plot*, which summarizes when faults are detected, which response variables are shifted for each OC observation, and in which direction the shift occurred for each variable simultaneously. This provides an easy and concise way to communicate the health of the system to operators in real-time. Using this information, operators can then refer to the individual time series plots to diagnose and assess the severity of the fault.

Variations of checkerboard plots have been used to visualize FD&I results [19,56], but to our knowledge, none have used color to indicate the direction of the fault. Including the direction of the fault makes it much easier for an operator to diagnose the fault. For example, a downward shift in the concentration of total suspended solids (TSS) could indicate an environmental fault in which the biomass is decaying faster than it is growing. However, an upward shift in TSS could indicate an operational fault such as a valve stuck in the closed position, causing activated sludge to continuously accumulate in the system.

The checkerboard plots for the three case studies are shown in Figs. 4, 5, and 7 . Each row represents a response variable over time, and the shifted variables that were identified by the operators *retrospectively* are highlighted in yellow. At each time point, a white box indicates that

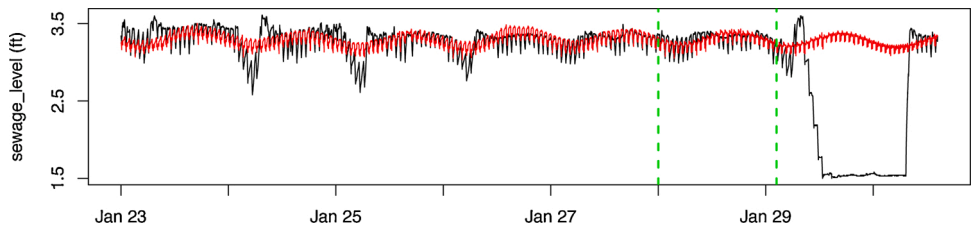

(a) Observed and fitted values

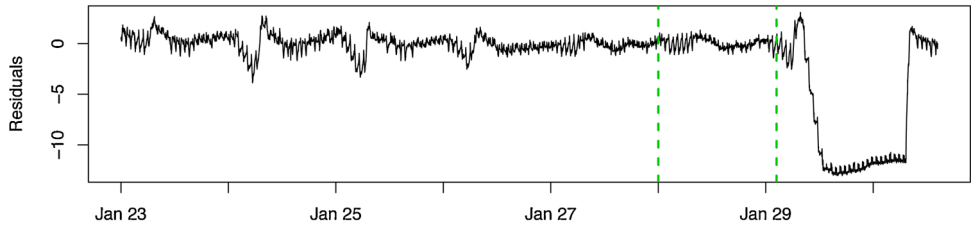

(b) Residuals

Fig. 3. Plot of (a) sewage_level in black with the fitted values in red, and (b) the residuals in black. The green vertical lines indicate when the model is updated, where the first five days of observations are the initial training window. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

the variable is classified as unshifted; a red box indicates that the variable has an upward shift detected; and a blue box indicates that the variable has a downward shift detected. The color gets darker as the number of consecutive alarms increases using a gradient color scale. The lightest blue indicates only a single alarm, and the darkest blue indicates that at least five alarms in a row have occurred. Note that zooming into Figs. 4, 5, and 7 will provide greater resolution and detail in the plots. Case studies #1 and #2 have two checkerboard plots, one for each of the two subsystems; case study #3 only has one checkerboard plot for the BR subsystem, and we will elaborate further in Section 3.5.

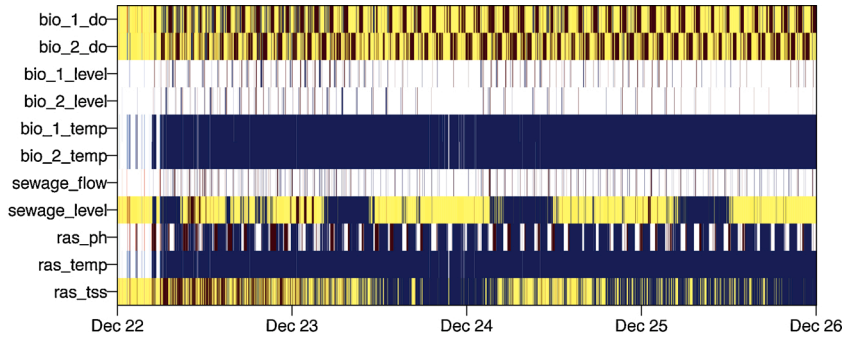

(a) BR subsystem

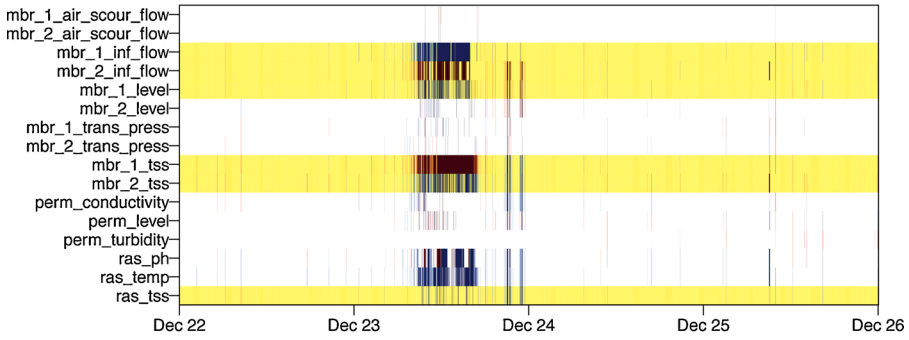

(b) MBR subsystem

Fig. 4. Checkerboard plot of monitored variables over time for case study #1, where the rows represent the monitored variables; red boxes represent a variable with an upward shift detected; blue boxes represent a variable with a downward shift detected; white boxes represent a variable classified as unshifted; and truly shifted variables are highlighted in yellow. The color of the red and blue boxes gets darker as the number of consecutive alarms increases. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

3.3. Case study #1

The first case study focuses on a fault at the end of December 2017 in which the system suffered from an operational fault and an environmental fault simultaneously that affected both the BR and MBR subsystems. The training window for case study #1 is December 17–21, 2017, followed by a monitoring period from December 22–25, 2017. First, we focus on the fault in the BR subsystem. In mid-December 2017, students at Mines finished taking final exams, and most left campus for the holiday break, which is a common situation for decentralized WWT facilities that serve seasonal communities. As a result, the hourly wastewater flow to replenish the sewage tank declined, and over the following week, the sewage_level began to drop, which is detected by the method beginning on December 22 as seen in Fig. 4(a).

These low flow conditions were compounded by a change in water quality. The wastewater was not only less concentrated, but the longer retention times in the collection system most likely led to preemptive carbon degradation. The carbon concentration of the raw wastewater is directly related to the oxygen demand, which is how much oxygen is required to fully treat the water of all oxidizeable material. At Mines Park, IC water quality conditions usually require almost a full hour of the air blowers at 100% speed to reach a DO concentration of 2 mg/L. However, during the monitoring period, when the blowers were turned on every 10-20 min, the little remaining carbon was quickly oxidized, and the DO concentration setpoint was easily achieved (and in some cases exceeded). This increase relative to IC conditions is evident in the intermittent upward shifts detected in bio_1_do and bio_2_do.

The shifts in bio_1_temp, bio_2_temp, and ras_temp are most likely false alarms due to the nonstationarity of the temperature variables, which is still present and not fully removed from the residuals. The downward shifts in ras_ph are also likely false alarms because the sensor for the pH concentration rounds the value to two significant digits, resulting in only four unique values observed over the course of the five-day training period. With so little variability in ras_ph, the FI method tends to classify the variable as shifted when it deviates from the most common observed value in the training period.

One day after the fault is detected by the method in the BRs, a fault is also detected in the MBRs, which is shown in Fig. 4(b). Mid-day on December 23, a clog developed in the inlet to MBR 1, causing a decrease in the mbr_1_inf_flow. Because the flow from the RAS pump splits into the two MBRs, the clog to MBR 1 caused a corresponding increase in the mbr_2_inf_flow. The decreased flow to MBR 1 and increased flow to MBR 2 resulted in a concentration and dilution of solids in MBR 1 and 2, respectively, which caused an increase in mbr_1_tss and decrease in mbr_2_tss. A few hours later, the low mbr_1_inf_flow caused a drop in the mbr_1_level, so very little (if any) water was overflowing into the trough that returns the activated sludge into the BRs. Therefore, the only activated sludge returned to the BRs was from MBR 2, which had been diluted by the increased flow rate, resulting in a decrease in ras_tss as well, which is detected in both subsystems before the end of the day on December 23.

When the low MBR 1 influent flow rate alarmed (i.e., fell below an operator-determined threshold), the operator artificially increased the flow to MBR 1 by exercising the valve where the clog was most likely located, restricting the flow to MBR 2, and increasing the pressure head on the RAS pump by increasing BR levels. This dislodged the clog and subsequently equalized the distribution of solids in the MBRs. As a result, the TSS concentration of the MBRs and RAS returned to normal operating conditions after temporary operational changes.

3.4. Case study #2

The training window for case study #2 begins shortly after the conclusion of the monitoring period for case study #1 and covers the period from December 30, 2017 to January 3, 2018, which is followed by a four-day monitoring period. The fault of interest occurred exclusively in the MBR, but we present the results from ADS-PCA in the BR as well to demonstrate the importance of the choice of the IC training period on the results. The ADS-PCA results for case study #2 are shown in Fig. 5.

The fault in the BR from case study #1 was not detected or corrected by operators at the time, so it carries over into the training and monitoring period for case study #2. Fig. 6 illustrates this behavior for bio_1_do, and both bio_2_do and sewage_level exhibit similar patterns. In the training window for case study #1, which is represented by the left green region, bio_1_do is truly IC but becomes OC towards the end of the training period. When bio_1_do continues to increase during the monitoring period of case study #1, the left blue region, the observations are flagged as OC by ADS-PCA. The fault was not detected by operators in the BR, so we used our method with a training window that contained OC data in case study #2, the right green region. As a result, Fig. 5(a) shows that our method does not flag any observations as OC during the monitoring period of case study #2, the right blue region, because the OC behavior in these variables is now considered to be IC by the method. This is to be expected given that the method is purely data-driven.

There is also OC behavior present in the training window for the MBR that was initially overlooked by operators. Around January 4, the beginning of the monitoring period, full-time operators returned to campus from the holidays and identified a decrease in perm_turbidity due to an equipment failure that had gone unrepaired over the break. Following a manual adjustment, the perm_turbidity increased to normal operating conditions. Because the method trains on OC data for perm_turbidity, the method flags IC perm_turbidity data as OC as seen in Fig. 5(b). Nevertheless, operators believe the OC behavior was isolated to perm_turbidity, so the FI results for other variables are still reliable and provide insight into the true cause of the fault.

In the MBR, a clog initially caused by hair, which is not completely removed by the 2 mm fine screen and is not digested by microorganisms, grew in the RAS pump, and this is detected by ADS-PCA late in the evening on January 4. The clog caused a slow decrease in the MBR influent flow, resulting in an initially intermittent shift and transitioning to a persistent downward shift detected in mbr_1_inf_flow and mbr_2_inf_flow. Additionally, the low strength of the wastewater due to the absence of students over the holidays led to a decline in biomass, which resulted in a decrease in mbr_1_tss, mbr_2_tss, and ras_tss. The downward shift in perm_conductivity detected at the beginning of January 5 was a result of a sensor fault.

The changes in the influent flow caused by the clog were very gradual and occurred over the course of multiple days, so this fault was a long, slow drift. As a result, the method only intermittently detects the downward shift in influent flow on January 4 and 5, but by January 6, the shift is consistently flagged. In comparison, in the absence of a real-time FI method, the clog went undiagnosed by operators for almost three weeks and was finally corrected on January 25. Therefore, this method would have identified the fault in the influent flow almost three weeks sooner than the operators did, which could have improved system performance and stability. The detection of the downward shift in TSS would have also pointed operators to a change in environmental pressures, which would have allowed them to alter the operating conditions to buffer against the decline.

3.5. Case study #3

The training window for case study #3 is January 23–27, 2018, followed by a two and a half day monitoring window. The fault in the MBR discussed in case study #2 was not corrected until January 25, so the training window for the MBR is not truly IC. Therefore, we focus solely on the BR, which is where the fault occurred for case study #3.

On January 29, a clog occurred in the raw wastewater diversion feeding the holding tank (which feeds the SB-MBR), causing the sewage_level to drop. This exposed the submerged pump responsible for feeding the SB-MBR, causing it to overheat and shut down. As a result,

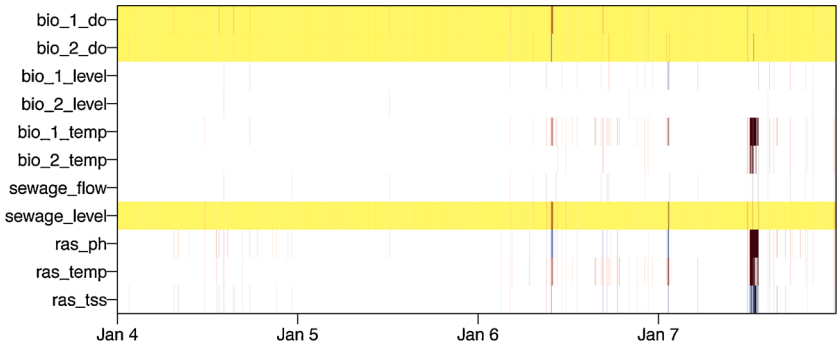

(a) BR subsystem

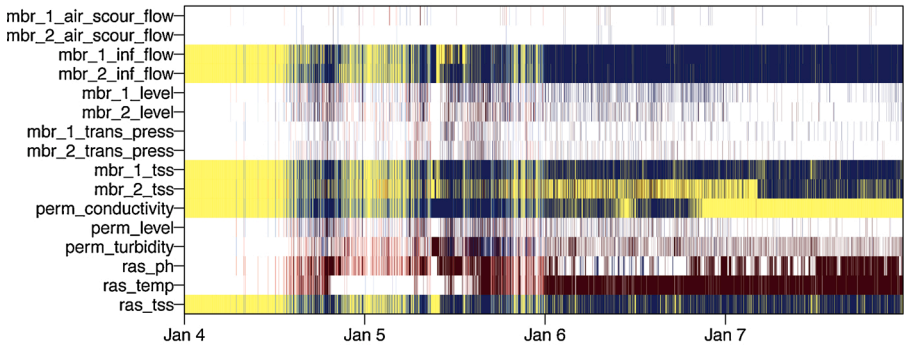

(b) MBR subsystem

Fig. 5. Checkerboard plot of monitored variables over time for case study #2, where the rows represent the monitored variables; red boxes represent a variable with an upward shift detected; blue boxes represent a variable with a downward shift detected; white boxes represent a variable classified as unshifted; and truly shifted variables are highlighted in yellow. The color of the red and blue boxes gets darker as the number of consecutive alarms increases. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

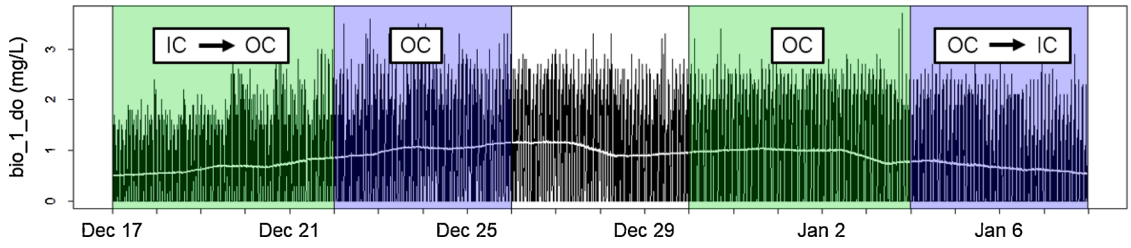

Fig. 6. Time series plot of bio_1_do for case studies #1 and #2. The two green regions are the initial training windows, and the two blue regions are the monitoring periods for case study #1 (left) and case study #2 (right). The white line is a one-day rolling average of bio_1_do. Each region is labeled as IC or OC based on whether we believe the process to be truly IC or OC at the time. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

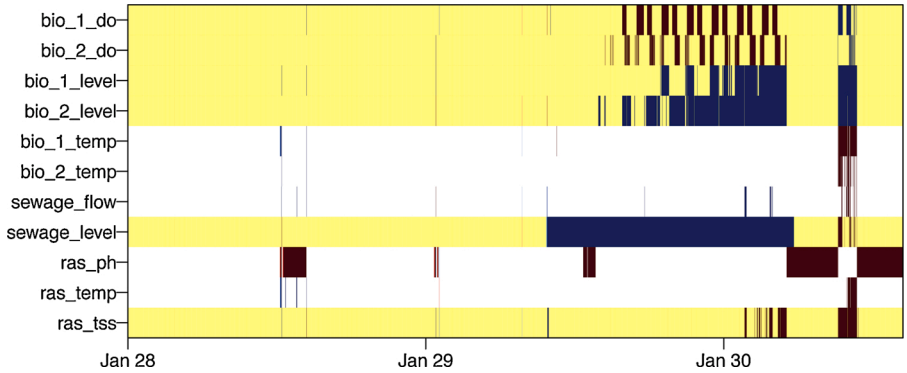

Fig. 7. Checkerboard plot of monitored variables over time for case study #3 for the BR subsystem, where the rows represent the monitored variables; red boxes represent a variable with an upward shift detected; blue boxes represent a variable with a downward shift detected; white boxes represent a variable classified as unshifted; and truly shifted variables are highlighted in yellow. The color of the red and blue boxes gets darker as the number of consecutive alarms increases. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

raw wastewater was not delivered to the BRs, causing the BR levels to decrease, which is evident in the downward shifts detected in bio_1_level and bio_2_level (see Fig. 7 mid-day on January 29). Without additional wastewater, there was no carbon or nutrients present to oxidize or provide microbial growth. Consequently, similar to case study #1, the DO concentrations immediately reached the setpoint without any readily oxidizeable material to consume the DO provided by the air blowers, which is consistent with a shift up in bio_1_do and bio_2_do detected by ADS-PCA. Operators were alerted to the problem when bio_1_level and bio_2_level dropped below their pre-specified threshold. Had the fault been detected sooner, the pump may not have been exposed, which would have prevented permanent equipment damage. During the fault event, no wastewater could be treated. For a facility with limited holding capacity, this would have been detrimental and would have led to untreated water being discharged to the environment. This fault could have easily been avoided with the installation of a more stringent level or level decline rate alarm, but the sensor was not accurately calibrated at the time of the fault, and the thresholds were not trusted by operators. Given the design of the system at the time of the fault, these results illustrate the power of a data-driven FD method to inexpensively complement existing WWT control systems and prevent severe consequences of simple faults.

4. Discussion

When incorporating a FD&I monitoring scheme into a WWT control system, it is important to remember that MSPM methods are purely data-driven. Even if the method performs as expected, it is possible for the results to be misleading or inconsistent with an operator's assessment of the fault. For example, the temperature variables were frequently flagged as shifted even when, upon further inspection by operators, the shifts were not substantial or were of no consequence. These false alarms were a result of a poor fit for the trend in Eq. (1) that resulted in non-stationarity in the residuals, leading to a shift detected in the mean when the temperatures increased or decreased. Secondly, ADS-PCA is very sensitive to the choice of the training data set. As highlighted in case study #2, if a data set with OC behavior is used to train the model, the monitoring scheme may flag IC data as OC. Therefore, it is imperative that operators and system experts thoroughly analyze the training data set prior to monitoring. Additionally, it is advisable to reset the monitoring scheme after a fault is detected, diagnosed, and corrected to ensure that the training window is not contaminated with observations from the OC period that were incorrectly classified as IC, which is especially common for long, slow drift faults.

In spite of some of the inherent limitations of ADS-PCA, it is able to account for changes in operating conditions by removing the effect of control settings, influent changes, and external factors on the monitored variables via the adaptive lasso linear regression. The general framework used for detrending makes this method very flexible for a variety of system configurations. Because lasso regression performs variable selection and parameter estimation simultaneously, it allows the user to include a multitude of predictor variables in the model without extensive model development or prior knowledge of the relationship between each response and predictor variable. This method can also accommodate stationary data by simply monitoring the standardized response variables rather than constructing a set of linear models and monitoring the residuals.

ADS-PCA is the first fully functional FD&I approach that can be used for WWT processes in real-time, and the approach can also be applied in other complex settings where nonstationarity, temporal dependence, and non-normality are present in the data. This method is particularly important for facilities of communities with limited means that might have outdated control systems. Instead of inspecting time series plots for each individual variable, which can lead to missed detection and delays in fault diagnosis, checkerboard plots allow operators to monitor all variables in the system simultaneously. They can be a very useful tool for operators to gain insight into the particular variables that are shifted. We demonstrate the utility of our approach in three complex fault settings at a decentralized WWT facility by comparing the operators' post-hoc analysis of the faults (along with their response in real-time) with the insight provided by ADS-PCA and the checkerboard plots, which confirms that our approach would have facilitated a more rapid detection and diagnosis had it been implemented in real-time. Without our fully integrated approach, operators are limited by the knowledge that a single variable is outside of the operator-determined thresholds that distinguish IC verses OC conditions. ADS-PCA and the checkerboard plots are intended to complement the expertise of WWT operators by synthesizing the wealth of data collected in WWT systems; narrowing their focus to specific components of the system that require attention; and facilitating more rapid fault diagnosis in a multivariate system.

Authors' contribution

Molly C. Klanderman: methodology, software, formal analysis, writing – original draft. Kathryn B. Newhart: data curation, writing – review & editing. Tzahi Y. Cath: conceptualization, writing – review & editing, funding acquisition. Amanda S. Hering: conceptualization, methodology, writing - review & editing, supervision, funding acquisition.

Acknowledgments

This work is supported by the National Science Foundation PFI:BIC Award No: 1632227; the National Science Foundation Engineering Research Center program under cooperative agreement EEC-1028968 (ReNUWIt); and a grant from the Colorado Higher Education Competitive Research Authority (CHECRA). We also acknowledge Rudy Maltos, Mike Veras, and Aqua-Aerobic Systems, Inc. for valuable contributions to this work.

Appendix A. Supplementary data

Supplementary data associated with this article can be found, in the online version, at https://doi.org/10.1016/j.jwpe.2020.101556.

References

[1] D. Gerrity, B. Pecson, R.S. Trussell, R.R. Trussell, Potable reuse treatment trains throughout the world, J. Water Supply Res. T. 62 (2013) 321–338.
[2] A.L. Prieto, D. Vuono, R. Holloway, J. Benecke, J. Henkel, T.Y. Cath, T. Reid, L. Johnson, J.E. Drewes, Decentralized wastewater treatment for distributed water reclamation and reuse: The good, the bad, and the ugly-Experience from a case study. Novel Solutions to Water Pollution, ACS Publications, 2013, pp. 251–266.
[3] D. Vuono, J. Henkel, J. Benecke, T.Y. Cath, T. Reid, L. Johnson, J.E. Drewes, Flexible hybrid membrane treatment systems for tailored nutrient management: a new paradigm in urban wastewater treatment, J. Membr. Sci. 446 (2013) 34–41.
[4] K.B. Newhart, R.W. Holloway, A.S. Hering, T.Y. Cath, Data-driven performance analyses of wastewater treatment plants: a review, Water Res. 157 (2019) 498–513.
[5] W.H. Woodall, D.C. Montgomery, Some current directions in the theory and application of statistical process monitoring, J. Qual. Technol. 46 (2014) 78–94.
[6] M.S. Reis, G. Gins, Industrial process monitoring in the big data/industry 4.0 era: from detection, to diagnosis, to prognosis, Processes 5 (2017) 1–16.
[7] R.B. Crosier, Multivariate generalizations of cumulative sum quality-control schemes, Technometrics 30 (1988) 291–303.
[8] H. Hotelling, The generalization of student's ratio, Ann. Math. Stat. 2 (1931) 360–378.
[9] C.A. Lowry, W.H. Woodall, C.W. Champ, S.E. Rigdon, A multivariate exponentially weighted moving average control chart, Technometrics 34 (1992) 46–53.
[10] A. Tang, P. Castagliola, J. Sun, X. Hu, An adaptive exponentially weighted moving average chart for the mean with variable sampling intervals, Qual. Reliab. Eng. Int. 33 (2017) 2023–2034.
[11] G.M. Abdella, K.N. Al-Khalifa, S. Kim, M.K. Jeong, E.A. Elsayed, A.M. Hamouda, Variable selection-based multivariate cumulative sum control chart, Qual. Reliab. Eng. Int. 33 (2017) 565–578.
[12] W. Liang, D. Xiang, X. Pu, A robust multivariate EWMA control chart for detecting sparse mean shifts, J. Qual. Technol. 48 (2016) 265–283.

[13] W. Li, X. Pu, F. Tsung, D. Xiang, A robust self-starting spatial rank multivariate EWMA chart based on forward variable selection, Comput. Ind. Eng. 103 (2017) 116–130.
[14] Y.K. Chen, K.L. Hsieh, Hotelling's T^2 charts with variable sample size and control limit, Eur. J. Oper. Res. 182 (2007) 1251–1262.
[15] Y. Dai, Y. Luo, Z. Li, Z. Wang, A new adaptive CUSUM control chart for detecting the multivariate process mean, Qual. Reliab. Eng. Int. 27 (2011) 877–884.
[16] S.J. Qin, Survey on data-driven industrial process monitoring and diagnosis, Annu. Rev. Control 36 (2012) 220–234.
[17] C.F. Alcala, S.J. Qin, Reconstruction-based contribution for process monitoring, Automatica 45 (2009) 1593–1600.
[18] C.F. Alcala, S.J. Qin, Reconstruction-based contribution for process monitoring with kernel principal component analysis, Ind. Eng. Chem. Res. 49 (2010) 7849–7857.
[19] Z. Yan, Y. Yao, Variable selection method for fault isolation using least absolute shrinkage and selection operator (LASSO), Chemometr. Intell. Lab. 146 (2015) 136–146.
[20] K. Wang, W. Jiang, High-dimensional process monitoring and fault isolation via variable selection, J. Qual. Technol. 41 (2009) 247–258.
[21] W. Jiang, K. Wang, F. Tsung, A variable-selection-based multivariate EWMA chart for process monitoring and diagnosis, J. Qual. Technol. 44 (2012) 209–230.
[22] K. Ghosh, M. Ramteke, R. Srinivasan, Optimal variable selection for effective statistical process monitoring, Comput. Chem. Eng. 60 (2014) 260–276.
[23] F. de Assis Boldt, T.W. Rauber, F.M. Varejão, Cascade feature selection and ELM for automatic fault diagnosis of the Tennessee Eastman process, Neurocomputing 239 (2017) 238–248.
[24] K. Wang, F. Tsung, An adaptive dimension reduction scheme for monitoring feedback-controlled processes, Qual. Reliab. Eng. Int. 25 (2009) 283–298.
[25] K. Nishimura, S. Matsuura, H. Suzuki, Multivariate EWMA control chart based on a variable selection using AIC for multivariate statistical process monitoring, Stat. Probabil. Lett. 104 (2015) 7–13.
[26] C. Zou, W. Jiang, F. Tsung, A LASSO-based diagnostic framework for multivariate statistical process control, Technometrics 53 (2011) 297–309.
[27] G. Capizzi, G. Masarotto, A least angle regression control chart for multidimensional data, Technometrics 53 (2011) 285–296.
[28] S. Ebrahimi, C. Ranjan, K. Paynabar, Large Multistream Data Analytics for Monitoring and Diagnostics in Manufacturing Systems, 2018. arXiv:1812.10430.
[29] H. Yan, K. Paynabar, J. Shi, Real-time monitoring of high-dimensional functional data streams via spatio-temporal smooth sparse decomposition, Technometrics 60 (2018) 181–197.
[30] K. Kazor, R.W. Holloway, T.Y. Cath, A.S. Hering, Comparison of linear and nonlinear dimension reduction techniques for automated process monitoring of a decentralized wastewater treatment facility, Stoch. Environ. Res. Risk. A 30 (2016) 1527–1544.
[31] P. Krupskii, F. Harrou, A.S. Hering, Y. Sun, Copula-based monitoring schemes for non-Gaussian multivariate processes, J. Qual. Technol. 52 (2019) 219–234.
[32] D.S. Lee, P.A. Vanrolleghem, Adaptive consensus principal component analysis for on-line batch process monitoring, Environ. Monit. Assess. 92 (2004) 119–135.
[33] G.J. Odom, K.B. Newhart, T.Y. Cath, A.S. Hering, Multistate multivariate statistical process control, Appl. Stoch. Model Bus. 34 (2018) 880–892.
[34] F. Baggiani, S. Marsili-Libelli, Real-time fault detection and isolation in biological wastewater treatment plants, Water Sci. Technol. 60 (2009) 2949–2961.
[35] F. Harrou, A. Dairi, Y. Sun, M. Senouci, Statistical monitoring of a wastewater treatment plant: a case study, J. Environ. Manage. 223 (2018) 807–814.
[36] K. Gernaey, U. Jeppsson, P. Vanrolleghem, J.B. Copp, Benchmarking of Control Strategies for Wastewater Treatment Plants. Number 23 in Scientific and Technical Report Series, IWA Publishing, London, 2014.
[37] D. Aguado, C. Rosen, Multivariate statistical monitoring of continuous wastewater treatment plants, Eng. Appl. Artif. Intel. 21 (2008) 1080–1091.
[38] L. Corominas, K. Villez, D. Aguado, L. Rieger, C. Rosén, P.A. Vanrolleghem, Performance evaluation of fault detection methods for wastewater treatment processes, Biotechnol. Bioeng. 108 (2011) 333–344.
[39] A. Sánchez-Fernández, F.J. Baldán, G.I. Sainz-Palmero, J.M. Benítez, M.J. Fuente, Fault detection based on time series modeling and multivariate statistical process control, Chemometr. Intell. Lab. 182 (2018) 57–69.
[40] Y. Liu, B. Liu, X. Zhao, M. Xie, A mixture of variational canonical correlation analysis for nonlinear and quality-relevant process monitoring, IEEE Trans. Ind. Electron. 65 (2018) 6478–6486.
[41] J.M. Lee, C. Yoo, S.W. Choi, P.A. Vanrolleghem, I.B. Lee, Nonlinear process monitoring using kernel principal component analysis, Chem. Eng. Sci. 59 (2004) 223–234.
[42] M.C. Klanderman, K.B. Newhart, T.Y. Cath, A.S. Hering, Fault isolation for a complex decentralized wastewater treatment facility, J. R. Stat. Soc. C Appl. 69 (2020) 931–951.
[43] P. Qiu, D. Xiang, Univariate dynamic screening system: an approach for identifying individuals with irregular longitudinal behavior, Technometrics 56 (2014) 248–260.
[44] P. Qiu, D. Xiang, Surveillance of cardiovascular diseases using a multivariate dynamic screening system, Stat. Med. 34 (2015) 2204–2221.
[45] J. Li, P. Qiu, Nonparametric dynamic screening system for monitoring correlated longitudinal data, IIE Trans. 48 (2016) 772–786.
[46] J. Li, P. Qiu, Construction of an efficient multivariate dynamic screening system: construction of an efficient multivariate dynamic screening system, Qual. Reliab. Eng. Int. 33 (2017) 1969–1981.
[47] L.A. Jones-Farmer, W.H. Woodall, S.H. Steiner, C.W. Champ, An overview of Phase I analysis for process improvement and monitoring, J. Qual. Technol. 46 (2014) 265–280.
[48] G. Capizzi, Recent advances in process monitoring: nonparametric and variable-selection methods for Phase I and Phase II, Qual. Eng. 27 (2015) 44–67.
[49] S. Chakraborti, S.W. Human, M.A. Graham, Phase I statistical process control charts: an overview and some results, Qual. Eng. 21 (2009).
[50] X. Wang, U. Kruger, G.W. Irwin, Process monitoring approach using fast moving window PCA, Ind. Eng. Chem. Res. 44 (2005) 5691–5702.
[51] N. Ayech, C. Chakour, M.F. Harkat, New adaptive moving window PCA for process monitoring, IFAC Proc. 45 (2012) 606–611.
[52] W. Ku, R.H. Storer, C. Georgakis, Disturbance detection and isolation by dynamic principal component analysis, Chemometr. Intell. Lab. 30 (1995) 179–196.
[53] S. Chakraborti, M.A. Graham, Nonparametric (distribution-free) control charts: an updated overview and some results, Qual. Eng. 31 (2019) 523–544.
[54] J. Friedman, T. Hastie, R. Tibshirani, Regularization paths for generalized linear models via coordinate descent, J. Stat. Softw. 33 (2010) 1–22.
[55] H. Zou, The adaptive lasso and its oracle properties, J. Am. Stat. Assoc. 101 (2006) 1418–1429.
[56] Z. Yan, Y. Yao, T.B. Huang, Y.S. Wong, Reconstruction-based multivariate process fault isolation using Bayesian lasso, Ind. Eng. Chem. Res. 57 (2018) 9779–9787.