Stochastic Resonance Enhancement for Leak Detection in Pipelines Using Fluid Transients and Convolutional Neural Networks

Jessica Bohorquez[1]; Martin F. Lambert, M.ASCE[2]; Bradley Alexander[3]; Angus R. Simpson, M.ASCE[4]; and Derek Abbott[5]

Abstract: Water losses through leakage represent a significant problem for asset management in water distribution systems. The interpretation of fluid transient pressure waves after the generation of a transient event has been previously used as a technique to locate and characterize leaks, but existing approaches are often both model-driven and limited to the existing knowledge of the system. The potential of using artificial neural networks (ANN) and fluid transient waves to detect, locate, and characterize anomalies in water pipelines has recently been proposed. However, its application in more realistic conditions (e.g., in the presence of background pressure fluctuations) has proven challenging. To address this, one alternative to enhance the response of any nonlinear system includes the introduction of artificial noise, a phenomenon known as stochastic resonance. In this paper, the enhanced detection of leaks in pressurized pipelines via the deployment of stochastic resonance is demonstrated. This paper harnesses this approach by presenting a methodology for the active inspection of pipelines using convolutional neural networks (CNNs). This methodology finds the optimal artificial noise intensity to be introduced into the training dataset for a set of CNNs. The methodology has been applied to a real pipeline in a laboratory at the University of Adelaide in which 14 transient experimental tests were conducted. The results indicated that the addition of noise to the transient pressure head training samples significantly enhances the CNN predictions for the leak location highlighting the existence of an optimum noise intensity to obtain both accurate and reliable results. When trained with the optimum noise intensity, the CNNs were able to locate leaks with an average error of 0.59% in terms of the actual location (in a 37.24-m long pipeline), demonstrating the promising potential of developing techniques based on CNNs to detect leaks and anomalies in water pipelines. **DOI: 10.1061/(ASCE)WR.1943-5452.0001504.**

Author keywords: Leak detection; Water pipelines; Fluid transients; Artificial neural networks (ANN); Stochastic resonance; Machine learning; Water distribution systems; Convolutional neural networks (CNNs).

Introduction

Population growth and urban expansion are a challenge for water distribution systems (WDSs) because these systems are responsible for the supply of a vital resource to society. In recent years, major cities have faced a serious water supply crisis (Ahmadi et al. 2020). One major challenge in addressing these crises is the detection of water losses and pipeline repairs, which has received attention considering that the percentage of water losses can reach values of 35% in cities such as Kolkata (Mukherjee et al. 2018). Different methodologies have been used to estimate, monitor, detect, and pinpoint the location of leaks as part of water loss management strategies (Mutikanga et al. 2013). One of these methodologies includes the use of fluid transients for leak detection that usually involves the generation of a transient event that travels along the pipeline, allowing its inspection in a way that is similar to the functioning of radar and sonar techniques (Puust et al. 2010).

Fluid transient-based techniques have proven successful in the detection, location, and characterization of leaks in pipelines using the information that can be retrieved from transient pressure data. However, in most cases, existing techniques are model-driven. Such model-driven approaches usually require extensive and accurate numerical modeling, a priori estimation of certain pipe parameters assuming an intact or original condition, or long processing times to obtain an estimate of the leak characteristics. These limitations motivate the need for data-driven techniques that can quickly interpret transient pressure data obtained from a test and locate leaks accurately. Existing literature indicated that artificial neural networks (ANNs) can be trained using numerically generated transient head pressure traces to locate leaks and changes in pipeline diameters in numerically modeled pipelines (Bohorquez et al. 2020). This framework demonstrated the potential of using numerical data to train ANNs; however, a need exists for a methodology to apply these principles to the active inspection of pipelines under more

[1]Postdoctoral Researcher, School of Civil, Environmental and Mining Engineering, Univ. of Adelaide, Adelaide, SA 5005, Australia (corresponding author). ORCID: https://orcid.org/0000-0001-5071-8676. Email: jessica.bohorquez@adelaide.edu.au

[2]Professor, School of Civil, Environmental and Mining Engineering, Univ. of Adelaide, Adelaide, SA 5005, Australia. ORCID: https://orcid.org/0000-0001-8272-6697. Email: martin.lambert@adelaide.edu.au

[3]Senior Lecturer, School of Computer Science, Univ. of Adelaide, Adelaide, SA 5005, Australia. Email: bradley.alexander@adelaide.edu.au

[4]Professor Emeritus, School of Civil, Environmental and Mining Engineering, Univ. of Adelaide, Adelaide, SA 5005, Australia; Honorary Professorial Fellow, Dept. of Infrastructure Engineering, Univ. of Melbourne, Parkville, VIC 3010, Australia. Email: angus.simpson@adelaide.edu.au

[5]Professor, School of Electrical and Electronic Engineering, Univ. of Adelaide, Adelaide, SA 5005, Australia. Email: derek.abbott@adelaide.edu.au

Note. This manuscript was submitted on March 9, 2021; approved on October 7, 2021; published online on January 4, 2022. Discussion period open until June 4, 2022; separate discussions must be submitted for individual papers. This paper is part of the ***Journal of Water Resources Planning and Management***, © ASCE, ISSN 0733-9496.

realistic conditions. In addition, Bohorquez et al. (2021) presented a methodology for the passive inspection of pipelines to detect the occurrence of bursts. This methodology proved to be successful in identifying the occurrence, locating and sizing a burst by interpreting the incoming transient pressure signal measured in a pipeline. Although this methodology has indicated that ANNs can be trained to detect bursts in laboratory pipelines, no methodology has been proposed for the active inspection of pipelines experiencing background pressure fluctuations.

This paper presents a methodology for the active inspection of pipelines as an important contribution to the development of a general technique to exploit convolutional neural networks (CNNs) for leak detection in pipelines using fluid transients. This methodology indicates that the performance of CNNs for the detection and location of leaks in water pipelines is enhanced via the deployment of stochastic resonance. In the following, we describe the background related to transient-based methods for leak detection focusing on recent applications. A summary of the new methodology for the detection of leaks in pipelines under more realistic conditions using a set of CNNs trained with different noise intensities is then presented. This methodology is split into two main stages: model development and model application. Finally, the proposed methodology is applied to a series of experimental transient tests conducted in a laboratory setting. The results demonstrate that, by training a group of CNNs with pressure transient traces with the optimum noise intensity, the accuracy of the leak location predictions can be significantly enhanced, thus providing more robust predictions.

Background

Transient-based leak detection techniques have been in development for more than two decades (Jönsson and Larson 1992; Liggett and Chen 1994). Different approaches have been explored and can be classified into five main groups: reflection-based methods, damping-based methods, frequency response-based methods, inverse transient techniques, and signal processing-based methods (Duan et al. 2020). More recently, frequency response methods have been combined with enumeration techniques for leak detection in pipelines with branches and loops by separating the effect of these known elements on the frequency response of the system and employing a GA-based optimization to find the leak characteristics (Duan 2017). This method has proven successful for a numerical application and has indicated the potential of transient-based methods for operation in more complex systems. Meniconi et al. (2019) examined the influence of the pipeline initial flow conditions on transient pressure traces after the generation of a transient event for the visual detection of a leak in the pipeline (as an example of a reflection-based method). This study concluded that, depending on the location of the transient generator device, the transient measured signal can be more sensitive to the initial conditions. If the generator is located close to the water source, the transient pressure signals obtained are almost indistinguishable. In contrast, locating the generator close to the end of the pipeline can produce different transient traces for the same leak in terms of the initial pressure increase (Meniconi et al. 2019).

Matched-field processing (MFP) has been explored as a frequency domain technique that can obtain satisfactory results even in noisy environments and has been developed for elastic (Wang and Ghidaoui 2018) and viscoelastic pipelines (Wang et al. 2019). The advantages of this method include its robustness under noisy conditions or uncertainty in the wave speed. It has proven effective in both numerical and experimental conditions when the noise has been assumed to be white noise with a zero-mean Gaussian distribution (Wang et al. 2019). Other recent applications have indicated that frequency-based techniques can detect leaks under realistic background noise scenarios using the paired impulse response function obtained from two measurement points in the system (Zeng et al. 2020). Although significant advances have been achieved in transient-based methods in recent years, most existing techniques still require testing under perfect conditions (without any leaks), detailed numerical modeling, significant computer resources, or extensive signal preprocessing. In addition, uncertainty in the parameters of a pipeline system to create numerical models and the presence of irregular pipeline anomalies (e.g., extended and irregular blockages) can make the anomaly detection process more challenging when applied to water pipelines in the field (Che et al. 2021).

A different group of techniques for leak detection in pipelines has proposed the use of machine learning algorithms to process the available information from a particular pipeline system. Some of these techniques have used surrogate features of the pipeline to predict the most likely location of a leak (Geem et al. 2007) or the remaining lifetime of a pipeline (Zangenehmadar and Moselhi 2016). However, more recent techniques have proposed the combined use of machine learning algorithms and hydraulic measurements in the pipeline. Romano et al. (2014) used different self-learning artificial intelligence, statistical analysis, and Bayesian inference tools to detect a burst at a DMA level in real water distribution systems using wavelets for the denoising of the obtained signal before its analysis. Roy (2017) proposed the use of pressure fluctuations with hybrid dense ANNs to locate leaks by classifying the status of the system to characterize a normal and abnormal condition in the pipeline. Mujtaba et al. (2020) introduced the use of adaptive thresholds to detect the occurrence of leaks in gas pipelines using pressure and mass measurements at the beginning of the pipeline as inputs for the machine learning model and the potential mass flows at the end of the pipeline as the output of the model for comparison with measured data.

Bohorquez et al. (2020) presented a methodology that uses the transient pressure trace after the generation of a transient event and convolutional neural networks (CNNs) to determine the location and size of a leak in a water pipeline. This merging of pressure transient traces and CNNs has been demonstrated in a numerical application and has great potential, given that this technique is data-driven and can provide immediate results for the characteristics of a given anomaly. Nonetheless, challenges can arise when pipelines under more realistic conditions are analyzed. Background pressure fluctuations due to system operations, such as changes in demand or unknown system components, are not reproduced by numerical models but affect the transient pressure traces.

To date, although no applications using noise to improve ANN performance have been reported in assessing the condition of water pipelines, several related strategies have been applied in other fields. Previous approaches have proposed the introduction of noise during the training in ANN training samples (Rifai et al. 2011), in activation functions (Ikemoto et al. 2018), in ANN weights (Goodfellow et al. 2016), or in the direction of update of ANN weights (Neelakantan et al. 2015). The most popular approach has been the introduction of noise directly into the ANN training samples to enhance model robustness and reduce overfitting (Bishop 1995). Rifai et al. (2011) demonstrated that the error of a multilayer perceptron for document recognition can be reduced by adding a Gaussian distributed noise in the input layer regardless of the standard deviation of the noise. Fukami et al. (2020) applied the same concept to different ANN architectures that were trained to estimate laminar wakes in a fluid field from limited measurements, demonstrating that the addition of noise in the training

samples for the analyzed architectures improved the performance of the ANNs when tested in noisy input measurement environments. Nonetheless, if the magnitude of the noise deviation was too large, the ANN performance was compromised.

These past applications demonstrate the potential of the use of noise during the training of an ANN. However, few studies have identified the potential of using an organized framework to introduce noise in the training of an ANN. The phenomenon for which the performance of a nonlinear system (i.e., in this case, an ANN) is optimally enhanced by the addition of a certain noise intensity is known as stochastic resonance. The concept was proposed for the first time by Benzi et al. (1981) as a nonlinear cooperative effect in which periodic signals can be greatly amplified by large environmental fluctuations; however, other researchers rapidly extended this to include any nonperiodic signals (Collins et al. 1995). In a nonlinear system, it can be indicated that there exists a nonzero value of noise that gives an optimal response to the system (Harmer et al. 2002). Stochastic resonance has been observed and applied in multiple fields (Benzi et al. 1982; Luchinsky et al. 1999; Wang and Santamarina 2002; Cheng et al. 2020) with only a few applying it to the training of an ANN.

Ikemoto et al. (2018) introduced a noise-modulated neural network as an application of stochastic resonance by perturbing the threshold units in the activation functions with different noise intensities (described by its standard deviation). Their application in benchmark artificial problems demonstrated that by adding noise to the threshold units, the standard deviation of the mean squared error (MSE) decreased as the standard deviation of the noise increased regardless of the structure of the neural network in terms of the hidden units. The advantages of using stochastic resonance in areas related to the development of new technologies, such as signal processing (Feng et al. 2019) or time series analysis (Falanga et al. 2020), is an active research area. However, previous applications using stochastic resonance in ANNs have been limited to artificial and numerical benchmark problems in computer science, and no applications have been reported for anomaly detection problems in real infrastructure, such as the detection of leaks in water pipelines.

Methodology

The methodology presented in this research paper to detect leaks in pipelines using a set of CNNs is outlined in Fig. 1. This methodology is divided into two stages: model development (Stage 1 in Fig. 1) and model application (Stage 2 in Fig. 1). The leak detection model development stage should be carried out first and can be repeated regularly to account for new transient pressure information data collected from the system. The model application stage comprises the processes required to analyze a transient pressure head trace to determine the location and the size of a leak in the pipeline.

Leak Detection Model Development (Stage 1)

The first stage of the proposed methodology is the development of a leak detection model. This stage comprises the training of a set of CNNs that can locate and size a leak when the analyzed pipeline experiences background pressure fluctuations. An appropriate CNN architecture needs to be designed, and transient pressure head traces are numerically generated for the training of these CNNs. The five steps presented in Stage 1 in Fig. 1 summarize the development of the leak detection model. It is important to highlight that these CNNs do not constitute a metamodel of the transient flow pressure response to the closure of a valve in a pipeline with a leak. These CNNs are trained to identify the transient pressure wave reflections created by the existence of a leak in the pipeline.

CNN Architecture Design

The first step in the leak detection model development is the definition of an appropriate ANN architecture. Bohorquez et al. (2020) concluded that one-dimensional (1D) convolutional networks with three convolutional layers had the potential to identify leaks in numerically modeled pipelines. However, it has been found that a more robust architecture is required for the application of this technique in pipelines under more realistic conditions. The design of this new architecture considered different alternatives, including variations from the architecture proposed in Bohorquez et al. (2020) to a 1D convolutional network with 5 layers and increasing filters in those layers.

For brevity, the results from this design process are not indicated here, but the resulting CNN architecture includes (1) four convolutional layers, (2) the use of a leaky rectified linear unit (Leaky ReLU) as the activation function, (3) 20 filters that increase in the last convolutional layer, and (4) three dense layers of sizes 14, 6,

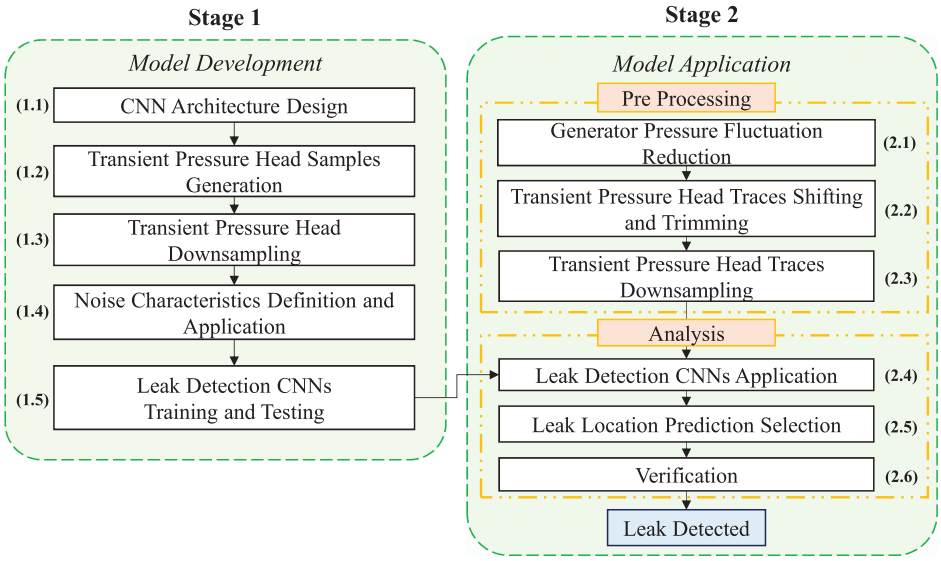

Fig. 1. (Color) Model development and application of active leak detection methodology.

J. Water Resour. Plann. Manage., 2022, 148(3): 04022001

and 2. The resulting number of weights for the 1D convolutional networks used in this work depends on the size and number of filters and the downsampling frequency selected in Step 1.3 in Fig. 1. Eq. (1) presents the total number of weights for a 1D convolutional network for which the first term represents the weights in the convolutional layers (n), and the second term represents the weights in the dense layers (j)

$$W = \sum_{1}^{n}[((w \times h \times f_{n-1}) + 1) \times f_n] + \sum_{1}^{j}[(c_j \times c_{j-1}) + c_j] \tag{1}$$

In this equation, w and h are the width and height of the filters; f_n = number of filters in the convolutional layer n; and c_j = number of neurons in the dense layer j. For the first dense layer (i.e., $j = 1$), c_{j-1} depends on the dimensions of the input layer defined by the downsampling frequency, thus affecting the total number of weights for the CNN to learn. In general, a larger input layer provides the CNN with more information regarding the transient pressure head trace, but the training of the CNN is more difficult because there are more weights to define.

Transient Pressure Head Samples Generation

The leak detection model development stage includes training and testing a set array of CNNs (Step 1.5 in Fig. 1). To train these CNNs, numerical transient pressure head data or available recorded transient data can be used. For the application presented in this paper, numerical transient head pressure traces have been used for the CNN training. Fig. 2 presents the hydraulic configuration of the single pipeline that has been used to generate the numerical transient head pressure traces. The pipeline is supplied by a reservoir with an upstream head H_0, an internal diameter D, and total length L_T. At the downstream end of the pipeline, there is a side discharge valve that is initially open with a flow Q_0. The transient pressure head information is obtained from a measurement point (M) at the same location as the side discharge valve. The specific characteristics of the pipeline that was analyzed in this paper are presented in Table 1.

A leak modeled as a circular orifice with a diameter D_L could be present at any point along this pipeline at a distance x from the upstream reservoir. As part of this step, a range of leak sizes must be defined for the generation of the training data. This range can be defined based only on the diameter of the leak, the flow that is going through the orifice when the leak is active, or previous knowledge of the system on past detected leaks. For the application indicated in this paper, the leak size range was defined using the available orifice diameters in the laboratory associated with the experimental pipeline apparatus and is presented in the "Results" section.

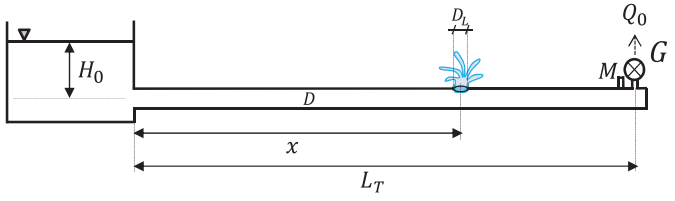

Fig. 2. (Color) Single pipeline with leak system configuration.

Table 1. Pipeline characteristics

Characteristic	Units	Value
Length of pipe (L_T)	m	37.24
Internal diameter of pipe (D)	mm	22.14
Wave speed of pipe (a)	m/s	1,305
Wall thickness (e)	mm	1.63
L_T/a time	s	0.029

Each sample of the CNN input dataset is a transient pressure head trace including the transient pressure response generated after the closure of the side discharge valve. A simulation time of at least $2L/a$ seconds after the valve closure is selected to cover the first set of transient wave reflections. To form the complete CNNs' training and testing dataset, a large number of transient pressure head traces are required (Bohorquez et al. 2020). A total of 50,000 different transient pressure head traces were generated using the method of characteristics (MOC). In ten different rounds, 5,000 locations along the complete pipeline are randomly sampled from fixed-length segments, and a random leak size is assigned.

Due to the short separation between different leak locations to cover the entire length obtaining 5,000 different transient pressure head traces, the time resolution required to guarantee that the MOC is computed along the characteristic lines is also very small. Therefore, the potential size for the training and testing dataset can be significant. For instance, for a pipeline that is 1,000 m long and has a wave speed of 1,000 m/s, generating 50,000 different transient pressure head traces for $2L/a$ seconds includes 500 million pressure head values. If these complete traces were used to train the CNNs, the total required parameters to train one CNN (following the architecture defined in Step 1.1) would be 454,000 according to Eq. (1). The potential size of this input dataset and its implications for the number of training weights indicates that a downsampling process is necessary.

For the application in a numerically modeled pipeline described in Bohorquez et al. (2020), the transient pressure head traces were obtained after modeling the sudden closure of a side discharge valve. However, a more realistic approach should consider that, regardless of the closure method (i.e., a mechanical actuator, a solenoid activated valve, or any other device), the injected transient wave is not completely sharp. As part of Step 1.2 in the leak detection model development, the closure curve of the side discharge valve is obtained and incorporated into the MOC modeling. This can be achieved by running preliminary tests in the analyzed pipeline to characterize this curve.

Transient Pressure Head Downsampling

As was previously mentioned, the potential size of the input dataset when using the MOC for the generation of the transient pressure head traces can be very large considering that all data points along the transient pressure head trace are required. Therefore, a timewise downsampling process is conducted in Step 1.3 of Fig. 1. Previous research indicated that ANNs trained with downsampled data have improved performance and are more computationally efficient (Bohorquez et al. 2020). This is partially because ANNs trained on downsampled data have a smaller size for the input layer for the ANNs, resulting in fewer weights and, thus, less of a tendency of overfitting. In addition, the use of downsampled data could potentially reduce data transfer requirements on applications of this methodology in the field. Depending on the analyzed pipeline, Step 1.3 includes the selection of the sampling frequency to which transient pressure head traces are going to be transformed into. This selected frequency will influence the final number of weights on which the CNNs will be trained because the frequency defines the size of the initial layer in the CNN. This frequency can be

selected based on the expected pressure sensor sampling frequency and the dimensions of the analyzed pipeline.

Noise Characteristics Definition and Application

Bohorquez et al. (2020) demonstrated the potential of using ANNs for leak detection in pipelines using fluid transient pressure waves. However, the application of this approach to pipelines under more realistic conditions, such as those containing background pressure fluctuations, has proven a challenge because the generated numerical pressure transient head traces cannot replicate these conditions. To address this issue, this paper introduces the deployment of stochastic resonance applied to the training of a set of CNNs. The addition of multiple noise intensities to the transient pressure head input dataset enhances the robustness and performance of the trained CNNs when the optimum noise intensity is applied. The addition of noise in training datasets has been explored in other numerical applications involving artificial intelligence in different fields. The addition of a noise distribution to the input of a model has been found to be translated into a better response of the model output (Murray and Edwards 1994; Rifai et al. 2011; Fukami et al. 2020). Considering this, Step 1.4 of the leak detection model development comprises the definition of the noise distribution and the selection of noise intensities to be added to the numerical transient pressure head traces samples (obtained in Step 1.3).

The transient pressure head noise has been characterized by a Gaussian distribution with zero mean and standard deviation of σ, similar to the concept presented by Duan (2017). The magnitude of the standard deviation has been defined with respect to the magnitude of the pressure drop in the transient pressure head trace when the smallest leak (from the range defined in Step 1.2) is present in the pipeline. To illustrate this, Fig. 3 presents a generic example of two transient pressure head traces.

The continuous blue line represents the transient pressure head trace after the closure of a side discharge valve for an intact pipeline. The dash-dotted blue line denotes the transient pressure head trace when a leak is present in the pipeline and for which a transient event has been generated. The initial pressure head increase after the closure of the valve is the same in both cases; however, differences arise when part of the transient wave reflects from the leak. The larger the leak present in the pipeline, the larger the drop in pressure (Bohorquez et al. 2018; Meniconi et al. 2019; Wang et al. 2019). A second y-axis is included on the right-hand side of Fig. 3 to present the differences between both transient pressure head traces (red line). This line indicates that a difference Δh exists during the first $2L/a$ seconds after the closure of the side discharge valve.

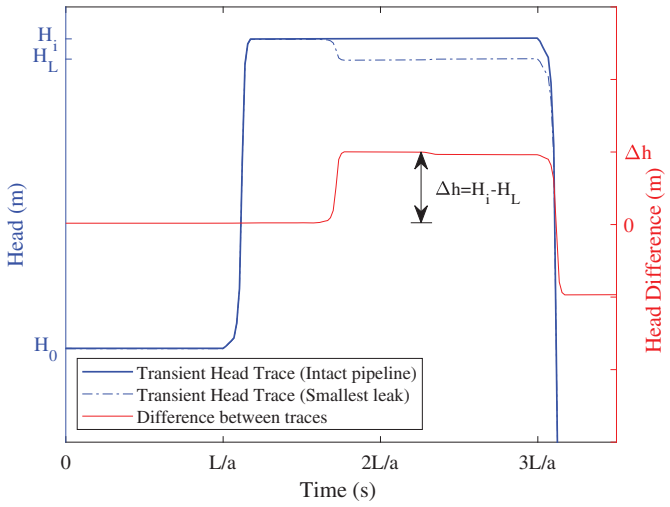

Fig. 3. (Color) Pressure head difference between transient pressure head trace for intact pipeline and pipeline with smallest leak.

Considering this, if the selected noise intensity has a standard deviation larger than this difference, the CNNs are expected to not perform well in identifying small leaks. In this case, the noise added to the transient pressure head trace hides the transient wave reflections from the leak. A total of n noise intensities are selected in Step 1.4, and the standard deviation for each intensity is defined in Eq. (2) as a proportion of Δh where k_i is the multiplier for noise intensity $i \in \{1, \ldots, n\}$ and k_i is a percentage of Δh

$$\sigma_i = k_i \times \Delta h \tag{2}$$

The selection of the number of noise intensities n depends on the computational resources available because each additional noise intensity represents more CNNs to be trained (as is explained in Step 1.5). The array of multiplier values $[k_1, \ldots, k_n]$ depends on the analyzed pipeline and the expected leak sizes to be present; however, if large values of k_i are selected, the ability of the CNN to identify certain leak sizes is expected to decrease.

Once the number of noise intensities and the array of standard deviations σ_i have been defined, multiple random transient pressure head traces are created. In this paper, five transient pressure head traces are created for each of the samples of the input dataset obtained from Step 1.3 by adding the noise to the original transient pressure head trace. This allows the CNNs to be exposed to different transient pressure head traces that correspond to the same leak location and size but with different values for the pressure noise. Thus, the input dataset for each noise intensity has 250,000 different transient pressure head samples.

Leak Detection CNNs' Training and Testing

The last step of the leak detection model development (Step 1.5 in Fig. 1) is the training and testing of a set of CNNs with the architecture defined in Step 1.1 using the $n + 1$ input datasets obtained from Step 1.4 (including the original dataset without any noise in the samples). A diagram presenting the set of CNNs to be trained is indicated in Fig. 4. Each leak detection CNN receives as input one transient pressure head trace and should be able to predict the correct location and size of the leak only based on this information.

This diagram indicates that for each noise intensity (σ_n) including a no-noise scenario (σ_0), m leak detection CNNs are trained using its corresponding dataset. Considering that the training of those CNNs is conducted using Stochastic Gradient Descent algorithms, a different final set of weights is obtained every time a CNN is trained. Similar to applying genetic algorithms starting from different random number seeds, training m leak detection CNNs can assist in testing the consistency of the CNN predictions. A good set of CNNs should provide very similar results when testing with the same data despite having different CNN weights. The number of possible CNNs to train for each noise intensity depends on the availability of computational resources.

Each input dataset is then randomly divided into two groups of equal size: a training and a testing dataset. For the training process, smaller groups of data are selected one at a time to find values for the CNN weights and then validated with the rest of the training data. This is known as batch training, which allows the CNN to learn from smaller groups of data to avoid overfitting (Nakama 2009). Considering that the separation of the input dataset into a training and a testing dataset is random and that batches for each training trial are different for each CNN, the resulting weights are different.

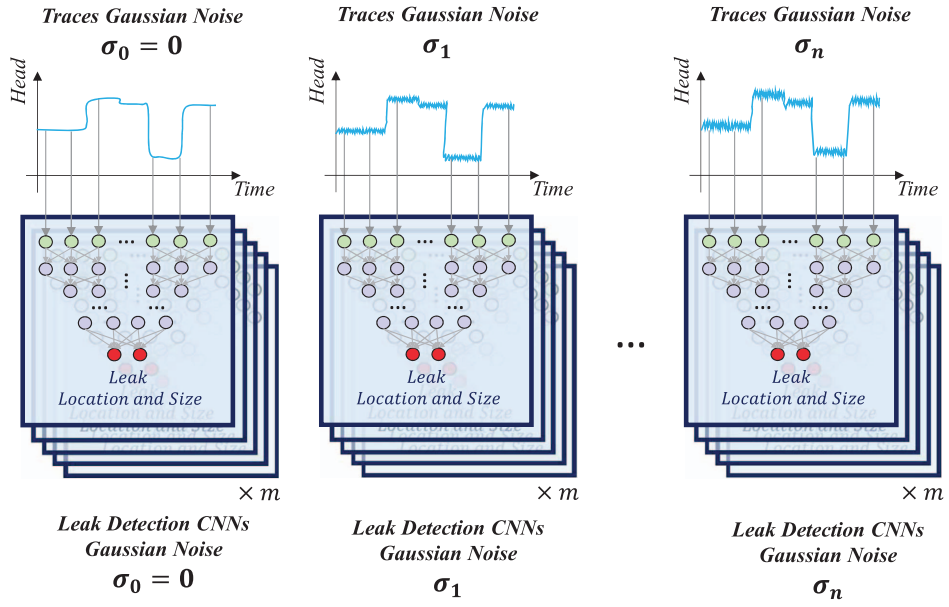

Fig. 4. (Color) Leak detection set of CNNs. Each group of CNNs is trained with samples with different noise intensities.

Once the training process is complete, the CNNs are tested with transient pressure head traces that have not been exposed to. These predictions are then compared with the real location and sizes of the testing samples. A CNN that has been successfully trained should present with a similar distribution of errors in the training and testing stages.

Leak Detection Model Application (Stage 2)

The second stage of the methodology presented in this paper is the leak detection model application (Stage 2 in Fig. 1). This stage includes a number of different steps that are necessary to process real measured transient pressure head data from a valve closure test in a pipeline with a leak to obtain a prediction of its location and size. A six-step process is described in Stage 2 in Fig. 1, which is divided into two substages: preprocessing and analysis. This section explains how each step might be carried out for any analyzed pipeline when the results from multiple valve closure tests are available. However, the same procedure could be conducted if only one test result is available.

Generator Pressure Fluctuation Reduction

The first part of the leak detection model application refers to the preprocessing of the tests that have been conducted. In this preprocessing stage, the first step corresponds to the reduction of pressure fluctuations caused by the transient generator device in the recorded transient pressure head signal (Step 2.1 in Fig. 1). Preliminary applications of the leak detection model, not indicated in the paper for brevity, demonstrated that the performance of the set of CNNs was not satisfactory for different transient tests under the same conditions. To understand the reasons for this apparent inconsistency in the leak detection model predictions, a *vulnerable region detection analysis* was conducted.

Vulnerable region detection analysis has been previously used in computer science to evaluate the performance of a classifier machine learning model to small perturbations in different regions of an image. This type of analysis indicated that deep neural networks are vulnerable to changes around the object of interest (Shu and Zhu 2019). Vulnerable region detection analysis is closely related to the study of adversarial examples for deep neural networks for which imperceptible perturbations (localized or distributed in the image) can disrupt the predictions of the models (Szegedy et al. 2013; Akhtar and Mian 2018).

For the CNNs developed in this paper, this analysis included the successive testing of the CNNs with perturbed transient pressure head traces. These traces were obtained by applying different magnitudes of perturbation to each point along the original transient pressure head trace. Fig. 5 presents the distribution of the predicted leak location error (at the top of Fig. 5) after successively applying a single 0.1 m perturbation in turn along a transient pressure head trace measured in the laboratory (indicated at the bottom of Fig. 5). The distribution of errors was obtained after testing the perturbed samples with five CNNs trained with the same noise intensity.

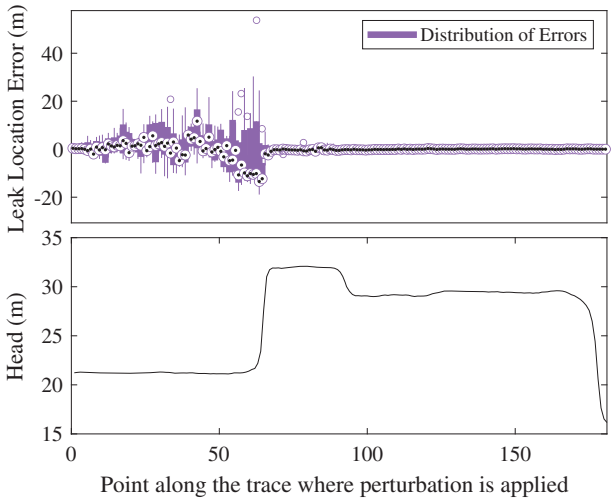

Fig. 5. (Color) Leak location error with application of 0.1 m perturbation at only one point along laboratory transient pressure head trace.

J. Water Resour. Plann. Manage., 2022, 148(3): 04022001

This figure indicates that perturbations at the first 60 points of the transient pressure head trace induce a considerably larger error in the distribution of the leak location predictions. These 60 points correspond to the steady-state pressure head before the valve closure. When the perturbation is applied after the valve closure, the CNN predictions are more consistent, although errors are present. This analysis demonstrated that features in the steady-state segment of a laboratory transient pressure head trace induce errors in the CNN performance. Thus, a more in-depth analysis of measured transient pressure head traces has been conducted.

Fig. 6(a) presents an example of a transient pressure head trace obtained in the laboratory (with characteristics presented in Table 1 and equivalent to the system presented in Fig. 2). In this figure, two segments of this trace are enlarged, preserving the same scale. Subplot (i) in Fig. 6(a) indicates the background pressure fluctuations before the valve closure, and subplot (ii) indicates the background pressure fluctuations after the dissipation of the transient event created by the valve closure. Clear differences in these pressure fluctuations are visible before and after the transient event. The background pressure fluctuations before the transient event are more prominent in magnitude, which is the result of the interaction that the transient generator (open valve) has with the pipeline itself that adds to the interaction that the leak orifice has with the pipeline. The background pressure fluctuations induced by the transient generator have been reported by Gong et al. (2018), for which leaks were simulated in real pipelines with the opening of a standpipe connected to an air valve or a fire hydrant. More recently, Brunone et al. (2021) analyzed the influence of transient pressure fluctuations at different locations along a pipeline before the introduction of a pressure wave in the minimum detectable reflected pressure wave due to the presence of a leak. Assessing the existing transient pressure fluctuations was indicated to be important to guarantee a successful leak detection outcome.

A background pressure fluctuation reduction step has been included to reduce the background pressure fluctuations caused by the combined effect of the transient generator and the leak to only the fluctuations due to the presence of the leak. To accomplish this, the distribution of pressure fluctuations needs to be studied in greater detail. Fig. 6(b) presents the distribution of the pressure head for the segments of the complete transient trace highlighted in Fig. 6(a). Step 2.1 in the proposed methodology includes fitting both pressure head series to a probabilistic distribution (in this case, a normal distribution), which in the examples presented indicate a reasonable agreement. The distribution of the pressure head after the transient test normally has a larger mean value because the total flow in the pipeline reduces. The parameters of the normal distribution for the background pressure fluctuation before the transient generator closure are denoted with the subscript b, and the fluctuations corresponding to the pressure after the transient generator closure are denoted with the subscript a.

The procedure for Step 2.1 consists of transforming the background pressure fluctuations *before* the transient event to have a similar distribution to the pressure fluctuations *after* the generator closure. Based on the parameters of the distributions obtained, a new probability function can be defined as a normal distribution with mean μ_b and standard deviation σ_a. Therefore, the new normal distribution preserves the mean value of the background pressure fluctuations before the transient test, but its standard deviation is modified to match the background pressure fluctuations after the transient test. For any pressure value from the measured transient pressure head trace before the transient event (h_t), a Z-score (Z_b) is computed using μ_b and σ_b. With this value, a modified value for the pressure (h_t') is found using Eq. (3). This is then repeated for each value of pressure before the generator closure to obtain a transient pressure head trace with reduced background pressure fluctuations

$$h_t' = Z_b \times \sigma_a + \mu_b \tag{3}$$

Transient Pressure Head Traces Shifting and Trimming

Once the background pressure fluctuations induced by the transient generator have been reduced, the transient pressure head traces are shifted and trimmed at Step 2.2 (Fig. 1). This is conducted to match the conditions used to generate the numerical transient pressure head samples in Step 1.2 of the model development stage. Vertical shifting of the transient pressure head traces obtained from measurements might also be required if the conducted tests had a different initial pressure. This shifting includes the computation of the difference between the mean pressure before the transient event and the steady-state pressure used in training the CNNs and the transformation of the transient pressure head traces by adding or subtracting this difference. If the steady-state pressure of the transient tests is significantly different, some variation in the transient response of the system can be expected, as reported by Meniconi et al. (2019). Trimming the time extent of the transient pressure head traces includes the selection of the length of interest from the complete trace. The specific length to trim the traces depends on the selected characteristics in Step 1.2; however, in general, the objective is to include some pressure information before the

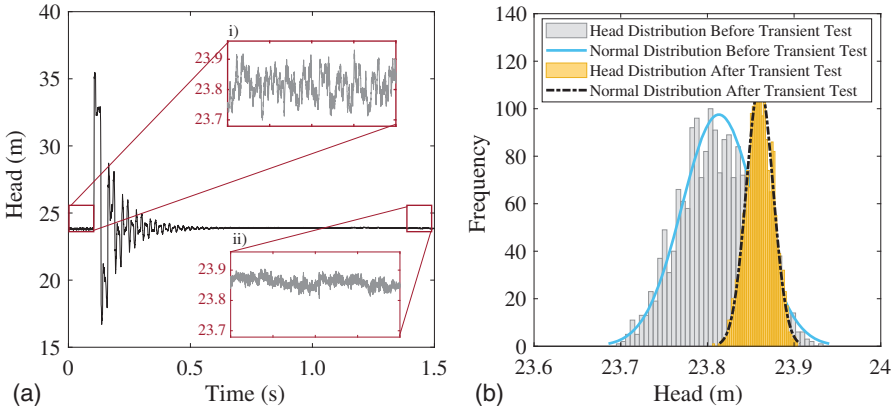

Fig. 6. (Color) Background pressure fluctuation analysis: (a) background pressure head fluctuation: (i) before transient event and (ii) after transient event; and (b) distribution of pressure head before and after transient event.

J. Water Resour. Plann. Manage., 2022, 148(3): 04022001

transient event and at least $2L/a$ seconds after the valve closure to cover the complete length of the pipeline.

Transient Pressure Head Traces Downsampling

To capture the reflections from anomalies such as leaks, high-frequency pressure transducers are required for realistic applications (Nguyen et al. 2018; Zeng et al. 2020). In addition, this sampling frequency depends on the pipeline dimensions and material. Step 2.3 in Fig. 1 includes the downsampling of the measured transient pressure head traces to the frequency selected in Step 1.3. This downsampling is necessary to match the measured transient pressure head trace with the input layer of the CNNs (as presented in Fig. 4). This downsampling process has been applied in different fields because signature recognition can be carried out using different interpolation methods including linear, polynomial, or spline interpolation (Martinez-Diaz et al. 2007).

Leak Detection CNNs' Application

The reduction in background pressure fluctuations, the shifting, trimming, and downsampling of the measured transient pressure head traces complete the preprocessing stage. A second stage comprises the following three steps: analysis of the measured transient pressure head traces using the available CNNs (Step 2.4 in Fig. 1), selection of a final prediction for the leak location and size (Step 2.5 in Fig. 1), and a verification process (Step 2.6 in Fig. 1).

The leak detection CNNs application step includes the analysis of all the recorded and preprocessed transient pressure head traces using the CNNs trained in Step 1.5. Therefore, multiple leak location and size predictions are obtained from this step depending on the number of noise intensities (n) selected in Step 1.4, the number of CNNs trained per noise intensity (m), and the number of available transient tests results (q). For each transient test (q), a box whisker plot can be created that summarizes the distribution of the predicted leak locations. This distribution is created from the results of m leak detection CNNs for each noise intensity (n) and the CNNs trained with transient pressure head samples without any added noise. The analysis of the transient tests through CNNs trained with different noise intensities constitutes the application of stochastic resonance. Therefore, an optimum noise intensity is expected to be identified and a final prediction selected in the following step of the methodology.

Leak Location Prediction Selection

According to Harmer et al. (2002), the addition of some noise to a nonlinear system can enhance its response. However, there is a point at which the addition of too much noise prevents further improvement. As previously mentioned, this phenomenon is known as stochastic resonance. One of the objectives of this paper is to demonstrate the application of this concept to the use of CNNs to detect leaks in pipelines under more realistic conditions, such as background pressure fluctuations. The selection of a leak location (and size) prediction is included as a separate step in Fig. 1 (Step 2.5) because it involves the analysis of the distribution of predictions obtained in Step 2.4.

The first part of this analysis is related to the scatter of the leak location predictions for the CNNs (m) of a particular noise intensity (n). If stochastic resonance is relevant, the predictions of the leak location in CNNs trained with larger noise intensities are expected to be more consistent and closer to the real location of the leak until the optimum noise intensity is reached. To test this, box whisker plots can be created using the predictions from the available transient tests to evaluate the effect of adding noise to the training samples of the CNNs.

In contrast, CNNs trained with noise intensities that are too large are expected to not perform well on the training because the reflections from small leaks are combined with the added Gaussian noise, and the overall performance of the CNNs decreases. To measure this, for each group of CNNs corresponding to each noise intensity (n) including the CNNs trained without any noise, the root mean squared error (RMSE) is computed for both the CNN training and testing datasets. If the resulting RMSE for a particular noise intensity exceeds a predefined threshold or the training and testing RMSE are considerably different, this noise intensity is considered too large.

By analyzing the scatter of the leak location predictions and the CNN RMSE for training and testing, one group of CNNs (m) trained with the optimum noise intensity (n_{opt}) can be selected. Using these CNNs, a final prediction for the location and size of the leak in the pipeline using the $m \times q$ available predictions can be obtained by computing the median leak location for each transient test available and then analyzing the distribution of those predictions. If for each transient test, the median location predictions are clustered around one possible location, only that prediction will be assessed in the verification process. However, if two or more clusters are identified, the multiple leak location predictions are used in the last step of the methodology.

Leak Detection Verification

The last step of the application stage considers the use of a numerical transient model of the analyzed pipeline to verify whether the predicted leak characteristics match the measured trace to a reasonable degree of accuracy (Step 2.6 in Fig. 1). This does not represent a verification of the complete methodology but a confirmation step including potential refinement of the CNN predictions for a particular pipeline. Using the same MOC numerical model used in Step 1.2 of the model development, new transient pressure head traces can be obtained using the predicted leak characteristics (size and location) obtained in Step 2.5. These numerically generated traces are then compared with the pressure head measured to assess its similarity using the normalized root mean squared error (NRMSE). Differences are expected between these transient pressure head traces due to different elements present in the pipeline that are not included in the numerical model. In addition, the fact that the final prediction is obtained from a distribution of predictions can also cause differences between these pressure head traces. However, a threshold can be defined to decide whether the CNNs' predictions are accurate enough and a final prediction has been reached.

Preliminary analysis indicated that cases in which the CNNs' predictions are not within the defined threshold are due to a discrepancy in the predicted leak size in a manner similar to what was reported by Bohorquez et al. (2021) for the detection of bursts. Thus, a potential leak size correction has been considered in this methodology through the generation of additional numerical transient pressure head traces covering the range of possible leak sizes to find the one that produced a trace with the lowest NRMSE.

Results

The proposed methodology for leak detection in pipelines as previously described has been applied to a series of tests conducted in the Robin Hydraulics Laboratory of The University of Adelaide. The objective was to demonstrate the feasibility of using CNNs to detect the location and size of leaks in pipelines under more realistic conditions. This section outlines the characteristics of the analyzed pipeline and the transient tests. A description of the application of the methodology for both stages following the steps presented in Fig. 1 is then provided.

Laboratory Tests

The pipeline in the laboratory has the configuration indicated in Fig. 2. The pipeline is connected at both ends to pressurized tanks. An inline valve was closed at the downstream end of the pipeline to allow flow only through a solenoid valve installed right before the end of the pipeline. The characteristics of the pipeline are provided in Table 1. A circular orifice of size 2.2 mm was installed 28.52 m downstream of the source tank to simulate a leak.

The transient event to detect the leak is generated by the fast closure of the solenoid valve with a closure time of 5 ms. The pressure was measured with a PDCR 810 pressure transducer (Druck, Leicester, UK) with a 10-kHz sampling rate. A total of 14 transient tests were conducted with the same configuration under similar initial conditions. The pressure head traces measured for the 14 tests at the downstream end of the pipeline are presented in Fig. 7, and each line represents a different test. The initial pressure head at the end of the pipeline was set to between 20.0 and 23.9 m. The pressure head was measured from 0.2 s before the valve closure and for a total of 3 s (although Fig. 7 indicates the pressure changes only until 1 s).

Using the results from these 14 laboratory tests and the known characteristics of the pipeline, the CNN leak location methodology presented in Fig. 1 was applied to this system. Multiple tests were analyzed to study the robustness of the leak location predictions to small differences in testing conditions, such as the initial pressure and background pressure fluctuations.

Model Development

First, a leak detection model was developed for the pipeline described in Table 1 following the steps described in Stage 1 of Fig. 1. The 1D convolutional neural networks created followed the architecture previously described with four convolutional layers, 20 filters, and three dense layers (Step 1.1). A total of 50,000 numerical transient pressure head traces were generated with a MOC numerical model. Ten different leaks were modeled at random locations within each 7.45-mm interval along the pipeline. Each of these 10 transient pressure head traces had a different randomly selected diameter varying between 0.4 and 3.5 mm. The total simulation time was set as 0.09 s, which corresponds to $3.15L/a$ seconds, L/a seconds before the closure of the valve, and $2.15L/a$ seconds after to account for the effects of the valve closure curve in the computed pressure head. To obtain different transient head pressure traces, the time resolution of the MOC numerical model needed to be at least 0.006 ms. Therefore, the total size of the CNNs' input dataset before the downsampling process is 788 million transient pressure head values (Step 1.2), and each trace has almost 16,000 head values.

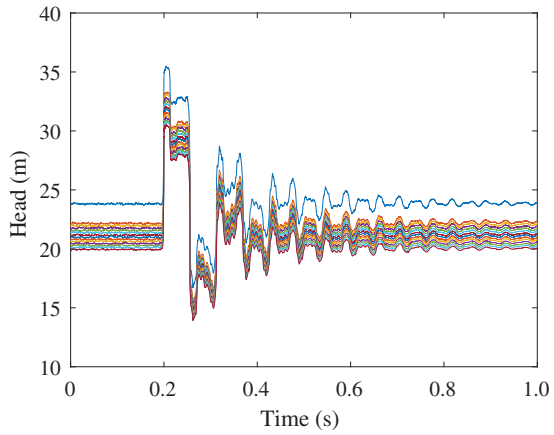

Fig. 7. (Color) Laboratory transient pressure head traces. Each series represents a different test with a different initial pressure.

According to Step 1.3 in Fig. 1, the obtained input dataset was then downsampled to a selected downsampling frequency of 5 kHz. This frequency was selected considering the dimensions of the pipeline and the potential number of weights to train in the resulting CNNs. A smaller downsampling frequency creates a very small CNN that cannot learn enough information from the transient pressure head traces. Smaller downsampling frequencies can be selected for larger pipelines with larger L/a characteristics. The resulting number of weights for the leak detection CNNs following Eq. (1) is 13,868.

After the downsampling process, the input dataset contains 8.55 million transient pressure head values for the 50,000 traces. This dataset was used in Step 1.4 to create additional CNN input datasets with the addition of noise in the transient pressure head traces. Following the definition of noise intensity previously presented, the smallest leak drop [Δh in Eq. (2)] corresponding to the smallest leak considered was 0.1238 m. Six different noise intensities were considered in this step, and the selected values of k_i and the derived standard deviations (σ_i) are presented in Table 2. These noise intensities were selected considering that the objective was to obtain CNNs with the ability to find leaks across the complete defined leak size range without significantly decreasing performance with the addition of noise. Important to note is that the values for the noise intensities to be considered are independent of the background noise present in the analyzed pipeline and are defined only based on the definition provided in Eq. (2).

The information presented in Table 2 was used to generate six additional input datasets. Each dataset contains a total of 250,000 transient pressure head traces given that five traces were created for each for the original numerical traces. Five CNNs were created for each defined noise intensity and five CNNs using the original training dataset, with no noise included. Each group of five CNNs have the same architecture but different resulting weights considering that Stochastic Gradient Descent algorithms were used in its training. As was previously explained, using these training algorithms is similar to applying Genetic Algorithms using different random number seeds. The resulting set of 35 CNNs were trained and tested simultaneously using graphics processing units (GPUs) on the University of Adelaide's High Performance Computer (HPC), Phoenix. The training process was conducted for a maximum of 24 h or less if the desired threshold of accuracy was achieved.

Figs. 8(a–g) present the percentage exceedance associated with the absolute average error in the location of leaks. This plot summarizes the results of the training and testing of seven of the 35 CNNs, where each plot corresponds to a different noise intensity CNN. Only one plot per noise intensity is included because the distribution of the errors obtained during training and testing was consistent across the five CNNs.

Table 2. Selected standard deviation for gaussian noise distribution

Standard deviation multiplier, k_i (%)	Resulting standard deviation, σ_i (mm)
5	6.19
10	12.4
25	31.0
50	61.9
100	123.8
150	185.6

Two series are included in each of the plots of Fig. 8. The blue solid line corresponds to the distribution of the absolute average leak location error for the samples used for the CNN *training*. The pink dotted line presents the leak location error for the samples used during the CNN *testing*. The percentage exceedance can be interpreted as the proportion of the total trained or tested samples for which the average leak location surpassed a certain error size. An average error in the predictions is presented because, in some cases, two or more traces with the same leak location and size were used for either the training or the testing.

Important to observe is that the maximum percentage indicated in the figure is 10% (y-axis). This means that 90% of the time that these CNNs are used with numerical transient pressure head traces, the absolute average leak location error obtained is smaller than the minimum absolute average error visible in these plots. In addition, the x-axes in Figs. 8(a–e) are presented at the same scale to facilitate its analysis.

As the standard deviation for the Gaussian distributed noise increases, the absolute average leak location error also increases due to the noise added to the training and testing samples. This figure also makes it evident that the CNNs trained and tested with transient pressure head traces without any noise performed better than the rest. However, for all of the considered noise intensities, 90% of the time, the absolute average leak location error is 0.12 m or smaller, which points to a successful result from the training of these CNNs.

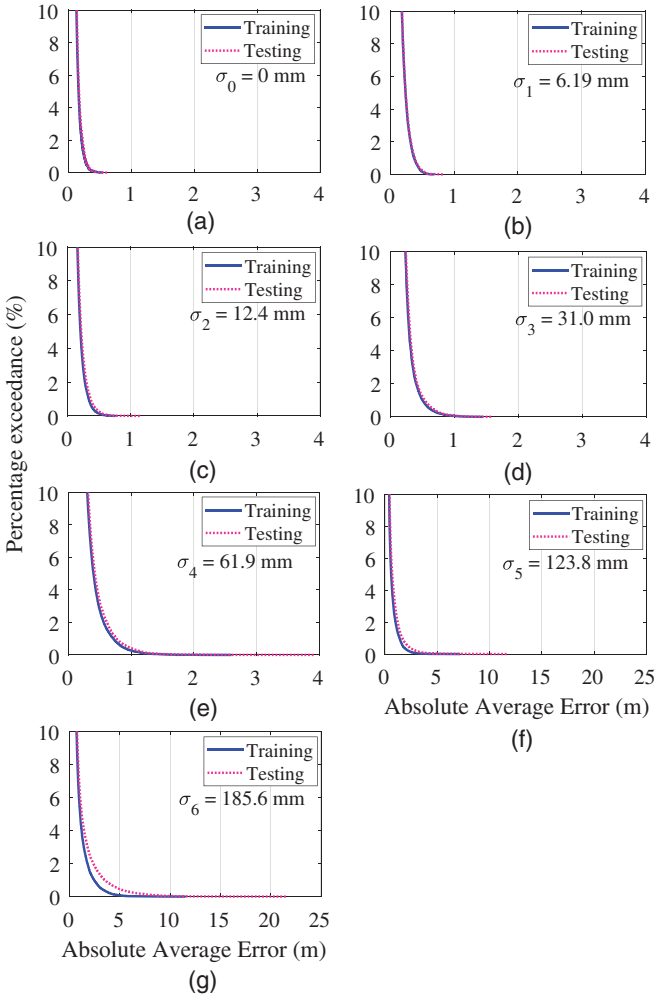

Fig. 8. (Color) Percentage exceedance for absolute average leak location error when CNNs are trained and tested with samples with different noise intensities: (a) noise intensity $\sigma_0 = 0$ mm; (b) noise intensity $\sigma_1 = 6.19$ mm; (c) noise intensity $\sigma_2 = 12.4$ mm; (d) noise intensity $\sigma_3 = 31.0$ mm; (e) noise intensity $\sigma_4 = 61.9$ mm; (f) noise intensity $\sigma_5 = 123.8$ mm; and (g) noise intensity $\sigma_6 = 185.6$ mm.

Model Application

With the successful training and testing of the leak detection CNNs, the model development stage is completed. The leak detection model application stage is now applied to identify the leak present in a real pipeline in a laboratory setting (Stage 2 in Fig. 1). The preprocessing of the obtained transient pressure head traces (Steps 2.1–2.3 in Fig. 1) started with a reduction of the background pressure fluctuations due to the flow through the solenoid valve installed at the end of the pipeline. Following the process described in Step 2.1, two 0.2-s segments were analyzed in each of the 14 measured pressure head traces before the solenoid valve closure and at the end of the 3-s recorded signal. Two normal distributions were obtained from the pressure fluctuations before and after the solenoid valve closure for each transient test. Average standard deviations before and after the transient test of 0.0392 m and 0.0097 m, respectively, were obtained. An example of the resulting transient pressure head traces after the background pressure fluctuation reduction step is presented in Fig. 9.

This figure indicates that the background pressure fluctuation reduction process does not dramatically change the transient pressure head traces because no differences are evident when a 20-m scale is used for the y-axis. However, when a different scale is analyzed (in the red subplot), clear differences in the pressure fluctuations are noticeable after the transformation of the pressure before the transient events. This step allows for a reduction in the background transient pressure head fluctuations and an improved application of the leak detection CNNs.

The resulting transient pressure head traces were further transformed to complete the preprocessing described in Fig. 1. First, the 14 measured transient pressure head traces were shifted to be aligned to one initial average steady-state pressure head. As indicated in Fig. 7, the initial pressure head of each test was slightly different within a 3.9 m range. All traces were aligned to an average steady-state pressure head of 21.16 m. This value corresponds to the initial pressure head considered for the generation of the numerical transient pressure head traces in Step 1.2. The resulting shifted traces were also trimmed to select only the segments of the transient pressure head of interest corresponding to L/a seconds before the closure of the solenoid valve and $2.15L/a$ seconds after this closure. The resulting transient pressure head traces are presented in Fig. 10.

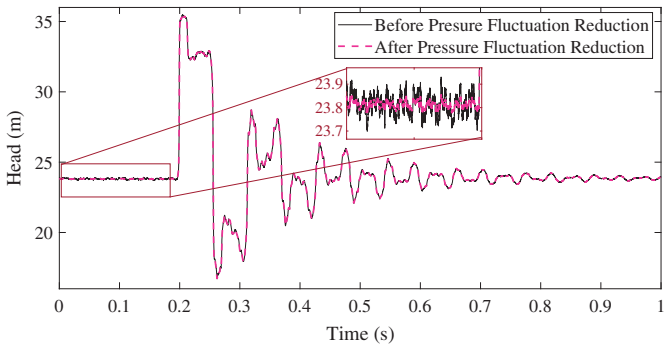

Fig. 9. (Color) Results of background pressure fluctuation reduction.

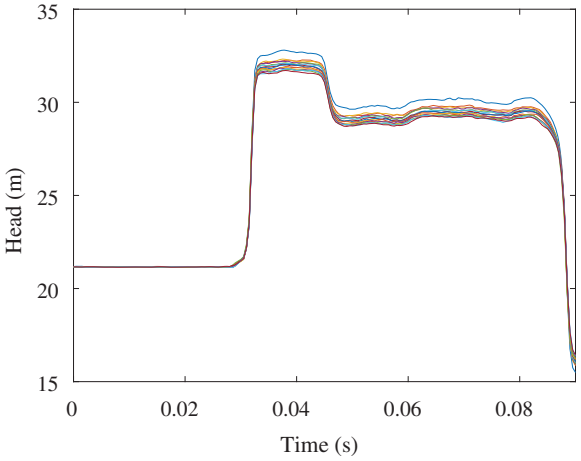

Fig. 10. (Color) Transient pressure head traces (measured at the downstream end of the pipeline) to process through leak detection CNN (each series corresponds to a laboratory test).

Important to observe is that because each transient test had a different steady-state pressure head, the initial pressure head increase after the solenoid valve closure is also different in every test. This difference is due to the small differences in the resulting flow in the pipeline given different initial pressures, in a similar manner as reported by Meniconi et al. (2019). However, the transient pressure head traces were not further transformed to test the leak detection CNNs' performance to predict accurate leak locations under these conditions. The last step of the preprocessing stage included the downsampling of the measured transient pressure head traces to a 5-kHz frequency to match the traces to the dimensions of the input for the leak detection CNNs.

The second part of the leak detection model application involved an analysis of the transient pressure head traces (Steps 2.4–2.6 in Fig. 1). All of the preprocessed transient pressure head traces were analyzed using a total of 35 trained CNNs in Step 1.5 (five CNNs per noise intensity level including the CNNs trained with samples without any noise). The distribution of leak locations predictions is indicated in Fig. 11. A seven-color scale was used in this figure to illustrate the distribution of leak location predictions on each noise intensity defined in Step 1.4 and the CNNs trained with samples without any noise. In addition, Fig. 11 presents an indication of the end of the pipeline (37.24 m) and in light blue the location of the leak in the pipeline (at 28.05 m).

This figure indicates the very large range of the leak location predictions when the CNNs were trained without any noise in the transient pressure head traces. Except for two outliers in the predictions for traces #13 and #14, none of the leak location predictions are within the physical limits of the pipeline. Therefore, these predictions are not visible in the figure. This result demonstrates the challenges of applying ANNs for the detection of anomalies in pipelines under more realistic conditions (Bohorquez et al. 2020). Because these CNNs were trained with theoretical numerical samples with perfect data, the predictions when the analyzed transient pressure head traces have background pressure fluctuations become illogical for the leak location.

Fig. 11 also presents the significant influence that the addition of noise in the numerical transient pressure head traces for the CNNs' training has in the resulting distribution of leak location predictions. The addition of a Gaussian distributed noise with a standard deviation of 6.2 mm (σ_1) has a significant effect on the resulting leak location predictions. Most of these predictions can now be found within the physical limits of the pipeline (orange series in Fig. 11). Important to note is that a background pressure fluctuation with a standard deviation of 6.2 mm is smaller in magnitude than the real background pressure fluctuations observed in this pipeline. However, its introduction in the CNN training dataset has proven to be highly effective in improving the obtained leak location predictions. This finding aligns with previous authors' findings that state that the addition of noise can benefit ANN performance (Fukami et al. 2020).

Despite the clear advantages of applying Gaussian distributed noise, the results presented in Fig. 11 also demonstrate that the addition of noise with a very small standard deviation is not enough for a satisfactory prediction of the leak location. This highlights the importance of deploying stochastic resonance to determine the optimum noise intensity that should be introduced in the ANN training samples (Harmer et al. 2002). Fig. 11 demonstrates that, as the noise intensity (σ_i) increases, the distribution of the leak locations is more compact, and the predictions are generally closer to the real leak location. Predictions from CNNs trained with noise intensities σ_2 and σ_3 (Table 2) are within the length of the pipeline but vary

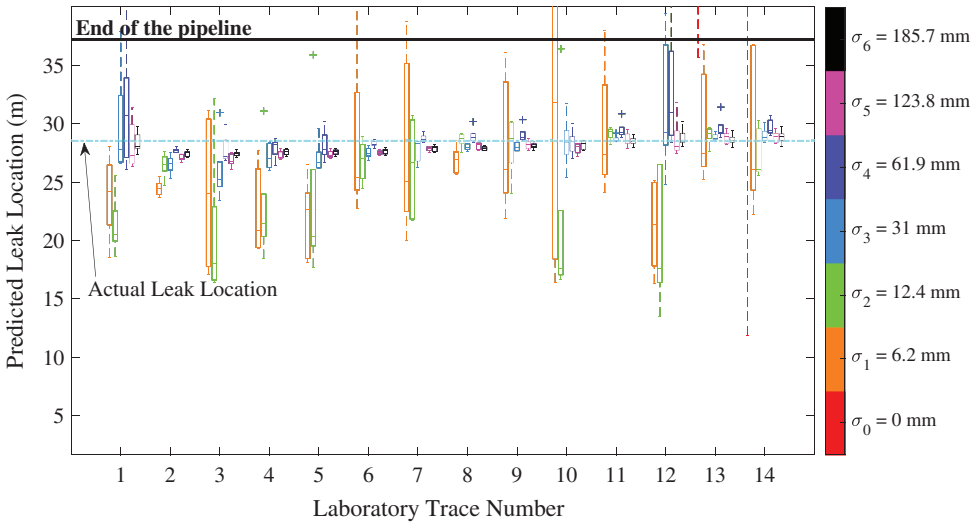

Fig. 11. (Color) Predicted leak location for sets of CNNs trained with samples with different noise intensities.

J. Water Resour. Plann. Manage., 2022, 148(3): 04022001

considerably between the different transient tests conducted. Leak location prediction errors obtained from the last three noise intensities (σ_{4-6}) range between 2 and 3.8 m, with a couple of predictions outside the physical length of the pipelines for σ_4.

Although most of the transient tests allow for a similar distribution of predictions for a particular noise intensity, transient tests #1 and #12 resulted in more scattered leak location predictions. Leak location predictions for transient test #1 are less satisfactory because this test had a more prominent difference in the steady-state pressure head. Thus, the difference in the resulting initial pressure head increase after the closure of the solenoid valve is more prominent, as is observed in Fig. 10. In contrast, even though transient test #12 does not present with any particular differences relative to the other transient tests, it produced less consistent results for all of the noise intensities. These results point out that additional background noise might have existed during this test. Considering this, conducting multiple tests in a range of similar initial conditions provides more information that the CNNs can process instantaneously and allows for a more confident prediction of the leak location. If only one test in the pipeline is used, the risk of considering an incorrect prediction as true information about an existing leak would exist.

A perfect distribution of leak location predictions implies that all CNNs trained for a particular noise intensity predict the correct leak location. However, given that each CNN has a different set of resulting weights after the training process, this result is very difficult to accomplish. Therefore, the effectiveness of CNNs should be measured by their ability to produce consistent predictions with a reasonable degree of accuracy for field applications of this technique.

To further analyze the results obtained from the leak detection CNNs, Fig. 12(a) presents the distribution of the absolute median error in the predicted leak location for each group of CNNs trained with different noise intensities. The median leak location prediction of each transient test in Fig. 11 was extracted, and the error between this prediction and the real leak location was computed. The distribution presented in blue in this box plot is obtained from the 14 median leak location errors. This distribution is presented as an absolute value to demonstrate the applicability of stochastic resonance, as was previously reported (Ikemoto et al. 2018).

The absolute median error in the leak location for the CNNs trained without any noise [i.e., noise standard deviation of zero in Fig. 12(a)] is not visible in the scale of the plot because almost all of the predictions are outside the length of the pipeline. Similarly, this plot demonstrates that the addition of a very small noise distribution in the training samples ($\sigma_1 = 6.2$ mm) drastically improves the performance of the CNNs. The resulting distribution of absolute mean location errors oscillates between 1 and 8 m. However, an 8 m error is still not acceptable for the location of a leak in a 37.24-m long pipeline (which represents a 21.48% error). As the noise standard deviation increases, the distribution of the absolute median error clearly narrows, in concordance with the concept of stochastic resonance. Absolute mean location errors vary between 0.02 and 1.09 m (0.05%–2.93% error) for the largest noise standard deviation considered.

An analysis of only the distribution of the absolute median location errors in Fig. 12(a) indicates that selecting the predictions of the CNNs trained with the largest noise intensity seems logical. However, the optimum noise intensity should be selected also by consideration of the performance of the CNNs during training and testing. Fig. 12(a) presents on the right-hand y-axis the distribution of the RMSE for the training (in light green) and the testing (in black) of the CNNs for each noise intensity (indicated in Table 2). The RMSE was computed using the leak location error

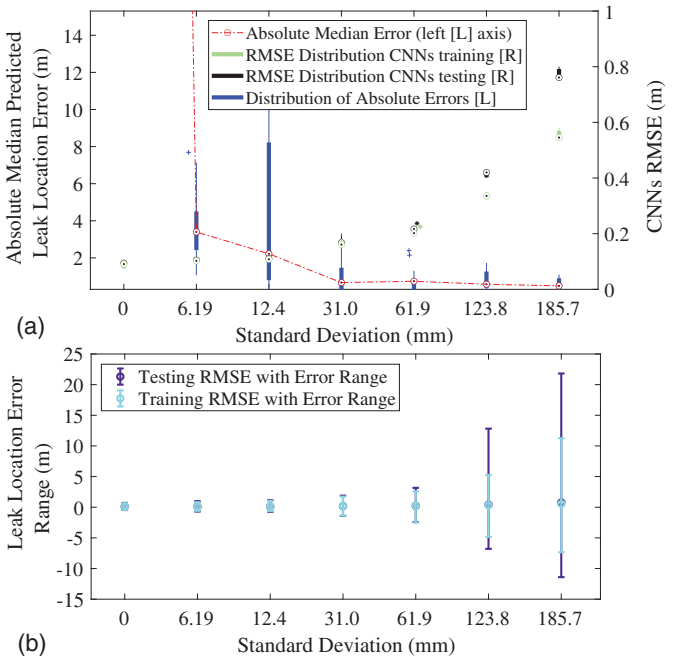

Fig. 12. (Color) Leak location prediction selection: (a) left axis: absolute median predicted leak location error for different noise intensities (standard deviation) obtained from laboratory tests. Error value for standard deviation = 0 is well above maximum extent of vertical axis (234.27 m). Right axis: training and testing CNN's root mean square error obtained using numerical samples; and (b) leak location error range during training and testing using numerical samples.

of each of the 125,000 samples used for the training or testing of the CNNs (or 25,000 for the case of the CNNs trained without any noise). A satisfactory CNN training process will have low RMSE values and similar RMSE magnitudes in both the training and the testing.

Fig. 12(b) presents error plots of the RMSE (in circles) and the complete range of errors for the leak location prediction (whiskers). These figures indicate that CNNs trained with samples with large noise intensities result in larger values of RMSE and significantly larger ranges of possible leak location errors. Both of these metrics are considerably larger for the last two sets of CNNs (corresponding to $\sigma_5 = 123.8$ mm and $\sigma_6 = 185.7$ mm) with significantly different results for the training and testing of these CNNs. These results point to a certain level of overfitting in these CNNs that is also visible in Figs. 8(f and g).

Although these results were presented as part of the model development stage (Step 1.5 in Fig. 1), they are relevant in the model application stage for the leak location prediction selection step. The final leak location prediction should be a robust prediction (in terms of consistency among the conducted tests) and be the product of a reliable set of CNNs. For this reason, it can be concluded that the optimum noise intensity for this application of the proposed leak detection model is obtained when the noise has a Gaussian noise distribution with a standard deviation of 61.9 mm (σ_4). The median leak location prediction for this group of CNNs was 28.74 m, and the median predicted leak size was 2.32 mm. These predictions represent a 0.58% error in the location of the leak and a 5.52% error in the size of the leak. Important to state is that the obtained value for the optimal noise intensity characterized by the standard deviation of a Gaussian noise distribution is unique for this pipeline and this set of experimental tests.

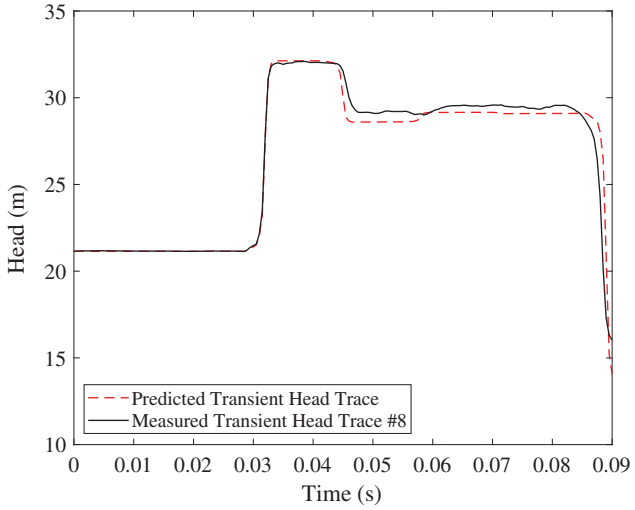

Fig. 13. (Color) Comparison between transient pressure trace numerically modeled from selected CNN prediction and laboratory transient pressure trace.

The last step of the model application consists of verifying the accuracy of the obtained prediction. This step comprises the generation of a numerical transient pressure head trace with the characteristics of the final prediction obtained in the previous step and its comparison with the measured transient pressure head traces. This comparison is presented in Fig. 13 using test #18 as an example.

A reasonable match between these two traces is observed in this figure, demonstrating the successful prediction of the location and size of the leak using a set of CNNs. The NRMSE was computed between these two transient pressure head traces to obtain a value of 2.06%, demonstrating again the accuracy of the methodology proposed.

Conclusions

This paper proposed a new, comprehensive technique for the location and characterization of leaks in pipelines using fluid transients and CNNs. This methodology has proven successful when applied to pipelines under more realistic conditions in the presence of background pressure fluctuations. A full methodology divided into two stages (model development and model application) was presented. The model development stage includes the design and training of CNNs capable of identifying leaks in transient pressure head traces with different noise intensities. The model application stage describes the preprocessing and analysis steps for any pressure transient measurements obtained from a transient event caused by the closure of a valve in a pipeline.

This technique was applied to a laboratory pipeline for which 14 transient events were generated with the closure of a side discharge solenoid valve. One circular orifice was installed in the pipeline to simulate a leak. A leak detection model with 35 different CNNs was developed for this pipeline for which six different noise intensities were considered with standard deviations between 6.2 and 185.7 mm. The application of the leak detection model to the available transient tests demonstrated the significant importance of the addition of noise in the performance of the CNNs for predicting the location of a leak in the pipeline. For CNNs trained using numerical transient pressure head traces with no noise added, the distribution of the leak locations was beyond the actual extremities of the pipeline in most cases. However, as the noise intensity increases, the distribution of the leak location predictions narrows around the real leak location (as indicated in Fig. 11).

The results obtained in this paper demonstrate that the deployment of stochastic resonance assists in detecting leaks in water pipelines. The addition of noise in the training samples of a CNN significantly improved its performance—to the point that consistent and accurate predictions can be obtained. To select the optimum noise intensity for the presented laboratory application, a combined analysis of the distribution of the predicted leak locations and the RMSE of the training and testing of the CNNs was conducted. The results from that analysis indicated that the optimum noise intensity for the analyzed pipeline and the set of experimental tests was found when the standard deviation was 61.9 mm. The final leak prediction was determined for the laboratory pipeline only 0.74 m from the real leak location. This prediction corresponds to an error of 0.59% with a very accurate prediction of the leak size.

The results obtained in this paper demonstrate that the use of CNNs trained with numerical samples with the addition of noise is a promising technique for leak detection in pipelines under more realistic conditions. Although expected differences are evident between the available numerical model and the measured transient tests, an accurate prediction of the location and size of the leak was obtained. The deployment of stochastic resonance improves the performance of CNNs in locating anomalies in these pipelines not by replicating the existing background noise but by enhancing the response of the CNNs when the optimum noise intensity is added. The specific value of optimum noise intensity is case-dependent, but this methodology can be expected to be applied to different pipelines with an appropriate range of noise intensities.

This paper represents an important contribution to the development of a fully automated technique for leak detection in pipelines using transient waves and artificial intelligence algorithms. However, more research is needed to analyze the performance of this technique when tests are conducted under different base conditions when the pipeline experiences significant background pressure fluctuations such as demand consumption, and in more complex systems such as pipeline looped networks. For these applications, a combination of more robust existing transient-based techniques and CNNs might be necessary.

Data Availability Statement

Some of the data, models, or code that support the findings of this study are available from the corresponding author on reasonable request. The data available include a subset of the convolutional neural network training and testing dataset, the data obtained from the laboratory tests, and the code developed to generate the convolutional neural networks input training and testing dataset.

Acknowledgments

This work has been supported with supercomputing resources provided by the Phoenix HPC service at the University of Adelaide and by the Australia Research Council through Discovery Project Grant No. DP190102484.

References

Ahmadi, M. S., J. Sušnik, W. Veerbeek, and C. Zevenbergen. 2020. "Towards a global day zero? Assessment of current and future water

supply and demand in 12 rapidly developing megacities." *Sustainable Cities Soc.* 61 (Oct): 102295. https://doi.org/10.1016/j.scs.2020.102295.
Akhtar, N., and A. Mian. 2018. "Threat of adversarial attacks on deep learning in computer vision: A survey." Accessed March 1, 2021. https://ui.adsabs.harvard.edu/abs/2018arXiv180100553A.
Benzi, R., G. Parisi, A. Sutera, and A. Vulpiani. 1982. "Stochastic resonance in climatic change." *Tellus* 34 (1): 10–15. https://doi.org/10.3402/tellusa.v34i1.10782.
Benzi, R., A. Sutera, and A. Vulpiani. 1981. "The mechanism of stochastic resonance." *J. Phys. A: Math. Gen.* 14 (11): L453–L457. https://doi.org/10.1088/0305-4470/14/11/006.
Bishop, C. M. 1995. *Neural networks for pattern recognition*. Oxford, UK: Oxford University Press.
Bohorquez, J., B. Alexander, A. R. Simpson, and M. F. Lambert. 2020. "Leak detection and topology identification in pipelines using both fluid transients and artificial neural networks." *J. Water Resour. Plann. Manage.* 146 (6): 04020040. https://doi.org/10.1061/(ASCE)WR.1943-5452.0001187.
Bohorquez, J., A. R. Simpson, and M. F. Lambert. 2018. "Characterization of transient pressure traces due to the effects of different anomalies and features in water pipelines." In *Proc., 13th Pressure Surges Conf.*, 151–169. Cranfield, UK: BHR Group.
Bohorquez, J., R. Simpson Angus, F. Lambert Martin, and B. Alexander. 2021. "Merging fluid transient waves and artificial neural networks for burst detection and identification in pipelines." *J. Water Resour. Plann. Manage.* 147 (1): 04020097. https://doi.org/10.1061/(ASCE)WR.1943-5452.0001296.
Brunone, B., C. Capponi, and S. Meniconi. 2021. "Design criteria and performance analysis of a smart portable device for leak detection in water transmission mains." *Measurement* 183 (Oct): 109844. https://doi.org/10.1016/j.measurement.2021.109844.
Che, T.-C., H.-F. Duan, and P. J. Lee. 2021. "Transient wave-based methods for anomaly detection in fluid pipes: A review." *Mech. Syst. Sig. Process.* 160 (Nov): 107874. https://doi.org/10.1016/j.ymssp.2021.107874.
Cheng, W., X. Xu, Y. Ding, and K. Sun. 2020. "Stochastic resonance in a single-well potential and its application in rolling bearing fault diagnosis." *Rev. Sci. Instrum.* 91 (6): 064701. https://doi.org/10.1063/1.5143050.
Collins, J. J., C. C. Chow, and T. T. Imhoff. 1995. "Aperiodic stochastic resonance in excitable systems." *Phys. Rev. E* 52 (4): R3321–R3324. https://doi.org/10.1103/PhysRevE.52.R3321.
Duan, H. F. 2017. "Transient frequency response based leak detection in water supply pipeline systems with branched and looped junctions." *J. Hydroinf.* 19 (1): 17–30. https://doi.org/10.2166/hydro.2016.008.
Duan, H.-F., B. Pan, M. Wang, L. Chen, F. Zheng, and Y. Zhang. 2020. "State-of-the-art review on the transient flow modeling and utilization for urban water supply system (UWSS) management." *J. Water Supply Res. Technol. AQUA* 69 (8): 858–893. https://doi.org/10.2166/aqua.2020.048.
Falanga, M., E. De Lauro, and S. de Martino. 2020. "Stochastic resonance observed in aerosol optical depth time series." *Atmosphere* 11 (5): 502. https://doi.org/10.3390/atmos11050502.
Feng, Z., S. Li, and Z. Xu. 2019. "A novel adaptive stochastic resonance scheme for underwater optical wireless communication." In *Proc., 2019 11th Int. Conf. on Wireless Communications and Signal Processing (WCSP)*. New York: IEEE. https://doi.org/10.1109/WCSP.2019.8927867.
Fukami, K., K. Fukagata, and K. Taira. 2020. "Assessment of supervised machine learning methods for fluid flows." Accessed March 1, 2021. https://ui.adsabs.harvard.edu/abs/2020arXiv200109618F.
Geem, Z. W., C.-L. Tseng, J. Kim, and C. Bae. 2007. "Trenchless water pipe condition assessment using artificial neural network." In *Proc., Pipelines 2007*. Reston, VA: ASCE.
Gong, J., S. T. N. Nguyen, M. L. Stephens, M. F. Lambert, A. Marchi, A. R. Simpson, and A. C. Zecchin. 2018. "Correlation of post-burst hydraulic transient noise for pipe burst/leak localisation in water distributions systems." In *Proc., 13th Int. Conf. on Pressure Surges*, 201–215. Cranfield, UK: BHR Group.
Goodfellow, I., Y. Bengio, and A. Courville. 2016. *Deep learning*. Cambridge, MA: MIT Press.
Harmer, G. P., B. R. Davis, and D. Abbott. 2002. "A review of stochastic resonance: Circuits and measurement." *IEEE Trans. Instrum. Meas.* 51 (2): 299–309. https://doi.org/10.1109/19.997828.
Ikemoto, S., F. DallaLibera, and K. Hosoda. 2018. "Noise-modulated neural networks as an application of stochastic resonance." *Neurocomputing* 277 (Feb): 29–37. https://doi.org/10.1016/j.neucom.2016.12.111.
Jönsson, L., and M. Larson. 1992. "Leak detection through hydraulic transient analysis." In *Pipeline systems*, edited by B. Coulbeck and E. P. Evans, 273–286. Dordrecht, Netherlands: Springer.
Liggett, J. A., and L. Chen. 1994. "Inverse transient analysis in pipe networks." *J. Hydraul. Eng.* 120 (8): 934–955. https://doi.org/10.1061/(ASCE)0733-9429(1994)120:8(934).
Luchinsky, D. G., R. Mannella, P. V. E. McClintock, and N. G. Stocks. 1999. "Stochastic resonance in electrical circuits. I. Conventional stochastic resonance." *IEEE Trans. Circuits Syst. II Analog Digital Signal Process.* 46 (9): 1205–1214. https://doi.org/10.1109/82.793710.
Martinez-Diaz, M., J. Fierrez, M. Freire, and J. Ortega-Garcia. 2007. "On the effects of sampling rate and interpolation in HMM-based dynamic signature verification." In *Proc., 9th Int. Conf. on Document Analysis and Recognition (ICDAR 2007)*, 1113–1117. New York: IEEE.
Meniconi, S., B. Brunone, M. Cifrodelli, C. Capponi, A. Rubin, L. Tirello, and P. Lucato. 2019. "Initial functioning conditions vs. transient test-based technique performance for leak detection." In *Proc., 38th IAHR World Congress*, 3345–3350. London: International Association for Hydro-Environment Engineering and Research.
Mujtaba, S. M., T. A. Lemma, S. A. A. Taqvi, T. N. Ofei, and S. K. Vandrangi. 2020. "Leak detection in gas mixture pipelines under transient conditions using Hammerstein model and adaptive thresholds." *Processes* 8 (4): 474. https://doi.org/10.3390/pr8040474.
Mukherjee, S., W. Bebermeier, and B. Schütt. 2018. "An overview of the impacts of land use land cover changes (1980–2014) on urban water security of Kolkata." *Land* 7 (3): 91. https://doi.org/10.3390/land7030091.
Murray, A. F., and P. J. Edwards. 1994. "Enhanced MLP performance and fault tolerance resulting from synaptic weight noise during training." *Trans. Neural Networks* 5 (5): 792–802. https://doi.org/10.1109/72.317730.
Mutikanga, H. E., S. K. Sharma, and K. Vairavamoorthy. 2013. "Methods and tools for managing losses in water distribution systems." *J. Water Resour. Plann. Manage.* 139 (2): 166–174. https://doi.org/10.1061/(ASCE)WR.1943-5452.0000245.
Nakama, T. 2009. "Theoretical analysis of batch and on-line training for gradient descent learning in neural networks." *Neurocomputing* 73 (1): 151–159. https://doi.org/10.1016/j.neucom.2009.05.017.
Neelakantan, A., L. Vilnis, Q. V. Le, I. Sutskever, L. Kaiser, K. Kurach, and J. Martens. 2015. "Adding gradient noise improves learning for very deep networks." Preprint, submitted November 21, 2015. https://arxiv.org/abs/1511.06807.
Nguyen, S., J. Gong, M. F. Lambert, A. C. Zecchin, and A. R. Simpson. 2018. "Least squares deconvolution for leak detection with a pseudo random binary sequence excitation." *Mech. Syst. Sig. Process.* 99 (Jan): 846–858. https://doi.org/10.1016/j.ymssp.2017.07.003.
Puust, R., Z. Kapelan, D. A. Savic, and T. Koppel. 2010. "A review of methods for leakage management in pipe networks." *Urban Water J.* 7 (1): 25–45. https://doi.org/10.1080/15730621003610878.
Rifai, S., X. Glorot, Y. Bengio, and P. Vincent. 2011. "Adding noise to the input of a model trained with a regularized objective." Accessed March 1, 2021. https://ui.adsabs.harvard.edu/abs/2011arXiv1104.3250R.
Romano, M., Z. Kapelan, and D. A. Savić. 2014. "Automated detection of pipe bursts and other events in water distribution systems." *J. Water Resour. Plann. Manage.* 140 (4): 457–467. https://doi.org/10.1061/(ASCE)WR.1943-5452.0000339.
Roy, U. 2017. "Leak detection in pipe networks using hybrid ANN method." *Water Conserv. Sci. Eng.* 2 (4): 145–152. https://doi.org/10.1007/s41101-017-0035-1.
Shu, H., and H. Zhu. 2019. "Sensitivity analysis of deep neural networks." Accessed March 1, 2021. https://ui.adsabs.harvard.edu/abs/2019arXiv190107152S.
Szegedy, C., W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus. 2013. "Intriguing properties of neural networks." Accessed March 1, 2021. https://ui.adsabs.harvard.edu/abs/2013arXiv1312.6199S.
Wang, X., and M. S. Ghidaoui. 2018. "Pipeline leak detection using the matched-field processing method." *J. Hydraul. Eng.* 144 (6): 04018030. https://doi.org/10.1061/(ASCE)HY.1943-7900.0001476.

Wang, X., J. Lin, A. Keramat, M. S. Ghidaoui, S. Meniconi, and B. Brunone. 2019. "Matched-field processing for leak localization in a viscoelastic pipe: An experimental study." *Mech. Syst. Sig. Process.* 124 (Jun): 459–478. https://doi.org/10.1016/j.ymssp.2019.02.004.

Wang, Y.-H., and J. C. Santamarina. 2002. "Dynamic coupling effects in frictional geomaterials—Stochastic resonance." *J. Geotech. Geoenviron. Eng.* 128 (11): 952–962. https://doi.org/10.1061/(ASCE)1090-0241(2002)128:11(952).

Zangenehmadar, Z., and O. Moselhi. 2016. "Assessment of remaining useful life of pipelines using different artificial neural networks models." *J. Perform. Constr. Facil.* 30 (5): 04016032. https://doi.org/10.1061/(ASCE)CF.1943-5509.0000886.

Zeng, W., J. Gong, A. R. Simpson, B. S. Cazzolato, A. C. Zecchin, and M. F. Lambert. 2020. "Paired-IRF method for detecting leaks in pipe networks." *J. Water Resour. Plann. Manage.* 146 (5): 04020021. https://doi.org/10.1061/(ASCE)WR.1943-5452.0001193.