Near real-time detection of blockages in the proximity of combined sewer overflows using evolutionary ANNs and statistical process control

Rosin, T. R.; Kapelan, Z.; Keedwell, E.; Romano, M.

DOI
10.2166/hydro.2022.036

Publication date
2022

Document Version
Final published version

Published in
Journal of Hydroinformatics

Citation (APA)
Rosin, T. R., Kapelan, Z., Keedwell, E., & Romano, M. (2022). Near real-time detection of blockages in the proximity of combined sewer overflows using evolutionary ANNs and statistical process control. *Journal of Hydroinformatics*, *24*(2), 259-273. https://doi.org/10.2166/hydro.2022.036

Journal of
Hydroinformatics

 Journal of Hydroinformatics Vol 24 No 2, 259 doi: 10.2166/hydro.2022.036

Near real-time detection of blockages in the proximity of combined sewer overflows using evolutionary ANNs and statistical process control

T. R. Rosin [a,*], Z. Kapelan [a,b], E. Keedwell[a] and M. Romano[c]

[a] Centre for Water Systems, University of Exeter, Exeter EX4 4QF, UK
[b] Department of Water Management, Delft University of Technology, Stevinweg 1, 2628 CN Delft, The Netherlands
[c] United Utilities Plc, Lingley Mere Business Park, Warrington WA5 3LP, UK
*Corresponding author. E-mail: talia.rosin1@gmail.com

TRR, 0000-0002-8956-3058; ZK, 0000-0002-0934-4470

ABSTRACT

Blockages are a major issue for wastewater utilities around the world, causing loss of service, environmental pollution, and significant clean-up costs. Increasing telemetry in combined sewer overflows (CSOs) provides the opportunity for near real-time data-driven modelling of wastewater networks. This paper presents a novel methodology, designed to detect blockages and other unusual events in the proximity of CSO chambers in near real-time. The methodology utilises an evolutionary artificial neural network (EANN) model for short-term CSO level predictions and statistical process control (SPC) techniques to analyse unusual level behaviour. The methodology was evaluated on historic blockage events from several CSOs in the UK and was demonstrated to detect blockage events quickly and reliably, with a low number of false alarms.

Key words: blockage detection, combined sewer overflow, evolutionary artificial neural network, radar rainfall nowcasts, statistical process control

HIGHLIGHTS

- Development of a novel real-time data-driven blockage detection system.
- Two blockage detection techniques: EANN discrepancy-based analysis and statistical trend-based analysis.
- Utilising only CSO level data, rainfall data and rainfall nowcasts.
- Testing on real case studies demonstrates timely and reliable performance.

1. INTRODUCTION

Blockages are responsible for a significant proportion of the failures which occur in wastewater networks (Jin & Mukherjee 2010). Blockages cause numerous issues such as unplanned maintenance, internal and external flooding, pollution, costly clean-up and compensation costs, and threats to public health due to water-borne pathogens (Rodríguez *et al.* 2012). In the UK alone, there are approximately 300,000 blockages every year, resulting in costs of £100 million (Water UK 2017). Ofwat, the regulatory agency in England and Wales, imposes a statutory duty on wastewater utilities to minimise the frequency of blockage events. However, with deteriorating sewer networks and increased water efficiency, the number of sewer blockages occurring in public sewers is increasing. The timely prediction of blockage events, therefore, plays an important role in the management of urban water systems.

Historically, wastewater utilities have relied on their customers to report blockage events, employing a reactive repair and maintenance approach. However, this causes increased loss of service and customer complaints, in turn affecting regulatory performance. Additionally, if the blockage does not cause any obvious external effects (e.g. flooding), it can remain undetected for many months, increasing the likelihood of combined sewer overflows (CSOs) into nearby watercourses. There are also several hardware-based techniques available to detect blockages, such as CCTV, acoustic techniques and laser profiling. However, these are generally expensive, intrusive and time-consuming to implement, and require the added expense of a human operator.

As part of the Event Duration Monitoring (EDM) programme implemented by the Environment Agency, wastewater utilities have been required to monitor sewer levels at the majority of CSOs in England and Wales since 2020. Consequently, large quantities of increasingly accurate level sensors have been installed in the sewer network. In conjunction with reductions in data storage costs and improved computer processing power, this has led to water utilities routinely collecting large volumes of accurate sewer level data in near real-time. This data is extremely valuable; however, it is not feasible for human operators to process this raw data manually. This motivates the need for intelligent data-driven technologies to assist in real-time modelling and management of the wastewater system, able to monitor the network and detect blockages as soon as they occur, thus allowing utilities to proactively remove obstructions before the customer and environment are affected.

2. BACKGROUND

Much of the recent literature published on sewer blockage management has focused on predicting blockage rates and representing the spatial distribution of sewer failures (Ugarelli & Taylor 2016) and flood hazards (Cherqui *et al.* 2015) using statistical methods. These methods are valuable in improving the planning and scheduling of sewer asset inspection, maintenance and replacement programmes, so that blockages are less likely to form. However, they are not capable of detecting actual blockage events in real-time.

Regarding real-time systems, there is little present in the literature. A small number of methodologies have recently been developed to automatically detect blockages and other sewer defects from CCTV data and can potentially be applied in real-time (e.g. Myrans *et al.* 2019). However, CCTV is generally limited to regular inspections of sections of the sewer network, in comparison to continuous real-time network-wide monitoring, which has the potential to identify blockages throughout the sewer system as soon as they occur and so prevent possible failures.

There are some recent commercially available products and services designed for sewer blockage detection. For example, the company Detectronic offers a service to predict blockages and provide early warnings of network failure by monitoring sewer data in real-time using human data analysts (Detectronic 2018, 2019). The system develops normalised data profiles for all monitored sites and uses predictive criteria for early preventative intervention. The system appears to work well, but it relies heavily on human input. Another example is the SMART Sewer™ that detects developing blockages autonomously in real-time using level data (EMS 2020). The system utilises fuzzy logic and is designed to be deployed in high incidence and high-risk areas. This system, however, seems to require a higher density of level sensors than is currently available in the UK and in most other countries. As both systems are commercially available products, there is little literature available describing the systems' methodologies or showing their performances in detail. These commercial solutions have been applied to case studies by wastewater utilities; Southern Water, for example, trialled the SMART sewer in hotspot areas in Brighton and East Worthing (Southern Water 2018). Such new technology has not yet become an industry standard. However, they demonstrate that utilities see the potential in such systems and are willing to deploy them in the future.

In addition, the area of fault detection in water distribution systems (WDS) exhibits many similarities. Although blockages like those found in sewer pipes are not common, many other abnormal events occur, including pipe bursts, leaks and equipment failures. Recently, computational intelligence techniques have gained prominence as tools for leakage and event detection using flow and pressure data (e.g. Mounce & Machell 2006; Mounce *et al.* 2010; Romano *et al.* 2014; Choi *et al.* 2016). The approach presented here builds on the useful information and insights gained during the development of these technologies for WDSs.

This paper outlines the development of a novel event detection system (EDS), designed to automatically and continually monitor the wastewater network for blockages and other unusual events in the proximity of CSOs in near real-time. The proposed system applies two different detection techniques: evolutionary artificial neural network (EANN) discrepancy-based analysis and statistical trend-based analysis. The system is designed to make use of level sensors already installed in the average UK wastewater utility and does not require the installation of any additional monitors. It is envisioned that the methodology will be beneficial to utilities, allowing proactive management of blockages before they affect the customer and the environment.

3. METHODOLOGY

3.1. Event detection system overview

During normal functioning of the wastewater network, the water level in a CSO exhibits predictable behaviour – a repeatable diurnal pattern during dry weather and an increase in level during rainfall events, resulting in an overflow if the sewer level exceeds the spill height of the chamber. However, when a blockage, or another abnormal event, occurs in the proximity of a CSO, this can cause a deviation from normal behaviour. Generally, this consists of a level increase if the blockage is downstream of the chamber and a decrease if the blockage is upstream. It is therefore possible to detect blockages by identifying this abnormal behaviour in measured CSO level data.

The EDS operates by monitoring CSO levels in real-time and identifying abnormalities using statistical and machine learning techniques to analyse the CSO data, with the aim of detecting blockages in a timely and reliable manner. Figure 1 provides a visual representation of the EDS methodology. The framework consists of three subsystems: (1) data retrieval and pre-processing, (2) event detection and (3) event inference.

The first subsystem is designed to retrieve the measured CSO level and radar rainfall data. The data is then pre-processed to construct a clean and uniform time series.

The second subsystem is designed to analyse the incoming processed data and generate evidence that a blockage has occurred. Two separate methodologies encapsulated in respective modules are employed here: the first EANN discrepancy-based analysis module is designed to monitor the discrepancy between EANN CSO level predictions and the measured CSO levels. The second statistical trend-based analysis module uses statistical analysis to monitor trends in the CSO level data. The aim of using a combination of statistical and AI techniques here is to collect more rigorous evidence that a blockage has occurred. Different types of blockage events have different effects on sewer levels. For example, sedimentation and build-up of fats, oils and grease (FOG) cause gradual level changes, whilst sewer collapse and snagging of objects cause sudden fluctuations. The application of different types of detection techniques increases the likelihood that all types of blockage events are detected quickly and reliably.

The third and final subsystem consists of an inference engine. The engine analyses the information from the above two evidence generation modules to determine if there is sufficient evidence to determine that a blockage event has occurred, and thus generate an alarm and notify the wastewater utility. Additional information concerning the blockage is also provided by

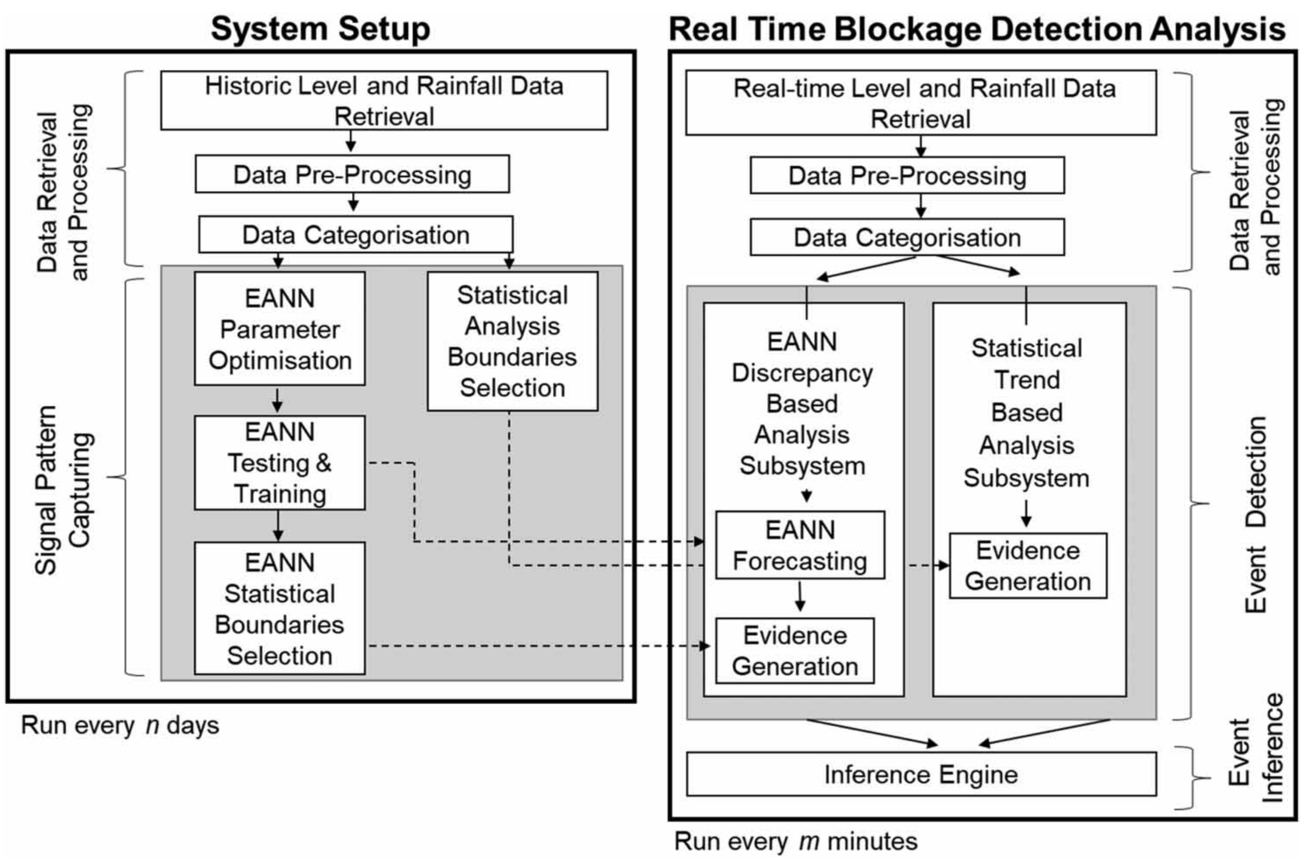

Figure 1 | Diagrammatic representation of the EDS methodology (left figure shows the offline system set-up and the right one shows the online application in near real-time).

the system to facilitate the user in deciding how to respond to the event. Each of the three steps will be discussed in detail in the following subsections.

The EDS has two modes of operation: a set-up (training, offline) mode and a real-time (deployment, online) mode (which consists of the three subsystems described above), as shown in Figure 1. During the set-up mode, the system 'learns' the behaviour of the CSO being analysed, using historical data to capture normal CSO levels. It is advised that this mode is employed periodically, and also following any significant sewer system configuration changes. This automatic retraining is necessary as the behaviour of the system may change over time, e.g. due to planned changes to the sewer network and wastewater treatment plants or the construction of new buildings connected to the network. This is a well-known problem known as concept drift (Widmer & Kubat 1996). The EDS must be able to adapt to this new definition of 'normal' behaviour in an unsupervised, automated fashion.

The second mode of operation is the real-time mode. This is the normal online operating mode used to detect events and raise alarms in near real-time. The system is designed to operate during this mode at the communication interval of the incoming level and rainfall data received by the system, e.g. operate at 15-min intervals in the case studies presented here. The system was coded and run in a MATLAB software environment. Further information is available upon request to the authors of this paper.

3.2. Data pre-processing and categorisation

The first subsystem of the detection process consists of data retrieval, pre-processing and categorisation. Both rainfall and CSO level data may contain bad or missing data points. Data pre-processing is performed to construct a clean, uniform dataset, i.e. with any anomalies removed and with regularly spaced timesteps. It is especially important that erroneous CSO level data is removed, as it could be mistaken by the system as abnormal behaviour caused by a blockage and therefore trigger a false alarm.

The first step involves identifying and correcting any erroneous data points. For both rainfall and level data, any negative values are removed. For level data only, values above the height of the level monitor are removed. Additionally, erroneous level data caused by 'benching' are identified and removed. Benching refers to bad data points caused by the level monitor recording the height of objects in the CSO chamber such as the side of the CSO or the bracket holding the monitor, rather than the water level. Not only does benching significantly affect data quality, benching which occurs above the spill level of the chamber can cause the CSO to mistakenly appear to be spilling frequently out of consent. The occurrence of benching is relatively common; an analysis of 100 monitors installed in CSO chambers located in North West England found that 36 suffered from some degree of benching.

There are numerous methods available for data cleansing; however, none are designed for the unique profile exhibited by benching. Thus, a methodology was developed specifically to identify benching in both historic and real-time CSO level data. Benching is characterised by sudden sharp peaks, unrelated to rainfall or the usual diurnal CSO pattern. These peaks can be distinguished from genuine increases in level caused by rainfall or unusual events such as blockages as these conditions generally cause level changes which occur progressively over multiple timesteps. Additionally, as benching is caused by the presence of objects in the CSO chamber, the incorrect data points largely occur consistently at the same height (within a few millimetres of tolerance). In addition, benching may also manifest as consecutive points of the same value, i.e. flat peaks. As this signature does not occur naturally in genuine level data, these points are unambiguously incorrect. The benching identification methodology detects local maxima in the level data occurring over single timesteps, as well as flat peaks. The identified peaks are analysed using a frequency histogram to determine if the rate of peaks at a certain height indicates the presence of benching in the CSO chamber.

Once any erroneous data has been removed, missing level and rainfall data are then infilled using linear interpolation and a uniform time series is constructed. Linear interpolation was selected here as it is a simple method that has been used previously to infill missing sensor data (Park & Kim 2020). More advanced techniques are available in the literature (Sattari *et al.* 2017). However, linear interpolation was deemed adequate as missing data points, especially consecutive missing points, were rare in the data utilised here.

The data is then categorised according to rainfall characteristics to construct a profile of normal sewer behaviour for different types of rainfall events. Sewer levels are highly dependent on the characteristics of the rainfall in the surrounding catchment, and different types of rainfall cause different sewer level dynamics. The categorisation process allows the detection methodology to be tailored to the various rainfall characteristics, thus improving the system performance, maximising the true positive rate, minimising the event detection time and decreasing the false-positive rate.

Rainfall varies considerably from one country to another, and as such no universal rainfall classifications exist. Rainfall is most commonly characterised using rainfall intensity–duration–frequency (IDF) curves (Koutsoyiannis *et al.* 1998). In this work, the data is categorised according to the rainfall duration and intensity only. Dividing the data into too many rainfall classes can result in categories with insufficient data to construct an adequate CSO level profile, whilst dividing into too few classes would result in ineffective detection. It was decided here to not use frequency (i.e. return period) as it resulted in too many categories.

An analysis was performed to determine the optimal number of categories which produced the best results when utilised by the EDS (not shown here due to space limitations). Based on this analysis, the rainfall is divided in this work into six principal categories described as follows:

1. dry weather (no rainfall),
2. short-duration, low-intensity rainfall (type A),
3. short-duration, high-intensity rainfall (type B),
4. long-duration, low-intensity rainfall (type C),
5. long-duration, high-intensity rainfall (type D) and
6. post-event period (the period immediately following a rainfall event).

The post-event period is designed to capture the period immediately following a rainfall event when the water level in the sewer remains high. This is due to the time it takes for water to flow from the surrounding catchment to the CSO chamber. An additional analysis was performed to identify the optimal duration and intensity threshold parameters for these categories. The six categories are designed to be generic and as such can be applied automatically to any CSO site, which was done in the approach used here.

3.3. Statistical trend-based evidence generation module

The statistical trend-based analysis module is the first module designed to identify unusual CSO level behaviour and identify evidence of a blockage. The module operates by determining if the incoming level data lies inside a data envelope of 'normal behaviour' defined by statistical analysis of the CSO system.

During the set-up of the EDS, historic CSO level and rainfall data during normal operations (i.e. time periods containing no blockages) are used to calculate the expected levels in the CSO chamber for each of the six rainfall categories. During real-time running of the system, the most recent pre-processed and categorised data are retrieved from the time-series database for a time window of d timesteps. A set of modified Western Electric (WE) control charts (Western Electric Company 1958) are then used to monitor the window of CSO level data to determine if any deviations from the normal operating behaviour are significant enough to provide evidence of a blockage.

Control charts are one of the most prominent statistical process control (SPC) techniques. The charts utilise upper and lower control limits to determine if, statistically, a process is behaving as expected or is 'out-of-control' and corrective action should be taken. The limits of the control chart are typically drawn at three standard deviations from the mean. Standard control charts are efficient at detecting large, fast deviations from the process average; however, they are generally insensitive to small changes. Therefore, other charts have been developed which detect small changes efficiently by analysing information from observations prior to the most recent timestep. The most widely known charts are the WE control rules which apply multiple run rules (Western Electric Company 1958). A set of modified WE rules were applied here, designed to minimise the number of false alarms generated, whilst still maintaining a true positive rate.

For each run rule i, the control chart limit L at timestep t is defined as follows:

If dry weather $L_{i,\ dry,w,t} = \mu_{dry,w,t} + M_i * N_{dry} * \sigma_{dry,w,t}$
(i.e. rainfall category (1))

If wet weather $L_{i,c} = \mu_c + M_i * N_c * \sigma_c$
(i.e. rainfall categories (2) – (6))

where μ_c and σ_c are the mean and standard deviation of the historic CSO level for the current rainfall category c, $\mu_{dry,w,t}$ and $\sigma_{dry,w,t}$ are the mean and standard deviation of the historic CSO level for the dry weather rainfall category at timestep t, w

denotes if the timestep is a weekday or weekend, M_i is a constant multiplier defined for each run rule i (e.g. 3 in the standard 3-sigma control chart), and N_c is an additional multiplier defined for each rainfall category c.

The various WE run rules i are shown in Table 1(a) and the run rules and corresponding parameter N selected for each rainfall category are presented in Table 1(b) (e.g. during dry weather periods, run rules 2, 3 and 4 are utilised, with N equal to 3). Rules 1–4 consist of the original WE rules. Rules 5 and 6 were added to accommodate gradually forming blockages, which cause CSO levels to increase slowly over a long period of time. Analysis indicated that these events were frequently overlooked by the original rules as they require a large deviation from normal levels over a small number of timesteps to be satisfied.

The values *of* μ, σ and N are determined separately for each rainfall category so that the control limits are tailored to the rainfall characteristics of the current timestep, therefore increasing the likelihood any blockage events are detected quickly and reducing false alarms. The limits are defined separately for wet and dry weather, as during dry weather the sewer level is highly dependent on the time of day and the day of the week (during dry weather, the level exhibits a strong diurnal pattern based on the water consumed by households). By calculating the dry weather μ and σ for each timestep of the day and for weekends and weekdays, the limits are adapted to the time-varying behaviour of the water level. During wet weather, the level is dependent primarily on rainfall only and the diurnal pattern is not significant. It is therefore not necessary to consider the time of day or the day of the week when calculating the wet weather limits.

There is no universal definition to categorise dry and wet weather. A binary threshold is used here to classify the data based on cumulative rainfall over a number of past timesteps. The threshold is defined as follows:

$$Z_i = \sum_{t=0}^{n} R_{i-t}$$

$$\text{Class} = \begin{cases} \text{Wet} & \text{if } Z_i > \theta \\ \text{Dry} & \text{otherwise} \end{cases}$$

where R is rainfall intensity, θ is the wet weather threshold and n is the number of past timesteps considered. θ and n are here defined as 0.5 mm and 10, respectively. These values were determined by analysing historical CSO level and rainfall data and identifying for which values rainfall during dry weather had a negligible effect on CSO level.

Table 1 | (a) WE rules for Shewhart control charts and (b) selected rules and parameters for each rainfall category for the statistical analysis module

(a)

Run rule *i*	**Description**	**Constant multiplier M_i**
1	1 out of 1 consecutive discrepancies fall outside the defined control limits	5
2	2 out of 3 consecutive discrepancies fall outside the defined control limits	4
3	4 out of 5 consecutive discrepancies fall outside the defined control limits	3
4	8 out of 8 consecutive discrepancies fall outside the defined control limits	2
5	15 out of 15 consecutive discrepancies fall outside the defined control limits	1
6	25 out of 25 consecutive discrepancies fall outside the defined control limits	0.5

(b)

Rainfall category *c*	**Selected run rules**	***N***
Dry weather	2–4	3
Rainfall type A	2–4	4
Rainfall type B	2–4	4
Rainfall type C	2–4	5
Rainfall type D	3–5	6
Post event	2–5	4

This definition appeared to work well for all the CSOs analysed in this catchment. However, it is possible that if the methodology were to be applied to other countries with different conditions, the definition should be revisited to ensure it still produces satisfactory results. However, note that the aim of the wet/dry threshold is just a means to partition the data in order to increase the effectiveness of the model. The dry and wet weather can then be accurately modelled using separate simple models, rather than using a single complex model. The goal of the definition is not to provide a method of accurately defining wet and dry weather. Therefore, the precise choice of the threshold value is not critical.

These rule sets and corresponding parameters were selected for each rainfall category via an extensive sensitivity analysis performed on 12 different CSO sites containing 12 sudden and gradual blockage events and located in both urban and rural catchments. These analyses included an ROC curve, the partial area under the curve (Ma *et al.* 2013), two methods to determine the optimal cut-off point of the ROC curve: the Youden index (Youden 1950) and a cost-based approach (Zweig & Campbell 1993), and an analysis of the blockage detection time. These tests were employed to select the optimal system sensitivity. There exists a trade-off when selecting system parameters between minimising the detection time and the false-negative rate, whilst also minimising the number of false alarms generated. An overly sensitive system results in wasting operational resources on investigating false alarms, whilst an insensitive system causes excessive losses due to missed events and long detection delays.

During real-time (i.e. online) running of the system, the rainfall category of the current timestep is identified and the appropriate run rules are applied. If the measured level falls outside the control chart limits and any of the rules are satisfied, evidence of a blockage is inferred. The selected parameters are designed to be generic, i.e. they can be applied to any CSO without further analysis as they are based only on the μ and σ of the CSO level data. Seasonal effects on the system are not directly addressed currently. However, the six rainfall classes mean the system is designed to accommodate different types of rainfall, and therefore adapt to increased rainfall in winter months.

3.4. EANN discrepancy-based evidence generation module

The EANN discrepancy-based module is the second module designed to identify blockage events. This module aims to detect blockages by analysing the CSO level forecasts produced by an EANN model. EANNs refer to a class of ANNs whereby evolutionary algorithms (EAs) are used in the ANN model design and/or training to automatically determine the optimum model parameters and architecture. The EANN model therefore has a greater prediction accuracy than an ANN with parameters selected via trial-and-error, and this increase in model performance when forecasting levels in a CSO chamber was demonstrated by Rosin *et al.* (2021). In addition, as an evolutionary strategy algorithm is employed to automatically select the optimal ANN structure and parameters, no significant human input is required to construct the model and the network is tailored specifically for the CSO location in question. Several ANN models have been demonstrated to successfully forecast CSO levels (e.g. Fernando *et al.* 2006; Mounce *et al.* 2014; Zhang *et al.* 2018; Rosin *et al.* 2021).

The EANN model utilised here was developed by Rosin *et al.* (2021). The model requires inputs of CSO level data, radar rainfall data and forecast rainfall data. The EANN is trained to forecast normal levels in the CSO chamber, assuming no blockage event. When a blockage occurs, the level in the sewer system exhibits abnormal behaviour and the EANN model is unable to produce accurate forecasts. As a result, the discrepancy between the model prediction and the measured level increases. By continually monitoring this discrepancy in real-time deviations from the norm can be identified and analysed to infer the presence of an event.

An example of the EANN model forecast discrepancy during a sudden and a gradual blockage is displayed in Figure 2. The sudden blockage causes a significant increase in the CSO level. Although the EANN model was able to forecast the change in level, it failed to capture the magnitude of the increase, resulting in a large discrepancy between the measured and forecast results. Gradually forming blockages are generally more difficult to detect as the change in level over time is smaller and the EANN model has more time to adapt to the change. However, a small, but still measurable, discrepancy between the forecast and measured level still occurs during the blockage which can be detected. As the control charts analyse the model discrepancy between EANN predictions and the corresponding observations using suitable discrepancy classes, the actual accuracy of the EANN predictions are not critical for the performance of the detection methodology.

Modified WE control charts tailored for each rainfall category continually monitor and assess the EANN discrepancy in real-time. The model discrepancy is defined here as $x_{\mathrm{EANN},t} - x_{\mathrm{obs},t}$, where $x_{\mathrm{obs},t}$ is the observed sensor level and $x_{\mathrm{EANN},t}$ is the EANN model prediction at time t.

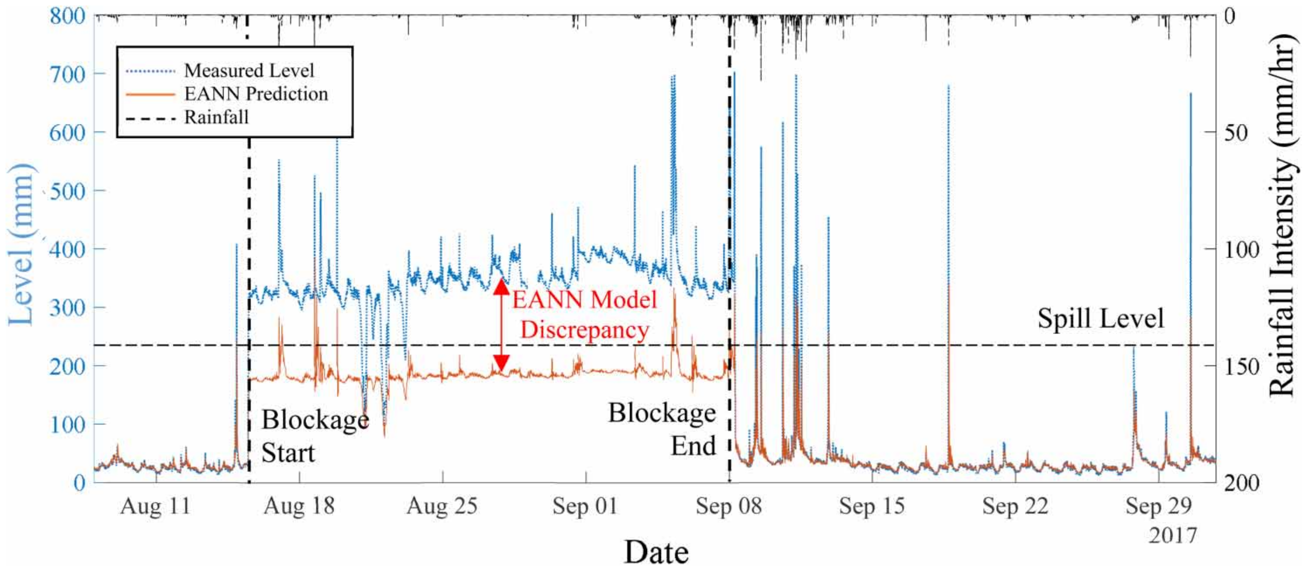

Figure 2 | Example of the EANN model-predicted CSO level (orange line) versus the measured level (blue line) before and after a sudden blockage event. Please refer to the online version of this paper to see this figure in colour: http://dx.doi.org/10.2166/hydro.2022.036.

The control chart limits $L_{i,c}$ for run rule i and rainfall category c are then defined as follows:

$$L_{i,c} = \mu_{\mathrm{EANN},c} + P_i * N_c * \sigma_{\mathrm{EANN},c}$$

where $\mu_{\mathrm{EANN},c}$ and $\sigma_{\mathrm{EANN},c}$ are the mean and standard deviation of the historic EANN discrepancy for rainfall category c, and P_i is a constant multiplier defined for each run rule i.

The accuracy of the EANN predictions are significantly influenced by rainfall intensity. The model performs extremely well during dry weather. During heavy rainfall, the CSO level is much more changeable and sudden in nature, and the model discrepancy increases (as shown in the Supplementary material). For more information on this, see Rosin *et al.* (2021). Therefore, calculating the μ and σ of the data separately for each rainfall category significantly improves the EDS performance.

The modified WE rules and corresponding parameters were selected using the same analyses utilised in the statistical analysis module, and are presented in Table 2.

3.5. Inference engine

The inference module constitutes the final processing stage of the EDS and enables the raising of alarms when a blockage is detected. When applied in near real-time (e.g. every 15 min), the EANN discrepancy-based analysis module and the statistical trend-based module described in previous sections are run in parallel, and the results are then fed to the inference module. At each timestep, the inference module analyses the incoming information from these two sub-systems and determines if there is enough evidence of a blockage event to raise an alarm.

The inference module is composed of a rule-based inference engine which applies pre-defined rules to the incoming data in order to deduce new information. A simple rule-based inference engine has been developed and applied here. More sophisticated inference methods are available, for example neural network-based engines, decision trees, and random forests. However, these generally require large amounts of event training data, which is not available here. Additionally, the advantage of the rule-based system is that as the rules are well defined and easily understandable, the engine can therefore be modified and extended by future users of the system who may not be experts in the field.

The inference module receives the evidence from the EANN discrepancy-based analysis module and the statistical trend-based module, and a simple set of If-Then rules are applied. If either module presents evidence at the current timestep that a blockage has occurred, a window of n past timesteps are retrieved from each module and analysed to determine if the total

Table 2 | (a) WE rules for Shewhart control charts and (b) selected rules and parameters for each rainfall category for the EANN discrepancy module

(a)

Run rule i	Description	Constant multiplier M_i
1	1 out of 1 consecutive discrepancies fall outside the defined control limits	5
2	2 out of 3 consecutive discrepancies fall outside the defined control limits	4
3	4 out of 5 consecutive discrepancies fall outside the defined control limits	3
4	8 out of 8 consecutive discrepancies fall outside the defined control limits	2
5	15 out of 15 consecutive discrepancies fall outside the defined control limits	1
6	25 out of 25 consecutive discrepancies fall outside the defined control limits	0.5

(b)

Rainfall category c	Selected run rules	N
Dry weather	3–6	1
Rainfall type A	3–6	0.25
Rainfall type B	3–6	0.5
Rainfall type C	4–6	2
Rainfall type D	4–6	3
Post event	3–6	2

number of timesteps containing evidence of a blockage event is over a pre-specified threshold m. If all these rules are satisfied, a blockage alarm is generated.

In addition to the above, the inference module applies a temporary alarm suppression period after every initial alarm is raised, which suppresses all subsequent alarms for a specified period of time. This avoids raising unnecessary alarms for the same event. The suppression period was implemented to prevent operator overloading from multiple alarms for the same blockage. Overloading can make it difficult for the operator to adequately assess simultaneous blockage events and alarms in the system, and increased stress can lead to poor judgments. The addition of the suppression period also has the benefit of significantly reducing the false-positive rate, as any false alarms which occur within the suppression period are removed.

It was hypothesised that the suppression period could also result in the system failing to detect genuine blockages or increase the EDS detection time if a blockage occurred during this suppression period. However, an analysis of the system on real blockage data found that the suppression period did not have a significant effect on either the true detection rate or the detection time.

Once an alarm has been generated and stored in the alarm database, the following information is recorded: (i) the blockage event start time, (ii) the CSO level, and (iii) whether an overflow has occurred. The inference module is designed to operate at the timestep of the incoming data. However, if desired, it can be modified to run over a user-specified time period, e.g. once per day, and then generate a list of any blockage alarms raised in the analysed period.

4. CASE STUDY

4.1. Case study description

To demonstrate and evaluate the performance of the methodology on unseen events, the EDS was applied to several real blockage events which occurred in a large sewer network located in North West England. Sixteen blockage events, from 10 different CSO sites, were designated for testing of the detection methodology. Information relating to the events is presented in Table 3. All the blockage events were identified via a manual visual inspection of the historic level data to identify anomalous behaviour with the characteristics of a blockage. This was found to be the most reliable method of identifying blockage events. The blockages were then cross-referenced with utility blockage removal data.

Table 3 | Summary of validation CSO sites and blockage events

CSO site ID	Location	Sewer network type	Blockage event number	Blockage ID	Blockage type	Blockage duration	Dry weather spill during blockage	Average dry weather CSO level before spill (mm)	Average dry weather CSO level after spill (mm)	Spill level (mm)	Blockage removed by utility
1	Wirral	Urban	1	1	Sudden	6 days 5 h	Yes	70	1,036	1,000	Yes
2	Cumbria	Urban	1	2	Sudden	32 days 4.5 h	Yes	55	164	147	Yes
			2	3	Sudden	66 days 12 h	Yes		160		No
			3	4	Sudden	11 days 19.25 h	Yes		159		Yes
3	Cheshire	Urban	1	5	Sudden	27 days 2 h	Yes	125	195	210	No
4	Wirral	Urban	1	6	Sudden	34 days 7 h	No	13	51	150	No
			2	7	Gradual	335 days 20 h	No		31		No
5	Liverpool	Urban	1	8	Sudden	2 days 20 h	Yes	58	467	900	No
			2	9	Sudden	5 days 6.5 h	No		499		No
6	Greater Manchester	Urban	1	10	Sudden	4 days 11.25 h	No	31	198	565	No
			2	11	Sudden	31 days 14.5 h	No		368		No
7	Cumbria	Urban	1	12	Gradual	229 days	No	21	70	150	No
8	Greater Manchester	Urban	1	13	Sudden	8 days 15.75 h	No	17	151	680	No
			2	14	Sudden	12 days 8 h	No		435		No
9	Carlisle	Urban	1	15	Sudden	43 days 8 h	Yes	88	869	550	Yes
10	Cumbria	Rural	1	16	Gradual	51 days 1 h	Yes	27	104	390	No

The events were selected to be representative of different blockage characteristics, i.e. a mixture of gradual and sudden blockages, various blockage causes (determined from the wastewater utility blockage removal report data), varying durations and sewer responses. The CSO chamber sites were also selected to be diverse in terms of location, catchment size, CSO chamber size and spill frequency. The blockages described here all resulted in an increase in the CSO level. Blockages that cause a level decrease do occur (e.g. blockages in sewers upstream of the CSO). However, these are significantly rarer and therefore have not been considered here.

The time-series water level data (mm) was obtained using an ultrasonic depth monitor installed in the CSO chambers, measured at a uniform resolution of 2 min. Observed radar rainfall intensity data (in mm/h, 5-min temporal resolution and 1×1 km spatial resolution) and forecast rainfall intensity data (mm/h, 15-min temporal resolution and 2×2 km spatial resolution) were obtained from the UK Met Office. During data pre-processing, all data series were interpolated to a common, uniform resolution of 15 min via subsampling.

For each case study site, approximately 60% of data was used to train the EANN required for the EANN discrepancy analysis module and for the detection system calibration. The remaining 40% of data (which included any blockage events) was used for testing (i.e. validation) of the system. However, the precise percentages varied for each site, due to variations in the CSO level monitors' installation dates and differences in the duration and timings of the blockages.

The combined test data from all 10 validation case study sites consisted of 2,240 days in total, of which 715 days occurred during a blockage event. Note that a systematic analysis of the effect of the amount of training data used by the system results has not been carried out yet. However, it has been noted that longer training datasets are beneficial as they increase the confidence of the control chart limits. This is especially important for high-intensity rainfall periods, which are less present in the dataset. To ensure there was sufficient data coverage of both dry and wet weather, the EANN models were required to have a minimum training data length of 2 months without blockages. A preliminary analysis showed that the precise length of data used to train the EANN model utilised in the EANN detection subsystem did not measurably affect the performance of the overall detection system.

During validation, the EDS was applied in a simulated real-time fashion to unseen data, i.e. as it would be used by a utility in real-time. The whole time-series dataset for each CSO site was fed as a continuous stream.

4.2. Example of a blockage event detection using the EDS

An example of the EDS when applied to validation blockage event 1 is presented here to demonstrate how the system would be used by a wastewater utility in near real-time. The CSO chamber is located in Merseyside, North West England in a predominantly residential catchment. The blockage lasted 6 days (the level data during this period is displayed in Figure 3), causing the CSO to overflow continuously for 4 days. As this was an unconsented spill, it had the potential to merit a penalty from the regulator.

The blockage event was cross-referenced with blockage removal data from the wastewater utility for the area in proximity to the CSO. A corresponding blockage removal was identified on 18 July 2017. Compared to the majority of blockage events analysed, this obstruction was removed unusually quickly – indicating that the blockage had an obvious impact which required rapid removal. The obstruction consisted of a soft blockage (i.e. a build-up of soft materials flushed through the system, such as hair or grease) and resulted in a surcharged system and flooding. Historical CSO level and rainfall data were collected from 11 June 2016 to 1 August 2017. The system was calibrated using data from 11 June 2016 to 1 April 2017, and the remaining data, from 2 April 2017 to 1 August 2017, which contained the blockage event, was designated for testing.

The results obtained by the EDS during the blockage period are displayed in Figure 3. The first vertical line indicates the timestep of the first alarm generated. The system detected the blockage event in 1 h 45 min, 3 days before the CSO began to overflow and 6 days before it was removed by the utility. The blockage was detected whilst the increase in level was still small – indicating that the system is sensitive to small deviations from normal behaviour. No false alarms occurred during the 4 months of data used to test the system. Note that Figure 3 also displays a spill event on 11 July 2017 caused by heavy rainfall. The EDS recognised that this was normal CSO behaviour and so did not generate an alarm.

The CSO first began spilling during rainfall. Generally, alarms are raised automatically in the operations centre of a wastewater utility when a spill event occurs. However, alarms are commonly dismissed (either automatically or manually) when they occur during rainfall as they are assumed to be caused by increased flow due to precipitation (and thus are spilling in consent). This alarm dismissal is performed as during a heavy storm there may be hundreds of CSOs overflowing across the network – responding to each alarm to determine the spill cause would overwhelm the operators. Therefore, when

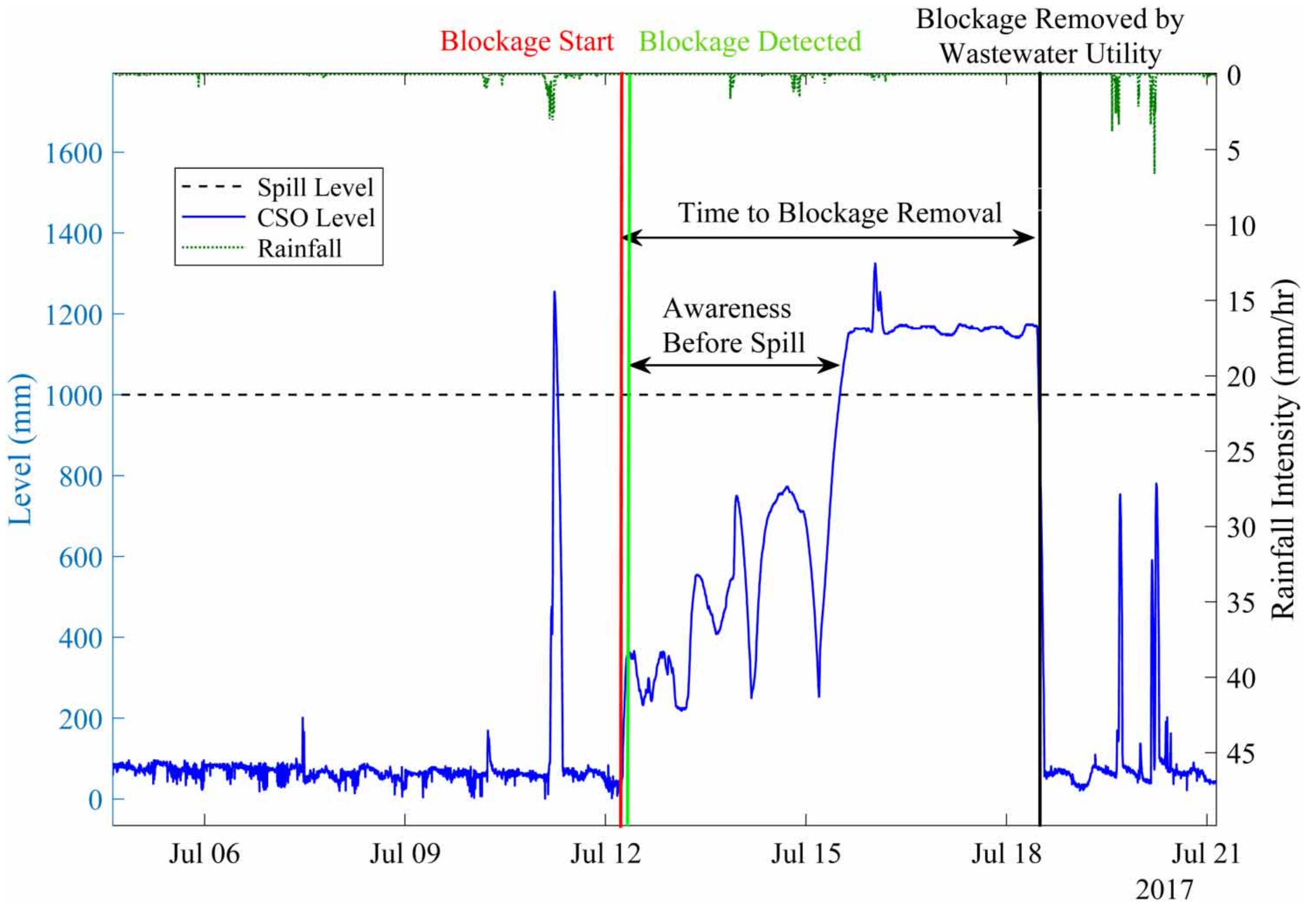

Figure 3 | Results obtained by the detection system during an identified blockage event.

this site began to spill, the utility was not aware of the blockage. It was only when the rainfall event ended and the CSO continued to spill that the utility was alerted to the problem. If the EDS had been employed, the utility would have been alerted to the blockage 3 days before it caused the unconsented overflow.

4.3. Summary of detection results

The overall results of the EDS when applied to all 10 case study sites, containing 16 blockages, are presented in Table 4. The results are given in terms of the detection time of any event(s) (i.e. the time elapsed between the visually identified start of the blockage and the generation of the first corresponding alarm), the total number of false alarms generated for that site, and the rate of false alarms for that site (calculated as the number of false alarms generated divided by the total number of 'non-event' timesteps).

Overall, the results shown in Table 4 indicate that the EDS can effectively detect both gradual and sudden blockages. All 16 events were detected, giving a true positive rate of 100%, and the blockage detection time is generally very good. Regarding the 13 sudden blockages, 12 were detected in under 24 h and 9 were detected in 6 h or less. This appears to be a good result considering the average duration of these sudden blockages (the average time between the blockage occurrence and the blockage removal, due either to removal by the utility or heavy sewer flow due to rainfall, is in the order of 22 days). Only the second blockage event from case study site 2 (i.e. blockage ID 3) has a comparatively long detection time of 163.75 h (i.e. almost 1 week). This occurred as initially the blockage resulted in only a very small increase in level. It therefore took longer for the detection subsystems to generate sufficient evidence to raise an alarm. However, it should be noted that the blockage was still detected a full 2 days before a long duration spill occurred.

Table 4 | Summary of validation blockage detection results

CSO site ID	Blockage ID	Blockage type	Actual blockage start date/time (ground truth)	Blockage detection date/time (by the EDS)	Detection time (h)	False alarm rate	Total false alarms
1	1	Sudden	11/07/2017 21:00	11/07/2017 22:45	1.75	0	0
2	2	Sudden	21/07/2016 11:30	21/07/2016 14:15	2.75	5.97×10^{-4}	4
	3	Sudden	8/11/2016 16:15	15/11/2016 12:00	163.75		
	4	Sudden	28/01/2017 23:45	29/01/2017 21:45	10		
3	5	Sudden	20/06/2016 16:00	20/06/2016 21:45	5.75	0	0
4	6	Sudden	21/08/2016 21:30	22/08/2016 07:15	5	0	0
	7	Gradual	20/10/2016	15/11/2016 21:00	645		
5	8	Sudden	26/05/2016 22:45	27/05/2016 07:00	8.25	0	0
	9	Sudden	13/08/2016 06:45	13/08/2016 22:15	5.5		
6	10	Sudden	19/07/2017 19:15	19/07/2017 23:15	4	6.18×10^{-4}	6
	11	Sudden	15/18/2017 12:25	15/18/2017 14:45	2.5		
7	12	Gradual	07/11/2018	12/11/2018 20:30	140.5	0	0
8	13	Sudden	27/10/2018 23:30	28/10/2018 12:00	12.5	1.12×10^{-4}	3
	14	Sudden	07/11/2018 11:45	07/11/2018 14:45	3		
9	15	Sudden	19/10/16 14:00	19/10/2016 16:15	2.25	2.05×10^{-4}	1
10	16	Gradual	12/09/2017	19/09/2017 21:00	189	0	0

All the sudden blockages were detected by the system before any dry weather spills occurred – therefore, if applied by a utility in near real-time, the EDS would have detected the presence of the blockage before any pollution incidents occurred and may have prevented any unconsented spills.

Regarding the gradual blockage events (blockage IDs 7, 12 and 16), these blockages took longer to detect – at 645, 140.5 and 189 h, respectively. However, this is expected as the change in the CSO level over time caused by gradual blockages is small. Discussion with industry personnel indicated that any preventative action (e.g. jetting) to remove a gradual blockage would only be undertaken once the sewer flow was demonstrably obstructed and could potentially cause flooding/spills during rainfall. Therefore, it is acceptable that the EDS is slower at identifying these types of events.

A total of 14 false alarms were generated over the 10 sites. This corresponds to a false alarm rate (i.e. the number of false-positive timesteps/the total number of timesteps without a blockage) of 0.0089% over the 2,240 days of validation data. The false alarms occurred on four sites only. An examination of the false alarms showed that they generally occurred during wet weather, when the level was higher than 'normal' for that rainfall classification. For some sites, it is possible that the sewer was partially obstructed by material during these periods causing a level increase; however, this cannot be confirmed.

The ratio between the total number of false positives and true positives can be seen as relatively high, at 14–16. This is due partly to the nature of the dataset: blockage events are very rare (the vast majority of data consists of non-blockage periods and, in addition, each blockage counts only as one event although it lasts for many timesteps). Therefore, the classification accuracy required to successfully detect the majority of events without incurring an equivalent number of false alarms would need to be incredibly high. Consequently, to ensure that the majority of blockage events are detected, the number of false alarms generated is relatively high compared to the number of true positives. However, as stated, this corresponds to an overall low error rate when considering the total amount of data the system was applied to. The various thresholds and rules used in the EDS can be modified to address this trade-off differently, resulting in a reduction of total false alarms but at a cost of reduced detection of actual blockage events. Indeed, the analysis performed to select the parameters utilised in this study could be used to aid wastewater personnel to quickly and easily adjust the parameters of the EDS to produce the desired sensitivity.

Overall, therefore, the detection system has been shown to have the capability to detect different types of blockage events reliably and in a timely manner, and with a low rate of false alarms. No work has yet been carried out to assess how well the detection system will work when transferred to other sewer regions. However, the case study CSOs were drawn from a large wastewater network covering the North West of England which contains 72,000 km of sewer pipe. It is hypothesised that the system will generalise well for CSOs beyond this area, as the system parameters are calculated using the CSO level data for each site, and therefore the system is adaptable to different environments. However, it may be that different sewer networks, differences in the built environment, climate, and geology will have an effect on the system performance. Also, it is important to note that the system is only able to detect events which cause a significant enough imprint on the CSO level signals to be identifiable. This limitation can only be overcome by installing additional monitors throughout the sewer system.

Future work will focus on testing and validating the methodology on additional CSO sites containing further variations in CSOs, blockage events and operating conditions, and in testing the system in real-time. There are also several different potential developments which could be investigated. For example, the presence of a blockage event in the sewer network can affect the level in multiple CSO chambers in proximity to each other; therefore, analysing data from multiple CSOs could increase the reliability and sensitivity of the system. Additionally, it is anticipated that additional data sources, such as level data from sewer pipes, wastewater temperature and pumping systems data, could be integrated into the system. The core philosophy of the detection system is detecting deviations from normal behaviour and is generic; therefore, transferring the existing methodology to other data sources is possible.

5. CONCLUSION

The increasingly widespread installation of real-time water level monitors in CSO chambers provides a useful data source for event detection in the sewer network. A novel methodology has been presented which is designed to detect blockage events in the vicinity of CSO chambers in near real-time. The detection system also provides additional information designed to aid the user in determining the appropriate response to the blockage event and, in the case of multiple alarms, help in ranking/prioritising the response.

The proposed methodology applies two different blockage detection techniques: EANN discrepancy-based analysis and statistical trend-based analysis. These systems analyse incoming data using SPC to identify evidence that a blockage has occurred. An inference module is then employed which combines the evidence from the blockage generation modules and enables the raising of alarms when a blockage is detected. The EDS requires only CSO level data, radar rainfall data and rainfall nowcast data, datasets which are currently routinely collected by the UK and many other wastewater utilities in near real-time. Additionally, the methodology is generic and does not require any physically based hydraulic models. The system can thus be applied to any CSO site without additional work, providing historic level and rainfall datasets are available for training (i.e. calibrating) the system.

The system has been tested and validated manually offline on 10 real-life CSO sites with a total of 16 historic blockage events. The system was run as if operating in real-time. The results obtained lead to the following conclusions:

1. The EDS presented in this article can detect blockage events at or in the proximity of a CSO site in a reliable manner, i.e. with a high true detection rate and a low false alarm rate. This is important as a low false alarm rate is a must for any successful detection system used in practice, especially in the case of a larger water utility where multiple events may be taking place simultaneously and so the limited human and other resources deployed to deal with these events need to be carefully utilised.
2. The EDS can detect blockages in a timely manner, typically within few hours of their occurrence. The relatively short detection times obtained in most cases coupled with automated near real-time alarms have the potential to allow the water utility to manage blockage events more proactively, thus reducing harm to the sewer network and damage to the surrounding environment and properties caused by flooding and pollution. The EDS also has the potential to reduce unconsented spills and hence help avoid the corresponding regulatory penalties.
3. The EDS is able to detect both sudden and gradually forming blockages, with the former being easier to detect than the latter. This is because sudden events result in sharper changes in water levels at the CSO chamber.

DATA AVAILABILITY STATEMENT

Data cannot be made publicly available; readers should contact the corresponding author for details.

REFERENCES

Cherqui, F., Belmeziti, A., Granger, D., Sourdril, A. & Le Gauffre, P. 2015 Assessing urban potential flooding risk and identifying effective risk-reduction measures. *Science of the Total Environment* **514**, 418–425. https://doi.org/10.1016/j.scitotenv.2015.02.027.

Choi, D. Y., Kim, S.-W., Choi, M.-A. & Geem, Z. W. 2016 Adaptive Kalman filter based on adjustable sampling interval in burst detection for water distribution system. *Water* **8** (4), 142. https://doi.org/10.3390/w8040142.

Detectronic 2018 *Predicting Blockages, Reducing Flooding and Improving Network Performance*. Available from: https://www.detectronic.org/wp-content/uploads/2018/09/MMS-B001-overview_bro.pdf.

Detectronic 2019 *What Is a Waste Water Flow Meter?* Available from: http://wpdev5.thelegaldc.co.uk/what-is-a-waste-water-flow-meter/.

EMS 2020 *SMART Sewer*. Available from: https://www.em-solutions.co.uk/innovation/smart-sewer-project/.

Fernando, A. K., Zhang, X. & Kinley, P. F. 2006 Combined sewer overflow forecasting with feed-forward back-propagation artificial neural network. *Transactions on Engineering, Computing and Technology* **1** (12), 58–64.

Jin, Y. & Mukherjee, A. 2010 Modeling blockage failures in sewer systems to support maintenance decision making. *Journal of Performance of Constructed Facilities* **24** (6), 622–633. https://doi.org/10.1061/(asce)cf.1943-5509.0000126.

Koutsoyiannis, D., Kozonis, D. & Manetas, A. 1998 A mathematical framework for studying rainfall intensity-duration-frequency relationships. *Journal of Hydrology* **206** (1–2), 118–135. https://doi.org/10.1016/S0022-1694(98)00097-3.

Ma, H., Bandos, A. I., Rockette, H. E. & Gur, D. 2013 On use of partial area under the ROC curve for evaluation of diagnostic performance. *Statistics in Medicine* **32** (20), 3449–3458. https://doi.org/10.1002/sim.5777.

Mounce, S. & Machell, J. 2006 Burst detection using hydraulic data from water distribution systems with artificial neural networks. *Urban Water Journal* **3** (1), 21–31. https://doi.org/10.1080/15730620600578538.

Mounce, S. R., Boxall, J. B. & Machell, J. 2010 Development and verification of an online artificial intelligence system for detection of bursts and other abnormal flows. *Journal of Water Resources Planning and Management* **136** (3), 309–318. https://doi.org/10.1061/(asce)wr.1943-5452.0000030.

Mounce, S., Shepherd, W., Sailor, G., Saul, A. & Boxhall, J. 2014 Application of artificial neural networks to assess CSO performance. In: *13th International Conference on Urban Drainage*.

Myrans, J., Everson, R. & Kapelan, Z. 2019 Automated detection of fault types in CCTV sewer surveys. *Journal of Hydroinformatics* **21** (1), 153–163.

Park, J. S. & Kim, S. 2020 Improved interpolation and anomaly detection for personal PM2.5 measurement. *Applied Sciences (Switzerland)* **10** (2). https://doi.org/10.3390/app10020543

Rodríguez, J. P., McIntyre, N., Díaz-Granados, M. & Maksimović, Č. 2012 A database and model to support proactive management of sediment-related sewer blockages. *Water Research* **46** (15), 4571–4586. https://doi.org/10.1016/j.watres.2012.06.037.

Romano, M., Woodward, K., Kapelan, Z. & Savić, D. 2014 Near real-time detection of pipe burst events in cascading district metered areas. In: *11th International Conference on Hydroinformatics, HIC 2014*.

Rosin, T. R., Romano, M., Keedwell, E. & Kapelan, Z. 2021 A committee evolutionary neural network for the prediction of combined sewer overflows. *Water Resources Management* **35** (4), 1273–1289. https://doi.org/10.1007/s11269-021-02780-z.

Sattari, M. T., Rezazadeh-Joudi, A. & Kusiak, A. 2017 Assessment of different methods for estimation of missing data in precipitation studies. *Hydrology Research* **48** (4), 1032–1044. https://doi.org/10.2166/nh.2016.364.

Southern Water 2018 *Making Fatbergs a Thing of the Past – SMART Sewer*. Available from: https://www.southernwater.co.uk/the-news-room/the-media-centre/2018/september/making-fatbergs-a-thing-of-the-past-smart-sewer.

Ugarelli, R. & Taylor, P. 2016 Historical analysis of blockages in wastewater pipelines in Oslo and diagnosis of causative pipeline characteristics. *Urban Water Journal*. https://doi.org/10.1080/1573062X.2010.526229.

Water UK 2017 *New Proof That Flushing Wipes Is a Major Cause of Sewer Blockages*. Available from: https://www.water.org.uk/news-item/new-proof-that-flushing-wipes-is-a-major-cause-of-sewer-blockages.

Western Electric Company 1958 *Statistical Quality Control Handbook*, 2nd edn. Western Electric Co, New York.

Widmer, G. & Kubat, M. 1996 Learning in the presence of concept drift and hidden contexts. *Machine Learning* **23** (1), 69–101. https://doi.org/10.1023/A:1018046501280.

Youden, W. J. 1950 Index for rating diagnostic tests. *Cancer* **3** (1), 32–35. https://doi.org/10.1002/1097-0142(1950)3:1 < 32::AID-CNCR2820030106 > 3.0.CO;2-3.

Zhang, D., Lindholm, G. & Ratnaweera, H. 2018 *DeepCSO: Forecasting of Combined Sewer Overflow at a Citywide Level Using Multi-Task Deep Learning*. arXiv:1811.06368.

Zweig, M. H. & Campbell, G. 1993 Receiver-operating characteristic (ROC) plots: a fundamental evaluation tool in clinical medicine. *Clinical Chemistry* **39** (4), 561–577. https://doi.org/10.1093/clinchem/39.4.561.

First received 9 March 2021; accepted in revised form 18 February 2022. Available online 2 March 2022