Article:

Shepherd, W. orcid.org/0000-0003-4434-9442, Mounce, S., Gaffney, J. et al. (5 more authors) (2023) Cloud-based artificial Iintelligence analytics to assess combined sewer overflow performance. Journal of Water Resources Planning and Management, 149 (10). 04023051. ISSN 0733-9496

https://doi.org/10.1061/JWRMD5.WRENG-5859

CLOUD BASED ARTIFICIAL INTELLIGENCE ANALYTICS TO ASSESS COMBINED SEWER OVERFLOW PERFORMANCE

[1]Will Shepherd, [2]Stephen Mounce, [3]Gavin Sailor, [4]John Gaffney, [5]Neeraj Shah, [6]Nigel Smith, [7]Adam Cartwright and [8]Joby Boxall.

[1]Affiliation: Pennine Water Group, Department of Civil and Structural Engineering, University of Sheffield, Sheffield, S1 3JD, UK. Email: w.shepherd@sheffield.ac.uk

[2]Affiliation: Pennine Water Group, Department of Civil and Structural Engineering, University of Sheffield, Sheffield, S1 3JD, UK. Email: s.r.mounce@sheffield.ac.uk

[3]Affiliation: Pennine Water Group, Department of Civil and Structural Engineering, University of Sheffield, Sheffield, S1 3JD, UK. Email: g.sailor@sheffield.ac.uk

[4]Affiliation: Digital Industries, Siemens plc, Princess Road, Manchester M20 2UR, UK. E-mail: john.gaffney@siemens.com

[5]Affiliation: Digital Industries, Siemens plc, Princess Road, Manchester M20 2UR, UK. E-mail: nss350@gmail.com

[6]Affiliation: Yorkshire Water Services Limited, Western House, Halifax Road, Bradford, BD6 2SZ, UK. E-mail: nigel.smith@yorkshirewater.co.uk

[7]Affiliation: Digital Industries, Siemens plc, Princess Road, Manchester M20 2UR, UK. E-mail: adam.cartwright@siemens.com

[8]Affiliation: Pennine Water Group, Department of Civil and Structural Engineering, University of Sheffield, Sheffield, S1 3JD, UK. Email: J.B.Boxall@sheffield.ac.uk

ABSTRACT

Discharges from combined sewer overflows (CSO) are increasingly unacceptable, particularly when these are not linked to wet weather. This paper presents evaluation of an online artificial intelligence based analytics system to give early warning of such overflows due to system degradation. It integrates a cloud based data driven system using artificial neural networks and fuzzy logic with near real-time communications, taking advantage of the increasingly available real-time monitoring of water depths in CSO chambers. The data driven system has been developed to be applicable to the vast majority of CSOs and requiring a minimum period of data for training. Results are presented for a live assessment of 50 CSO assets over a six-month period demonstrating continuous assessment of performance and reduction of CSO discharges. The system achieved a high true positive rate (86.7% on confirmed positives) and low false positive rate (3.4%). Such early warnings of CSO performance degradation are vital to proactively manage our ageing water infrastructure, and to achieving acceptable environmental, regulatory and reputational performance. The system enables improved performance from legacy infrastructure without gross capital investment.

Key words: Combined sewer overflows, Artificial neural networks, Fuzzy inference system, Cloud computing, Internet of Things, Rainfall radar, Depth prediction.

Practical Applications

Combined sewerage networks convey both wastewater from residential and commercial properties as well as rainfall runoff from the urban catchment. Combined Sewer Overflows, widely known as CSOs, provide a relief valve when runoff from rainfall would overwhelm the downstream network and treatment works. Excess water is spilt to a nearby watercourse, ideally when the watercourse flow has increased to provide additional dilution and thus minimise impacts. If a blockage or other defect downstream of a CSO results in a decrease in discharge capacity, the CSO can spill earlier than it is designed to, or even in dry weather. Prior to the deployment of level sensors, such premature spills could only be identified through a visible spill or water quality impact. Sensors allow water utilities to monitor depths in CSO chambers, however each utility will have a large number of CSOs, thus an automated system is needed to identify premature spills. This paper discusses the development and validation results obtained from a pilot deployment of a data analytics solution to identify abnormal water depths in a CSO.

INTRODUCTION

Drainage systems in urban areas across the western world have developed over long periods of time, expanding in a piecemeal manner with the populations that they serve. In many locations, combined sewer systems were installed, capturing both foul flows and runoff from rainfall. Excess flows during heavy rainfall would overwhelm downstream networks and treatment works and potentially cause flooding, thus combined sewer overflows (CSOs) are used to divert excess flows to a receiving watercourse.

Legislation, such as the EU Water Framework Directive (Council Directive (EC) 2000/60/EC) means that pollution of water courses is under increasing scrutiny, and discharges from CSOs have the potential to be a significant source of pollution, especially if they are not operating as designed. Poor performance of a CSO is often a function of both the variety of waste discharged through the sewer system and also the age and condition of the assets. These two factors can either in isolation or together reduce the capacity of the system through partial or even complete blockages. Where these blockages occur downstream of a CSO it is likely that the CSO will discharge to the watercourse prematurely, causing pollution. In the past, there was little option but to assume CSOs were operating as designed unless evidence suggested otherwise. Due to infrequent and spatially sparse sampling of rivers, the most likely source of this evidence would be from members of the public noticing the pollution or even a discharge from the outflow pipe itself. There has always been the potential for water service providers (WSP) in the UK to be fined for unconsented discharges, but this has tended to occur relatively infrequently due to the lack of definitive evidence.

Increased monitoring of CSOs has become feasible with advances in technology. In some cases uptake has been politically driven, for example in the UK, Richard Benyon MP (2013) wrote to water company chief executives asking for monitoring of the vast majority of their CSOs by 2020. This led to the Environment Agency requiring installation of event duration monitoring (EDM), which was a significant feature in AMP6 (five year Asset Management Periods used in the UK water industry) running from 2015-2020 in order to assess CSO performance. While EDM provides broad data on CSO performance, the potential to use such water depth data for the day to day management has also been recognised (e.g. Sumer et al., 2007). Ofwat has posed significant efficiency challenges to UK WSPs for AMP 7 (running until 2025), including cutting pollution incidents by more than a third (Ofwat, 2019). The UK context is mirrored worldwide, to one extent or another, regulations concerning the operation of CSOs vary significantly and are often linked to annual spill counts, or impacts, as discussed by Botturi et al (2021). In the EU there is an ongoing review of the Urban Waste Water Treatment Directive which regulates the discharge of waste water (EU, 2022). This provides a strong driver for the water industry to make proactive use of the available data in order to meet or exceed these targets.

The presence of CSOs within sewerage networks is a somewhat contentious issue. One school of thought is that they should not exist and that sewer networks should be dual systems with separate pipe for foul water and for storm water. Separated sewer systems have their own problems in that mis-connections are common when properties are extended or upgraded, potentially resulting in untreated discharges occurring continuously in dry weather from storm systems. Furthermore, the washoff from urban catchments cannot truly be considered clean with washoff of hydro-carbons, heavy metals and bacterial pathogens. Beyond this is the consideration that replacement of existing combined sewer systems with separated systems would, in the majority of cases, be prohibitively expensive. Looking forward, changing climate will compound the challenges our already struggling sewer systems face in meeting the demand of growing populations and urbanisation. Recent findings suggest that co-occurrence of rising sea levels, storm surges, and increased precipitation will lead to an increase in 'compound flooding' and increased pressure on sewer systems, and therefore the likelihood in the need to use CSO infrastructure (Fortier and Mailhot, 2015). It is hence reasonable to assume that CSOs, and their potential discharges, are a feature we have to accept in many places and that the challenge is to minimise unintended discharges from them, ideally with the minimum of investment in new built infrastructure. This research presents a likely key technology to achieve this.

This paper presents evidence of how an online artificial intelligence system can be an effective advance warning system of degradation in CSO performance, providing information that can be acted on proactively to help avoid unintended or premature CSO discharges. The contributions include development of a fuzzy logic system for classification, cloud based implementation to enable scalability, and historic and live validation to evidence the veracity of the information derived.

BACKGROUND

Internet of Things (IoT) objects and sensors connect to the cloud giving rise to the concept of 'smartness' and the development of 'Smart cities' and 'Smart water.' The sensing of data that could not be gathered in the past and collecting them on IoT platforms enables new value to be created. As these technological capabilities advance, so does the ability to collect information from remote devices and correlate that information across diverse systems. An infrastructure that can connect the monitoring and control systems to an IoT platform allows the effective use of the operational information the systems hold, and helps to achieve near-real time situational awareness based on digital performance twins. Hence, a new generation of smart and connected urban sewer systems will be enabled by emerging wireless technologies and data algorithms.

IoT enabled urban drainage systems can play an essential role in the "smart water cities" of the future, where sewerage infrastructure evolves from being passive to adaptive units that can proactively respond depending on any given situation (Lund et al. 2018). Water utilities are starting to take advantage of this, for example deploying arrays of sensors that capture and generate time-series data in real time. Transforming this data into timely, relevant insight using rich analytics is a key goal of any cloud-based, open IoT operating system.

Data from CSOs can be very valuable in understanding the performance of that asset and the immediately adjacent sewer system (Bachmann-Machnik et al. 2021), however with a large number of assets (in the thousands for many WSPs) it is not feasible (or affordable) to manually interpret this data and deterministic centralised modelling is often too complex, uncertain and time consuming. Data driven Artificial Intelligence (AI) systems are an option to address this, offering a way to incorporate the data without resorting to detailed physically-based mathematical models with their inherent high computational and calibration requirements. One recent study showed a near 5-fold performance improvement (ratio of overflows to precipitation) was achieved after commissioning of a real-time sensing (and subsequently) control system (Kerkez et al. 2016). Data driven software sensors have been used to estimate CSO emission flow rates from complex CSO structures by utilising correlation analyses between physical water depth sensors and discharge measurements (Ahm et al. 2016).

Data-driven modelling seeks to provide a mapping between the inputs and outputs of a given system, with little prior process knowledge – and is now being widely adopted for prediction and classification in water systems. More complex control algorithms have been shown to outperform more simple control strategies (van der Werf et al. 2022). Artificial Neural Networks (ANN) are one such approach, being universal computing machines capable of arbitrary non-linear function approximation (Hornick et al. 1989) for pattern recognition, classification, generalisation and abstraction, and the interpretation of incomplete or noisy data (Lingireddy and Brion 2005). Recent research for urban drainage systems has explored the utilisation of rainfall radar data, hydraulic models and data-driven modelling approaches for the prediction of urban flooding in real-time (for example Duncan et al. 2013, Garcia et al. 2015). Fernando et al. (2006) applied a standard feed-forward, back-propagation ANN model to forecast the occurrences of wastewater overflows in a combined sewerage system. The data used included the traditional model predicted overflow rates for one overflow structure and artificially generated rainfall for the rain-gauge in the closest proximity. Sumer et al. (2007) researched the feasibility of real-time detection of sanitary sewer overflows (SSOs) using time series analysis and ANN techniques in two case studies in Arizona, USA. An ANN was developed to estimate the 6-hour component of the forecast. In order to identify whether an SSO was occurring, control limit theory was used to detect important deviations between measured and expected depth and flow data. Kurth et al. (2008) demonstrated that a three hidden-layer Multilayer Perceptron ANN trained with back-propagation is capable of learning the underlying relationship between local rainfall occurrence and CSO response. In order to predict water depths 3 time steps into the future (fifteen minutes), lags of twelve previous values of two rain gauges and a lag of five of recent water depths for a CSO chamber were used. In Guo and Saul (2011) the concept of CSO Analytics was introduced in which an ANN (adaptive linear) was used to predict, at times of dry weather and in response to rainfall (measured using in catchment rain gauges), the hydraulic performance of a CSO in terms of flow depth. Mounce et al. (2014a) further developed this approach to incorporate rainfall radar data and demonstrated a prediction of CSO depth with less than 5% error for predictions more than one hour ahead for unseen data. Cross correlation was used to explore the spatial (rainfall radar cells) and temporal (time lags) i.e. the time of concentration and hence to inform the ANN inputs for a number of models. Whilst Mounce et al. (2014a) showed ANNs could be used to accurately predict future water depths in CSOs, based on radar rainfall (rather than rain gauges) and recent water depths, this methodology was unable to indicate when the performance of the CSO changes due to its reliance on recent depths. When a blockage occurs, it was found that the predicted water depth very rapidly followed the measured trend. Subsequent work (Mounce et al. 2014b) addressed this issue and provided a performance assessment by further classification of model outputs in order to provide a per asset state on a daily basis by developing a fuzzy logic based ‘traffic light’ evaluation system.

Other authors have explored similar approaches. Rosin et al. (2018, 2021) applied evolutionary ANN models to predict water depth in several CSO chambers up to 6 hours ahead using inputs of past CSO depth, radar rainfall and rainfall forecast data. This system was applied offline to four CSOs and the authors note the potential for future online operation for blockage detection. They found that it is more difficult to model major rainfall events precisely at higher forecast horizon values. In their most recent work, Rosin et al. (2022) further tested the system by incorporating Statistical Process Control for blockage detection, and validated it manually offline on 10 real world CSO sites with a total of 16 historic blockages. Bizer and Kirchhoff (2022) developed performance indicators based on regression modelling and applied these to a historic dataset for 11 CSOs (using CSO data and hourly precipitation in Cumberland, Maryland for the years 2005–2020). Annual thresholds of 1-hour precipitation intensity above which CSO incidence is predicted (and below which it is not predicted) in each year were identified. Subsequently they built a regression model to predict CSO volume from the precipitation depth and average intensity of the preceding rainfall event. Some initial work has explored using Deep Learning (Lecun et al. 2015) for a multi-step-ahead (close horizon) prediction of CSO water depth collected by IoT (Zhang et al. 2018a, 2018b). Kanneganti et al. (2022) applied a random forest model to predict sewer flow rates in 3 separate sewer systems with an accuracy of 91.7%, albeit for daily flows and with a short 5.5 week test period. It was demonstrated for case studies that fairly precise time series predictions could be produced for sewer system management, however there are few examples of the application of such predictions for system management and those that do are based on small validation case studies (e.g. Bailey et al., 2018). Any system that is practicably deployable by a water utility needs to be demonstrably robust and scalable.

METHODOLOGY

Previous work described in the literature has indicated the potential of Artificial Intelligence (AI) techniques to be able to predict water depths within CSOs. The work described here has further developed the approach described by Mounce et al. (2014b). An important advance in this research was to take into consideration the effects of rainfall, this was done by combining the ANN for prediction with a Fuzzy Inference System (FIS) which flagged significant changes in CSO performance in near real time. In order to demonstrate transformational scalable capability, the tool was re-written in Python and deployed on the MindSphere IoT open operating system, and utilised for a real world 50 asset case study. To prove the value of the information derived, validation from a two year historic period (2017-2019) and a 6-month live period (2020) is presented based on manual data interpretation, integration with the water utility control room and operational teams. Evaluation also included comparison to a moving average based legacy system.

System overview

The WSP partner has used a suite of tools, termed pollution tracker (PT) to analyse CSO depth data based on moving averages and rates of change of depths. This information is processed on a daily basis and ranked, the top fifty ranked CSOs are considered to be potentially underperforming and thus flagged for further investigation. During dry weather PT is able to flag assets which may be performing badly. However, the analysis is based purely on the data from the CSO depth monitors, so is unable to correctly understand the difference between water depths changing legitimately as a function of rainfall or when the water depths are changing because of a drop in performance of the sewer network. A CSO may therefore rank highly either because it has poor performance, due for example, to a blockage in a downstream pipe, or because there has been rainfall in its catchment.

The CSOA (CSO Analytics) cloud based methodology consists of a number of processes, as shown in Fig. 1. Each CSO has a unique response to rainfall, which is a function of the characteristics of the catchment and sewer networks upstream and downstream of the CSO, as well as the design of the CSO itself. It is therefore necessary to train each ANN to each asset's individual performance using historic CSO depth and rainfall radar data. However, in order to ensure scalability, training overheads need to be kept to a minimum, this has been achieved through the use of a transferable ANN architecture. Once the model is trained, it can be used by the live cloud system, where rainfall radar data is input to the ANN to produce a predicted depth. A FIS is then utilised to determine whether the measured data is within an acceptable range of the prediction.

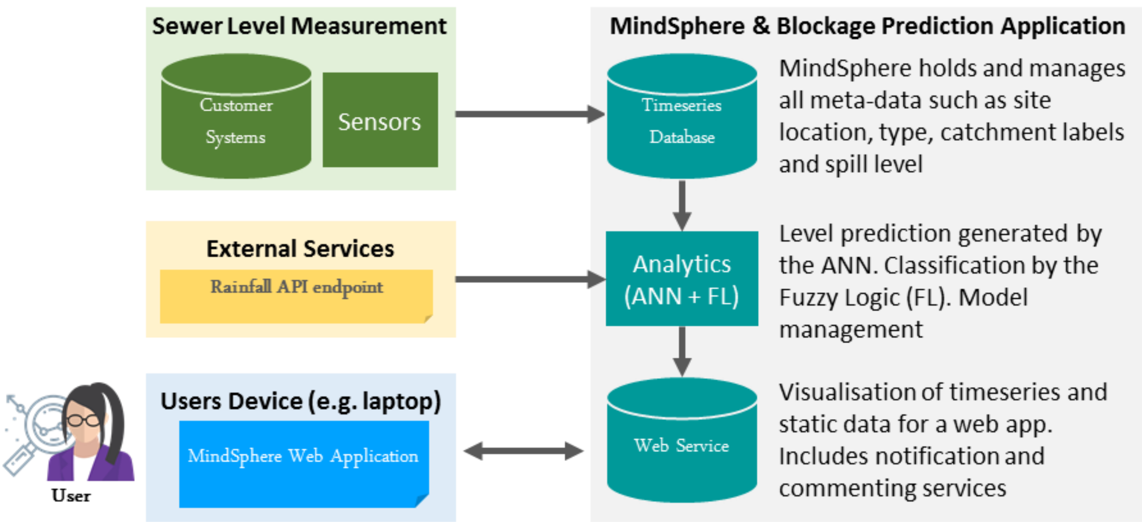

Fig. 1: System Architecture

A live system ran on the IoT platform providing continually updated results, as well as the facility for historic replay (see Fig. 2). Data from the live system and site summaries were available in the water company's operation control room through this live dashboard.

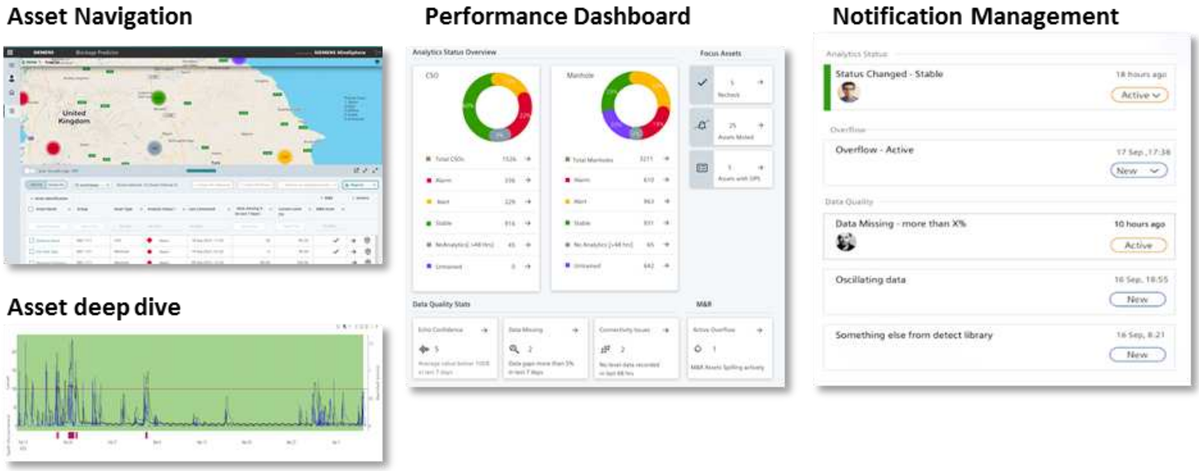

Fig. 2: IoT platform for system access

Data processing and quality

Water depth data from CSOs from the WSP telemetry system was used at 15-minute resolutions, the data units were a percentage value (in some cases 100% does not equate to the spill level). Rainfall radar data was obtained from WeatherOnline (the original producer being the MetOffice) also at 15-minute resolution. The data units were in mm (amount of rainfall in the last 15 minutes) with a spatial resolution of 1 km cells. Cells within a site specific radius (determined as part of the training) of the CSO were used in the analysis. Water depth data has missing values linearly interpolated.

Historic data was used in the training process, this covered the period from July 2017 to July 2019, although data from some sites was not available from the beginning, the average number of historic data days was 583, within this there were an average of 14 days of missing data. The data selected for training did not need to be continuous, but was selected based on data quality and consistency. Data would be rejected where it was missing for a significant period (e.g. > 1 day in dry weather, > 0.5 day in wet weather); periods of time where the data was noisier than usual for the site; periods where the background depth varied significantly from normal and was not due to seasonal variability (e.g. where there was a potential performance issue (e.g. a blockage), the sensor had been moved, was obstructed, or otherwise providing inconsistent data). Where the data contained rejected periods, date ranges with good data were collated. If the total period of good data was greater than 60% of the whole data set then the good data was split 60-40 into training and test sets. If the good data period was less than 60% of the whole data set then all of the good data was used for training. The rationale was to capture variation in weather and asset behaviour that occurred over the course of the year.

Artificial Neural Network

The ANN was implemented in Python (using the PyTorch machine learning framework). The model uses rainfall radar data as the predictor and the depth data input as a target, depth is predicted at the current time step. In order to capture diurnal variations, the ANN training uses two time features (sine and cosine of the hour) to replicate the diurnal dry weather flow, along with rainfall data. A standard one layer feed- forward ANN with no hidden layers was utilised to minimise training data requirements and for transferability (i.e. the ANN architecture as regards the structure was not changed from one asset to the next). The goal was not point wise accuracy for the depth prediction itself but a learning of the rainfall/depth response (for the depth at the current time following a rainfall lag of values), thus the overhead of recurrent networks or multi-time step predictions was avoided.

The spatiality and temporality of rainfall data as a predictor of CSO depth is a complex function of the sewer network. Mounce et al (2014b) used manual selection of radar cells based on assessment of the upstream sewer network, but this is not scalable. Hence for each CSO a total of 12 models were created using different rainfall data combinations, these being four spatial windows with 1, 2, 3 and 5 km radius and 3 temporal windows of 6, 12 and 24 hours (prior to the prediction time step). These temporal window periods were based on extensive autocorrelation tests (cf. Mounce et al., 2014a). This method is a better candidate for future automation removing the requirement for subjective human assessment.

The number of inputs was dependent on the amount of rainfall cells being included and the temporal window length. Since 15 minute data was used, total inputs were equal to rainfall_hours * 4 * num_cells depending on which of the 12 model types (therefore 6 hours rainfall for one cell would result in 24 inputs). The activation function on the output layer was rectified linear units (ReLU). This was used for better convergence and avoiding outputs less than 0, which should not in reality be possible in CSOs (except for calibration problems with instrumentation). General Matrix Manipulation with gradient descent was used as the optimiser and RMS (Root Mean Squared) as the cost function on the training data. Empirical trials revealed standardised values of 0.01 for learning rate and 50 epochs of training provided good performance for training with the datasets used.

For each asset, multiple models were trained and the best models were selected by the following criteria: loss on test data, with a general bias for selection of models that slightly over predict rather than under predict to avoid false positives (based on visual assessment to ensure a good overall fit, especially to the wet weather periods and with particular emphasis on predicting delayed runoff response).

The ANN used in this work is intended to work with a minimum number of parameters and a minimum period of historic data. This architecture reduces the overhead required for training the ANN and thus helps to ensure scalability, at the expense of some accuracy. However, the degree of accuracy is appropriate to work with the FIS system to enable assessment of performance with regard to rainfall influences.

Fuzzy Inference System

Having an acceptably accurate prediction of the CSO depth is the first stage of the system, the second stage is to identify whether the actual CSO depth is within a reasonable range of the predicted depth or not. This is a relatively easy task for an experienced professional, but these come at a premium and are not scalable. It is challenging for an automated system to capture the processes of such experts. Fuzzy logic (FL) is a useful technique for building systems that can incorporate the impreciseness associated with human reasoning and can be used to determine whether the measured data is within an acceptable range of the prediction.

The FL was implemented using the Python library Scikit-fuzzy using the difference between the measured and predicted CSO depths and the recent rainfall. Some of the complexities that the expert is judging, and which the FL system is capturing, are as follows: during dry weather, including insignificant rainfall, the depth in most CSOs follows a diurnal pattern around a mean value. At some locations this mean value varies according to the catchment wetness which affects the volume of infiltration into the sewer network. Depending on local topography and soil types this infiltration may be just for a period of hours or days after significant rainfall, or may follow a seasonal pattern (i.e. higher baseline water depths in winter than summer). In dry weather it is expected that the depth residual (predicted depth minus measured depth) will vary within a consistent range. During heavier rainfall however it is more difficult to make an accurate prediction, hence the absolute value of the residual will increase. Following a rainfall event, the CSO depth should gradually return to the dry weather condition.

The FL uses a moving window of data, Mounce et al. (2014b) used an 8 hour moving window, and this period was a function of the data and analysis only being updated on a daily basis. This new work is designed to benefit from regularly updating field data by using a shorter moving window. Testing showed that a 1 hour window provided the best trade-off between response to performance issues and stability of generated alerts (i.e. a shorter window provides a quicker response, but is significantly affected by short term fluctuations in the data). The data input to the FL is the 1 hour total rainfall and the 1 hour mean depth residual (predicted – measured). Four key values calculated from the training data period, were used in the FL, these were: the n^{th} percentile of the 1 hour mean depth residual (Depth n^{th} percentile, where n is between 1 and 3, determined on a per asset basis) representing a measured depth that is considered significantly higher than the predicted value, this percentile is variable to adjust the sensitivity of the FL; the 50^{th} percentile of the 1 hour mean depth residual (Depth 50^{th} percentile) representing an average residual; the 80^{th} percentile of the 1 hour rainfall depth (Rainfall 80^{th} percentile) to represent insignificant rainfall; and the 90^{th} percentile of the 1 hour rainfall (Rainfall 90^{th} percentile) represents the rainfall depth above which rainfall is considered significant (wet).

The FL has four input membership functions, each of which uses the number of times that the data for the past three hourly time steps is less than the lower percentile or higher than the high percentile. A degree of membership to each membership function is calculated resulting in a value between 0 and 1. These feed into three rules, which can be summarised as:

> If the residual is not significantly negative, status is normal, i.e. regardless of rainfall, if the predicted depth is not significantly lower than measured;

> If the rainfall is low and residual is significantly negative, status is abnormal high, i.e. the weather can be considered dry, but the measured depth is significantly higher than predicted;

> If the rainfall is high and residual is significantly negative, status is normal, i.e. in wet weather the quality of the prediction is lower, thus to avoid unwarranted alerts, a measured depth that is higher than predicted can be considered normal.

The fuzzy inference methodology used is the Mamdani method (Mamdani and Assilian, 1975), this is applied to the above rules to generate the output membership functions shown in Fig. 3. This was then de-fuzzified to produce a 3 category ‘traffic light’ status system, termed stable, alert and alarm, the output variable boundaries being: Stable < 0.6; $0.6 <$ Alert < 0.8; $0.8 <$ Alarm.

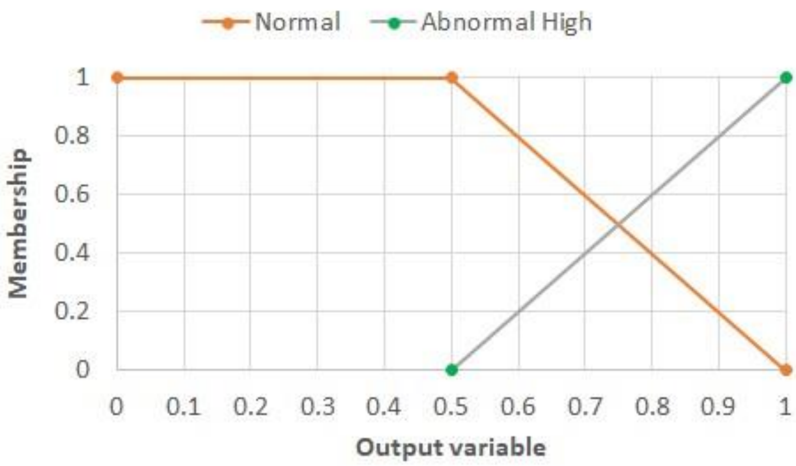

Fig. 3: Output Membership Function for CSOA

Case study and validation

The case study reported here concerns a collaborative project with a UK water utility company. The WSP partner had rolled out IoT ready level sensors in a large number of its CSOs and developed a number of systems and practices that provided an overview of their CSO performance based on recorded depths, termed pollution tracker (PT). However, whilst these systems adequately captured trends in dry weather flow, they did not incorporate rainfall data, and therefore offered little benefit during or immediately after any rainfall events within the WSP's area. The validation of the AI system reported here was conducted on a subset of fifty of these CSOs. The selection criteria were designed to obtain a cross section of CSOs covering a range of performance and location characteristics within the company's catchment area, and to include some with recorded historic unconsented discharges. The cloud based system was compared against the pre-existing baseline of the PT system, where the daily top fifty ranked sites are considered as alerts. Each CSO has data within two discrete periods, historic data from July 2017 to July 2019 as used in training and testing, and live data, when the CSO was on MindSphere, from January to July 2020. Both of these periods are sufficiently long to cover multiple seasons, thus the system is tested against a range of rainfall types and catchment wetnesses. The operations team provided their records of when maintenance teams were sent to a CSO (termed as a job raised) to investigate and resolve any issues, there was provision for maintenance teams to provide useful feedback (e.g. 'blockage cleared'), but often no feedback is received, sometimes raised jobs are precautionary and with hindsight were not required, hence a job being raised is not proof of a performance issue. Conversely there are times when, with hindsight, it is clear that there was a performance issue with the CSO, but no job was raised, or there was a delay in a job being raised. It should be noted that an apparent performance issue can also be a sensor issue (e.g. the sensor was accidentally moved during maintenance, or there is debris obstructing the sensor), but without further information it is impossible to identify whether the problem is with the sensor or with water levels. The sensor data is generally robust and any issue with the sensors is worthy of being flagged by an alert as it detracts from the sensor's ability to identify a performance issue. Validation was therefore carried out manually using the best available information from the operations team, combined with expert judgement as to whether the CSO was performing acceptably. This 'best available information' classification denotes whether an alert should reasonably be returned and is divided into three categories. Normal (green) indicates that the CSO appears to be performing normally. Probable performance issue (amber alert) indicates when the CSO seems to have a higher than expected water depth and it is not currently, nor recently has been raining significantly; this could be interpreted as there possibly being a blockage. Definite performance issue (red alarm) is only returned if a job report confirms a blockage was present, or if the observed water depth is significantly higher than normal and there is no apparent reason for this. The validation methodology was discussed and agreed with the operations team using example data, it was not feasible to discuss and agree a classification for each event.

RESULTS

Examples of individual events are provided to show the operation of the system and how validation was performed, then per site summaries are illustrated to show the range of operational performance over time and finally an overview of results for the full group of CSOs over a longer time period are presented to evaluate overall system performance.

Example events

Three examples are presented to exhibit the difference in performance of different assets and the analytics tool. Each example shows a compilation of four plots, the top being rainfall intensity from radar data, the mean, minimum and maximum across the cells input to the analytics is shown. The second plot shows the measured and predicted depths of water in the CSO. The third plot shows the residuals, calculated as predicted minus measured, thus a negative value (i.e. measured depth is greater than predicted) indicates a potential blockage downstream of the CSO. The bottom plot shows the warnings output by CSOA and PT, 0 represents normal stable operation; 0.5 an alert (CSOA only); and 1 an alarm (for PT this is taken as the CSO appearing in the top 50 ranking).

Fig. 4 shows a well-defined event at CSO 20 which, based purely on the time series data was very likely a blockage. The event starts around 16:00 on 30th January in dry weather, as indicated by a sudden increase in water depth and decrease in the residual. CSOA flags this event almost immediately after it occurs. The site is ranked 1st in PT on 1st and 2nd February and drops to 15th on 3rd February. A high priority job was raised on 1st February, the job was completed with an outcome of ‘Blockage Removed’ on 1st February. The depth plot shows that water depths returned to normal on 1st February in the early afternoon, the final CSOA alert appears shortly after the blockage was removed. The CSOA alert occurs significantly earlier than that of PT and the alert status also returns to normal more quickly.

Fig. 4: Analysis of a blockage event at CSO 20 during January/February 2019 showing rainfall, resulting depth and the response of both PT and CSOA analytics.

Fig. 5 shows an event at CSO 10, with a smaller change in depth and noisier data. The depth plot shows a sudden significant increase on 30th April in dry weather, this drops down to a depth that is still higher than normal on the same day. The CSOA responds rapidly to the initial event, but alarming red over a short period initially, then having many short periods at amber alert and occasionally red alarm. This periodic alerting is due to the short time window used and also likely to be a function of the sensitivity. PT ranks the site in the top 50 worst sites 13 times between 1st and 19th May, four jobs were raised on 3rd, 7th, 13th and 18th May.

The water depths can be seen to reduce on 14th May and finally return to normal on 18th May. The jobs sheet contains no comment as to whether a blockage was removed on any occasion.

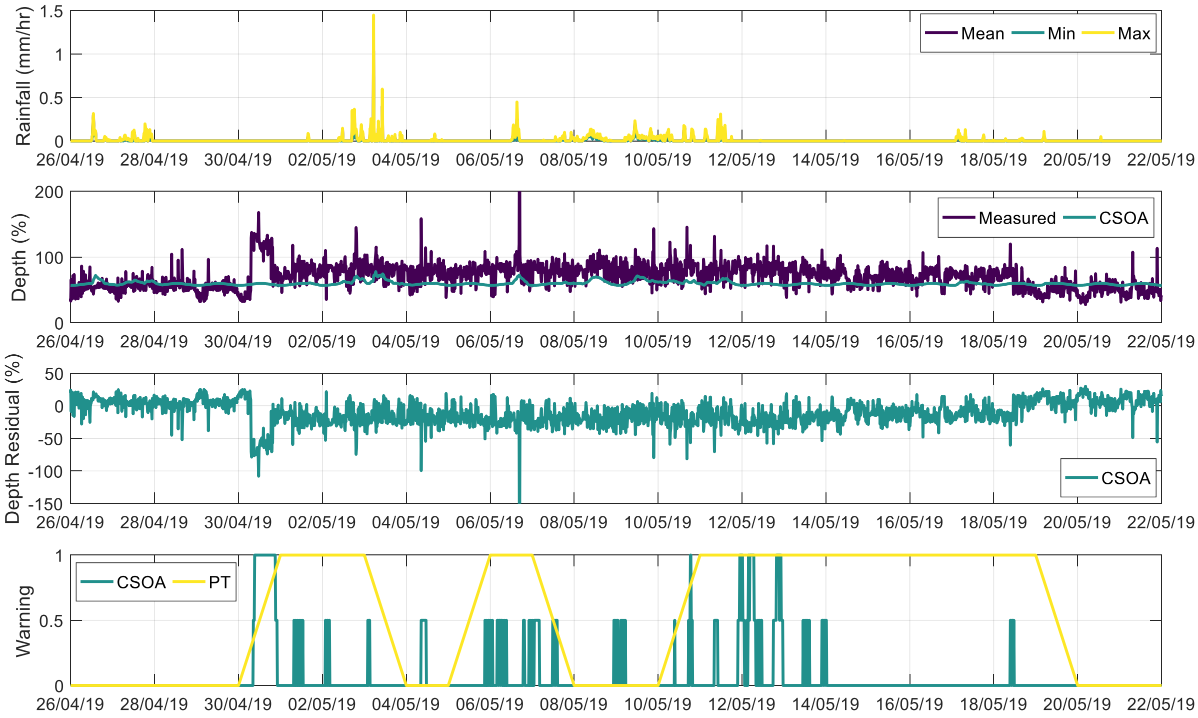

Fig. 5: Analysis of a longer period of high depth at CSO 10 during April/May 2019 showing rainfall, resulting depth and the response of both PT and CSOA analytics.

Fig. 6 is an example of an alert occurring and coinciding with rainfall at CSO 47. A sudden increase in depth is seen in the afternoon on 9th June, CSOA Alerts immediately and PT ranks the site as the 17th worst the following morning. There was however no job raised and the water depth returns to normal at around 10:00, it is likely that the water level had returned to normal before the operations team were able to visually check the data, having higher priority alerts at the time. The water depth remained below spill level, hence a pollution event did not occur, but it is highly likely that a partial blockage occurred, but this self-cleared during the morning high flows. Rainfall can be seen to occur overnight, during this rainfall, CSOA cannot be sure that the high levels are caused by a performance issue, hence the alert is suspended, but reappears immediately after the rainfall.

Fig. 6: Analysis of a probable self-clearing blockage at CSO 47 during June 2020 showing rainfall, resulting depth and the response of both PT and CSOA analytics.

Example per site daily analysis

The daily analysis takes the maximum value of the warning in each day in order to make longer periods of comparison feasible and able to be interpreted statistically. The coloured tables of data (cf. Fig. 7) have 5 rows for each day. The first row is the 'best available information' judgement as to whether an alert would reasonably be returned: green indicates that the CSO appears to be performing normally; amber indicates a probable performance or sensor issue; red is only returned if there is definitely a performance or sensor issue. Best available information is left blank if data is missing for a significant proportion of the day (at least 8 hours).

The next two rows summarise the alerts for the analytics. When an asset appears in PT (i.e. in the top 50), it is classed as a red warning on that day, otherwise it is green. The CSOA tool outputs warnings upon receipt of new data, these are red (alarm), amber (alert), or green (stable), consequently the highest severity warning occurring in the day is used, red being highest.

The final two rows are an evaluation of the analytics alert for PT and CSOA respectively, compared to the best available information. This has five possible outcomes: 1) True negative, coloured light blue, the best available information and the analytics agree that the status of the CSO is normal; 2) False negative, coloured red, the analytics returns a normal status (green), but the best available information is red or amber; 3) True positive, coloured green, the best available information and the analytics agree that the status of the CSO is a red or amber warning; 4) False positive, coloured amber, the analytics returns an abnormal (i.e. amber or red) status, but the best available information suggests the CSO is performing normally (green); 5) there is no data, the cell is not coloured, the analytics will return a value because pre-processing has interpolated missing data.

Finally, a table of summary statistics is provided, this includes data for the CSO for the whole period. Four percentages are given: True Positive rate (TP, also termed sensitivity); False Negative rate (FN); True Negative rate (TN, also termed specificity); False Positive rate (FP). TP and FN results are split into 'All Positives', where both amber alerts and red alarms are considered a positive; and 'Confirmed (Conf.) Positives' where only red alarms are considered positives, amber alerts are excluded as being neither conclusively positive nor negative.

Fig. 7 and Table 1 show results for CSO 20, here PT rarely alerts and predictability is low as the false positive rate is almost 50% of the true positive rate for all positives, performance improves for confirmed positives. CSOA identifies more higher than normal depths than Pollution Tracker, although performance in identifying all positives is far from perfect, mainly due to a long 'amber' event in 2018 where CSOA only alerts for a small proportion of the time, however CSOA maintains a low false positive rate. The performance increases for both PT and CSOA when only confirmed positives are considered, this would be expected as confirmed positives are likely to have a stronger response than unconfirmed ones.

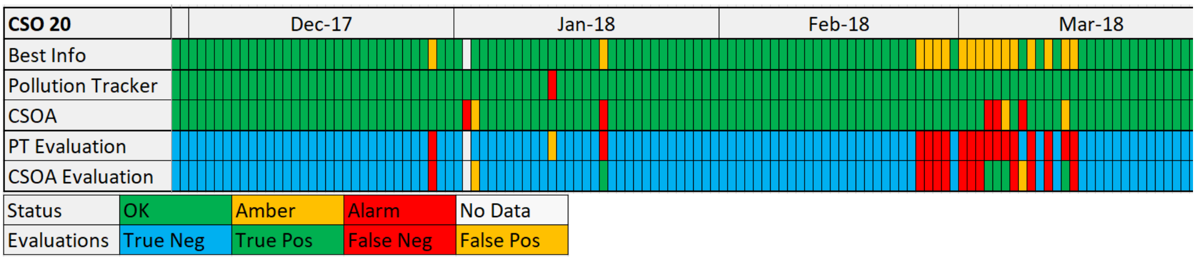

Fig. 7: Daily summary of best available information, status prediction from PT and CSOA, and evaluation of PT and CSOA predictions for CSO 20.

Table 1: Whole period summary of data and analytics evaluations for CSO 20.

	Days of data	All Positives	Conf. Positives	All Positives		Confirmed Positives		Negatives	
				TP(%)	FN(%)	TP(%)	FN(%)	TN(%)	FP(%)
PT	795	25	3	4.0%	96.0%	33.3%	66.7%	98.2%	1.8%
CSOA	795	25	3	32.0%	68.0%	100.0%	0.0%	97.7%	2.3%

Fig. 8 shows the data for CSO 5, it can be seen that this site never appears on PT, confirmed in Table 2 where there are zero true or false positives. CSOA does produce warnings and there is a reasonably good correlation between these and higher than expected water depths. Table 2 shows a reduced TP performance for confirmed positives, this is due to the majority of these occurring in the 6 month pilot in a period where the model required retraining due to a clear change in response of the CSO or in the monitor calibration, thus the magnitude of the residuals was too small to cause an alert. Table 2 also shows that there were issues with this monitor as significantly less data was available through the historic and live periods than other sites.

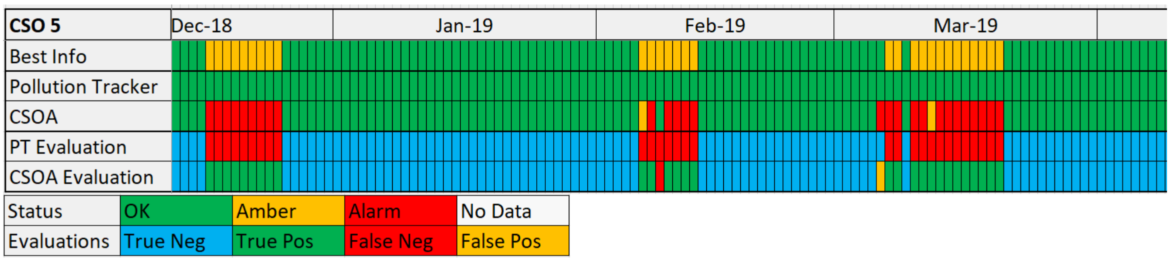

Fig. 8: Daily summary of best available information, status prediction from PT and CSOA, and evaluation of PT and CSOA predictions for CSO 5.

Table 2: Whole period summary of data and analytics evaluations for CSO 5.

	Days of data	All Positives	Conf. Positives	All Positives		Confirmed Positives		Negatives	
				TP(%)	FN(%)	TP(%)	FN(%)	TN(%)	FP(%)
PT	306	83	14	0.0%	100.0%	0.0%	100.0%	100.0%	0.0%
CSOA	306	83	14	57.8%	42.2%	21.4%	78.6%	92.8%	7.2%

Fig. 9 and Table 3 show the performance for CSO 15, which has one of the highest amounts of data. CSOA performs well in predicting performance issues. In this example, PT returns 7.5% false positives, as can be seen in Fig. 9, while CSOA only returns 1.7% false positives and correctly identifies more than twice as many positives in both categories.

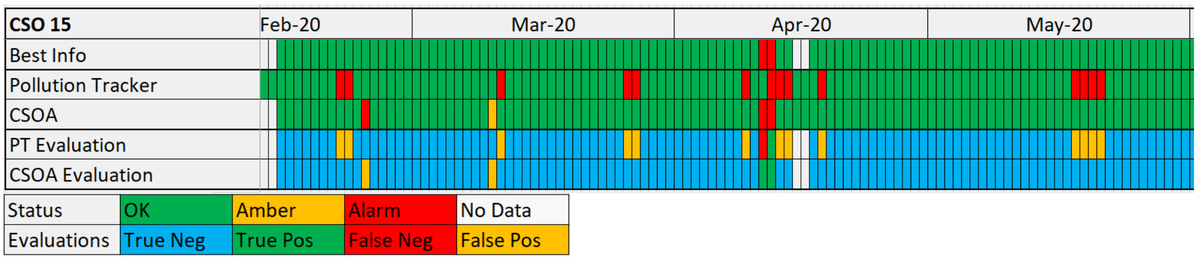

Fig. 9: Daily summary of best available information, status prediction from PT and CSOA, and evaluation of PT and CSOA predictions for CSO 15.

Table 3: Whole period summary of data and analytics evaluations for CSO 15.

	Days of data	All Positives	Conf. Positives	All Positives		Confirmed Positives		Negatives	
				TP(%)	FN(%)	TP(%)	FN(%)	TN(%)	FP(%)
PT	889	14	7	28.6%	71.4%	42.9%	57.1%	92.5%	7.5%
CSOA	889	14	7	85.7%	14.3%	100.0%	0.0%	98.3%	1.7%

Summary of performance across all sites

Next, information is summarised across all sites for both the historical and live periods. Distinction is made between the two continuous time periods – the 'historic' data from July 2017 to July 2019, and the 'live' data from late January to July 2020. Table 4 shows the total days of data and reveals that there are a higher proportion of confirmed positives in the 2020 data, this is due to a 4 month period where one site was continuously positive, and thus accounts for 126 confirmed positives. Table 4 also highlights the difficulty of being able to confidently confirm a performance issue, with the confirmed positive count only accounting for 15% of all positives. Table 5 gives a comprehensive overview of the performance of both analytics solutions. The CSOA tool provides strong performance across all data, with a True Positive rate of 68.2% or 88.4% for confirmed positives only, and a False Positive rate of 4.3%. This high degree of precision in the positive warnings is important for trust in the analytical solution. In contrast, PT only correctly identifies 16.0% of unexpectedly high depths (True Positives) within the data set, improving to 26.6% if only confirmed positives are considered. PT's False Positive rate is also higher at 5.4%.

Table 4: Summary of data for sites across both historical and live data periods.

	Total days of data	Missing data days	'Best Info' All positives	'Best Info' Confirmed positives
2017-19	28429	702	2101	221
2020	7238	632	551	177
Total	35667	1334	2652	398

Table 5: Summary of analytics performance for all sites and periods, both all positives and confirmed (conf.) positives.

	2017-19 data				2020 data				All data			
	TP (%)	TN (%)	FP (%)	FN (%)	TP (%)	TN (%)	FP (%)	FN (%)	TP (%)	TN (%)	FP (%)	FN (%)
PT All	17.1	94.6	5.4	82.9	11.6	94.4	5.6	88.4	16.0	94.6	5.4	84.0
CSOA All	65.9	96.0	4.0	34.1	77.2	94.5	5.5	22.8	68.2	95.7	4.3	31.8
PT Conf.	39.8	94.6	5.4	60.2	10.2	94.4	5.6	89.8	26.6	94.6	5.4	73.4
CSOA Conf.	89.6	96.0	4.0	10.4	86.9	94.5	5.5	13.1	88.4	95.7	4.3	11.6

DISCUSSION

Most sewerage systems are not managed in real time. They have little or no facilities for warning of service failure before it has impacted on customers and/or the environment. The innovation presented here enables a step change in this advance warning of the degradation of CSO performance which could result in a premature spill or pollution event.

While no direct comparison is conducted, it appears that performance is comparable to similar models in the literature, but achieved with an ANN architecture that has been developed to be transferable between CSOs and hence appropriate for applying at scale, rather than focussing only on predictive performance. The benefits of the architecture also include a relatively short training dataset to give acceptably accurate predictions. A short training dataset has benefits when the catchment response changes significantly (e.g. additional upstream urbanisation), or a monitor calibration changes for any reason (e.g. when the monitor or CSO is maintained, the head is either inadvertently moved, or is repositioned differently), then it is possible to retrain the ANN after only 2 or 3 months (depending on the number and range of events included). In comparison, deeply structured ANN models (e.g. Wu et al. 2015), may be able to better represent the performance of the CSO under many more conditions such as delayed response inflows in wet winter periods. However, these would require the training dataset to include a much wider range of conditions and events, which would inevitably mean a longer training period of potentially many years of data.

The period of historic data required to train the ANN is difficult to precisely define as it depends on the available data, in particular the number and range of rainfall events (and also climate variability is a factor). The data selected for training did not need to be continuous, but was selected based on data quality and consistency in order to include a variety of different rainfall events as well as periods of dry weather. With sufficiently long periods of data and automated data quality assessment, sites could be retrained at regular intervals with slightly updated data sets over time as has been implemented in other applications areas, for example water main burst detection (Mounce et al., 2010).

Supervised learning techniques require datasets where examples of good and degraded performance are clearly and accurately labelled. As with most water network applications, such labelling is not routinely carried out and it is not feasible to do so accurately with large historic datasets. Unsupervised AI techniques are therefore appropriate to apply to such datasets. A key example of this was the fuzzy inference system developed and validated here. The FIS captured expert judgement and assessment of residuals between measured and predicted depth data. Unlike the expert, the AI system enables repetitive error free operation at scale. The system is shown to capture a sufficient degree of the expert behaviour to yield high true positive and low false negative classifications.

The performance of the FL based control systems is a function of the rules and Membership Functions (MF) which in effect capture expert appreciation of system operation which are explicit and more human readable than opaque black box solutions, particularly for waste water. In Ostojin et al. (2011), an automatic control methodology for sewer pumping stations in dry weather conditions utilising FL was proposed. Simulation results indicated that cost savings of around 5% were achievable and that the number of pump runs was reduced by 20%. In Mounce et al. (2020) a Genetic Algorithm software tool was coded to optimise a FL control system which uses local water-depth sensing and a flow control gate to adjust the spatial distribution of the in-pipe water volume to reduce the local flood risk. The optimised FL MFs result in an average 25% decrease in the flood volume compared to those selected by experts for unseen rainfall events.

The UK WSP partner with whom this work has been carried out is a market leader, having been installing water depth monitors in CSOs for almost two decades. The WSP now has the majority of their CSOs monitored with data transferred to a central database by telemetry on at least a daily basis, or when a set alarm depth is breached. The system presented was designed and implemented on a cloud based architecture to take advantage of such data. Deployment on 50 assets for a 6 month demonstration showed that the analytics architecture could be effectively transferred to cloud operation enabling rapid scale up of deployment. Following the successful validation of the pilot, the WSP has conducted a full roll-out of CSOA, with over 2000 assets having an AI model deployed as of 2023. The close relationships with control room and operational functions within the water utility were vital in building trust and acceptance of the system, as well as providing the high level of ‘ground truth’ to the events that are detected. Thus the validation exercise has made it possible to understand the potential impact/benefit of integrating the tool into daily operations and replacing existing processes.

CSOA alerts are produced from either online or from archived data. CSOA has the potential to provide updates on a sub-daily basis when data is available, whereas the PT ranking methodology is designed to use daily data. The data from the sensors goes into the YW database and then on into MindSphere at a frequency of at least once a day. A shorter time window is used by CSOA for raising warnings, this means that a warning can be raised far more quickly after a change in performance, however it can also result in inconsistent warnings if depths fluctuate. The sensitivity of CSOA alerts is adjustable, overall this is a useful feature because different CSOs react in different ways and a fixed sensitivity can result in wrongly produced warnings, conversely if the sensitivity is too low then warnings might not appear when they should, or appear later and hence provide less time for a blockage to be cleared.

CONCLUSIONS

This paper demonstrates how cloud based analytics can be applied to transform data from on-line CSO depth monitors combined with rainfall radar data into information about how the sewer system is performing.

The system is based on a hybrid artificial neural network and fuzzy logic approach that enables incorporation of rainfall data. This allows expected high water depths to be ignored and is key to the predictive capability for identifying unexpected high water depths which could indicate a blockage which could result in a pollution incident. Examples are presented showing that the new system is resilient to wet weather, allowing performance deterioration to be identified during minor rainfall or shortly after more significant rainfall.

The ANN architecture was a one layer feed- forward structure which is transferable (i.e. the ANN architecture was not site specific) and has a comparatively low data requirement for training. Each asset has multiple models capturing different spatial and temporal windows of rainfall data, a grid search was used to identify the best structure as regards to ANN inputs. This approach was found to sufficiently capture the response to rainfall across a large number of real world locations.

The AI tool performance results for 50 assets over a 6-month live trial when high quality confirmatory analysis was possible show CSOA correctly identifying 86.9% of confirmed positives with only 5.5% false positives. This is significantly better than the legacy system with 10.2% and 5.6% respectively.

The system can deliver tangible benefits by producing reliable advance warnings of adverse performance of a CSO which may lead to a pollution event. These advance warnings mean that maintenance can be scheduled and prioritised in order to prevent the problem from escalating into a pollution incident. The cloud based architecture readily enables scaling, with the system now being rolled out across the water utility, with over 2000 assets on board.

The system presented here will help enable improved delivery of water services from existing infrastructure in the face of climate change, population growth and increasing urbanisation, without the need to build more infrastructure:

- The AI tool provided performance equal to or superior than (especially during wet weather) an existing analysis system based only on moving averages of depths.
- The system is resilient to wet weather, allowing performance deterioration to be identified during minor rainfall or shortly after more significant rainfall.
- The system is largely resilient to annual fluctuations in dry weather flow depths.
- The system can deliver tangible benefits by producing reliable warnings of decreases in performance of a CSO which may lead to a pollution event. By receiving these warnings as early as possible, maintenance can be scheduled and prioritised in order to prevent the problem from escalating into a pollution incident.

DATA AVAILABILITY STATEMENT

Some or all data, models, or code generated or used during the study are proprietary or confidential in nature and may only be provided with restrictions (e.g., anonymized data). Data supplied by the WSP is subject to a non-disclosure agreement; the code is commercially confidential.

ACKNOWLEDGEMENTS

The authors are grateful for the support, access to data and funding provided by Siemens plc and Yorkshire Water Services Ltd. and their permission to publish the details included herein.

REFERENCES

Ahm, M., Thorndahl, S., Nielsen, J. E. and Rasmussen, M. R. (2016). Estimation of combined sewer overflow discharge: a software sensor approach based on local water level measurements. Water Science & Technology, 74 (11), pp. 2683-2696.

Bachmann-Machnik, A., Bruning, Y., Bakhshipour, A. E., Krauss, M., and Dittmer, U. (2021). Evaluation of Combined Sewer System Operation Strategies Based on Highly Resolved Online Data. Water 2021, 13, 751. https://doi.org/10.3390/w13060751.

Bailey, J., E. Harris, E. Keedwell, S. Djordjevic, and Z. Kapelan. (2016). "The Use of Telemetry Data for the Identification of Issues at Combined Sewer Overflows." Procedia Engineering, 12th International Conference on Hydroinformatics (HIC 2016) - Smart Water for the Future, 154: 1201–1208. https://doi.org/10.1016/j.proeng.2016.07.524.

Benyon, R. (2013). Letter to Water Company Chief Executives. Available at: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/364435/letter_2013_07_18_RB_to_CEOs_-_CSO_spills__2_.pdf [Accessed 15th October 2021]

Bizer, M. A., and Kirchhoff, C. J. (2022). Regression modeling of combined sewer overflows to assess system performance. Water Science & Technology, 86 (11), 2848-2860, doi: 10.2166/wst.2022.362.

Botturi, A., E.G. Ozbayram, K. Tondera, N.I. Gilbert, P. Rouault, N. Caradot, O. Gutierrez, S. Daneshgar, N. Frison, Ç. Akyol, A. Foglia, A.L. Eusebi & F. Fatone (2021). Combined sewer overflows: A critical review on best practice and innovative solutions to mitigate impacts on environment and human health, Critical Reviews in Environmental Science and Technology, 51:15, 1585-1618, DOI: 10.1080/10643389.2020.1757957

Council Directive (EC) 2000/60/EC of 23 October 2000. Establishing A Framework For Community Action In The Field Of Water Policy. Available from:

http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=CONSLEG:2000L0060:20011216:EN:PDF.

Duncan, A. P., Chen, A. S., Keedwell, E. C., Djordjević, S. and Savić, D. A. (2013). RAPIDS: Early Warning System for Urban Flooding and Water Quality Hazards. MaLWaS Sympsium, AISB-IACA conference, University of Exeter, April 2013. ISBN: 978-1-908187-33-8.

EU (2022).

https://ec.europa.eu/environment/water/water-urbanwaste/evaluation/index_en.htm

Fernando, A. K., Zhang, X. and Kinley, P. F. (2006). Combined sewer overflow forecasting with feed-forward, back-propagation artificial neural network. Transactions on Engineering, Computing and Technology V12, ISSN 1305-5313, 58–64.

Fortier, C. and Mailhot, A. (2015). Climate Change Impact on Combined Sewer Overflows. ASCE Water Resour. Plann. Manage. 141 (5), 1-7, DOI:10.1061/(ASCE)WR.1943-5452.0000468.

Garcia, L. Barreiro-Gomez, J., Escobar E., Téllez, D., Quijano, N. amd Ocampo-Martinez, C. (2015). “Modeling and real-time control of urban drainage systems: A review.” Advances in Water Resources, 85, 120-132.

Guo, N. and Saul, A.J. (2011). Improving the operation and maintenance of CSO structures. In Proceedings of 12th International Conference on Urban Drainage, Porto Alegre, Brazil.

Hornick, K., Maxwell, S. & Halbert, W. (1989) Multilayer feedforward networks are universal approximators. Neural Networks 2: 359-366.

Kanneganti, D., Reinersman, L.E., Holm, R.H., Smith, T., (2022). Estimating sewage flow rate in Jefferson County, Kentucky, using machine learning for wastewater-based epidemiology applications. Water Supply 22, 8434–8439. DOI: 10.2166/ws.2022.395

Kerkez, B., Gruden, C., Lewis, M., Montestruque, L., Quigley, M., Wong, B., Bedig, A., Kertesz, R., Braun, T., Cadwalader, O., Poresky, A. and Pak, C. (2016). Smarter Stormwater Systems. Environmental Science & Technology, vol. 50 (14), pp. 7267-7273. DOI: 10.1021/acs.est.5b05870

Kurth, A., Saul, A., Mounce, S.R., Shepherd, W and Hanson, D. (2008). Application of Artificial Neural Networks (ANNs) for the prediction of CSO discharges. In Proceedings of 11th International Conference on Urban Drainage, Edinburgh, Scotland 31 August-5 September.

LeCun, Y., Bengio, Y. and Hinton, G. (2015) “Deep learning,” Nature, vol. 521, pp. 436-444.

Lingireddy S and Brion GM (2005) Artificial Neural Networks in Water Supply Engineering. ASCE.

Lund, N. S. V., Falk, A. K. V., Borup, M., Madsen, H. and Mikkelsen, P. S. (2018). Model predictive control of urban drainage systems: A review and perspective towards smart real-time water management, Critical Reviews in Environmental Science and Technology, 1-61, DOI: 10.1080/10643389.2018.1455484.

Mamdani, E.H. and Assilian, S. (1975). An experiment in linguistic synthesis with a fuzzy logic controller. International Journal of Man-Machine Studies, 7(1), 1-13.

Mounce, S. R., Boxall, J. B. and Machell, J. (2010). Development and Verification of an Online Artificial Intelligence System for Burst Detection in Water Distribution Systems. ASCE Water Resour. Plann. Manage. 309, 309–318.

Mounce, S. R., Shepherd, W., Sailor, G., Shucksmith, J. and Saul, A. J. (2014a). Predicting CSO chamber depth using Artificial Neural Networks with rainfall radar data. IWA Water Science and Technology. Vol 69 (6), pp. 1326-1333.

Mounce, S. R., Shepherd, W., Sailor, G., Saul, A. and Boxall, J. B. (2014b). Application of Artificial Neural Networks to Assess CSO Performance. 13th International Conference on Urban Drainage, Sarawak, Malaysia, 7th-12th September, 2014.

Mounce, S. R., Shepherd, W., Ostojin, S., Abdel-Aal, M., Schellart, A., Shucksmith, J., and Tait, S. (2020). Optimisation of a fuzzy logic based local real-time control system for mitigation of sewer flooding using genetic algorithms. IWA Journal of HydroInformatics. Vol 22 (2), pp. 281-295. https://doi.org/10.2166/hydro.2019.058.

Oftwat (2019). https://www.ofwat.gov.uk/pn-16-19-2019-price-review-ofwat-unveils-programme-of-huge-investment-service-improvements-and-lower-bills-for-water-customers/

Ostojin, S., Mounce, S. R. and Boxall, J. B. (2011). An artificial intelligence approach for optimising pumping in sewer systems. Journal of HydroInformatics, vol. 13, no. 3, pp. 295-306.

Rosin, T., Romano, M., Woodward, K., Keedwell, E., and Kapelan, Z. (2018). Prediction of CSO Chamber Level using Evolutionary Artificial Neural Networks. Proc. 13th International Conference on Hydroinformatics, Palermo, Italy.

Rosin, T.,Kapelan, Z., Keedwell, E. and 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, Vol 24 (2), pp. 259-273.

Rosin, T., Romano, M., Keedwell, E., and Kapelan, Z. (2021). A Committee Evolutionary Neural Network for the Prediction of Combined Sewer Overflows. Water Resources Management, Vol (35), pp. 1273-1289. https://doi.org/10.1007/s11269-021-02780-z

Sumer, D., Gonzalez, J., Lansey, K. (2007). Real-Time Detection of Sanitary Sewer Overflows Using Neural Networks and Time Series Analysis, Journal of Environmental Engineering, Vol. 133, No. 4, pp. 353-363.

Van der Werf, J. A., Kapelan, Z. and Langeveld, J. (2022). Towards the long term implementation of real time control of combined sewer systems: a review of performance and influencing factors. Water Science and Technology, Vol.85, pp. 1295-1320. doi: 10.2166/wst.2022.038.

Wu, Z. Y., El-Maghraby, M. E. and Pathak, S. (2015). Applications of deep learning for smart water networks. CCWI 2015. Procedia Engineering 119 (2015) 479 – 485.

Zhang, D., Holland, E. S., Lindholm, G. and Ratnaweera, H. (2018a). Hydraulic modeling and deep learning based flow forecasting for optimizing inter catchment wastewater transfer. Journal of Hydrology. Vol. 567, pp. 792-802.

Zhang, D., Lindholm, G. and Ratnaweera, H. (2018b). Use long short-term memory to enhance Internet of Things for combined sewer overflow monitoring. Journal of Hydrology. Vol. 556, pp. 409-418.