

TOWARDS RESILIENT WATER NETWORKS AND PROACTIVE WATER MANAGEMENT: A HOLISTIC MACHINE LEARNING FRAMEWORK FOR FORECASTING, DETECTION, AND LOCALIZATION OF LEAKAGES

JAWAD FAYAZ^{1,2}, LIZ VARGA^{3,4} AND YIJUN XIANG⁵

¹ Department of Computer Science
University of Exeter, Exeter, UK
j.fayaz@exeter.ac.uk

² UKCRIC Limited
London, UK
j.fayaz@ukcric.ac.uk

³ Department of Civil, Environmental, and Geomatic Engineering
University College London, London, UK
l.varga@ucl.ac.uk

⁴ UKCRIC Limited
London, UK
limited@ukcric.ac.uk

⁵ Department of Computer Science
Georgia Institute of Technology, Atlanta, USA
yxiang63@gatech.edu

Key words: Variational autoencoders, Transformer neural networks, Graph attention networks, Leakage prediction, Detection, Localization, Class-specific attention

Abstract. Efficient and sustainable water distribution is critical to addressing the global challenges of ageing infrastructure, increasing demand, and climate variability. Modern water distribution systems suffer from significant inefficiencies, with up to 25% of treated water lost to leakages in the UK and a yearly global loss of 32 billion m³ of treated water. These losses not only strain water resources but also lead to increased energy consumption and elevated greenhouse gas emissions, as additional energy is required to treat and pump water that ultimately does not reach consumers. Overall, this results in resource wastage, operational disruptions, increased energy costs, avoidable carbon emissions, and diminished public trust. To mitigate these issues, an integrated framework incorporating three pillars - forecasting, detection, and localization of leakages - is essential. For forecasting, the proposed framework includes a class-specific attention-based deep learning model which utilizes historical flow data, weather variables (including temperature, humidity, solar radiation, evaporation, and precipitation), and recent consumption patterns of the preceding five days to unifiedly predict future water demand and probability of leakages in a given location aggregated by district-metered area (DMA). These predictions enable utilities to preemptively allocate resources and reduce the burden on overworked infrastructure. For leakage detection and complementing the forecasting, a domain-informed variational autoencoder (VAE) leverages the recorded flow data to associate irregular flow

patterns with leakage events. The VAE reduces the dimensions of the water flow time series data into two-dimensional surrogate latent variable mapping which sufficiently and efficiently captures the distinct characteristics of leakage and regular (non-leakage) flow. Localization of leakages, the third critical component, relies on combining sensor inputs with spatial analysis to pinpoint affected pipeline segments. This is done using specially designed graph attention neural networks that can predict flow and leakages in a spatiotemporal manner, refining the accuracy of leak location estimates, enabling targeted interventions and reducing downtime. The models lead to an average accuracy of $\geq 85\%$ in leakage monitoring. These models can process data at fine temporal scales, offering near real-time insights into system health. This holistic framework minimizes water losses, optimizes repair strategies, and enhances the resilience of water networks against future challenges. Furthermore, incorporating weather and environmental data, such as precipitation, temperature, and soil moisture, enhances the accuracy and reliability of these systems by accounting for external factors influencing water demand and leakage risks. Explainability is another vital aspect, as interpretability tools such as Shapley additive explanations provide transparency into model decisions, fostering operator trust and informed decision-making. By integrating this framework with the three pillars—forecasting, detection, and localization—water utilities can transition from reactive to proactive leakage management and utilities can ensure the sustainability and reliability of this critical infrastructure, safeguarding water resources for future generations

1 EXTENDED ABSTRACT

Efficient and sustainable water distribution systems are critical in mitigating global challenges associated with ageing infrastructure, escalating water demands, and increasing climate variability. Modern water distribution networks exhibit substantial inefficiencies, evidenced by up to 25% leakage rates in the UK and global annual losses estimated at 32 billion cubic meters of treated water [1, 11]. These leakages exacerbate water scarcity, increase the energy consumption required for water pumping and treatment, and amplify greenhouse gas emissions, resulting in substantial resource wastage, operational disruptions, elevated operational costs, and reduced public trust [10].

In response, this study presents a scientifically robust, integrated framework systematically addressing leakage management through three interdependent stages: leakage forecasting, leakage detection, and leakage localization. The methodology facilitates proactive water distribution management and enhances operational decision-making through optimized resource allocation.

The dataset utilized comprises water flow data and repair logs from over 2000 district-metered areas (DMAs) across the UK, spanning multiple years, with measurements recorded at 15-minute intervals. Rigorous preprocessing techniques were employed, addressing critical data quality issues. Kalman smoothing [6] was applied for imputing missing data points or erroneous readings ensuring continuity and temporal consistency. Seasonal segmentation was implemented to accurately reflect seasonal demand fluctuations, enhancing the reliability of model inputs. Isolation forest algorithms [8] effectively identified and removed outliers, thus reducing data contamination and improving analytical precision [10]. Spatial dependencies were rigorously defined through detailed spatial metrics such as topological distances computed using breadth-first search (BFS), and Haversine distances derived from geographic coordinates. A robust data workflow was employed for sample groupings generation for both leakage and non-leakage cases, creating historical sequences of seven days (672 data points) utilized for training predictive models with a target of subsequent 24-hour data (96 data

points), including both water demands and leakage labels. Random forest analysis [3] rigorously evaluated feature importance for more than 10 weather parameters, statistically identifying temperature, specific humidity, and solar radiation as significant predictors for water demand and leakage risk.

The leakage forecasting stage is operationalized using FLOWAID (Flow and Leakage Forecasting using Weather-Adaptive neural network for Intelligent Decision-making). This model combines convolutional residual networks, bidirectional long short-term memory (BiLSTM) networks [12], and a dual-branch multi-task prediction head for jointly forecasting water demand and leakage probability. The convolutional residual network block extracts the spatial dependencies from intricate weather-flow interactions. This employs residual connections to maintain effective gradient flow, facilitating deeper network training and robust feature extraction. The convolutional residual block is expressed in Equation (1), where x represents the input tensor (flow or weather features), and y is the output after residual convolution. Temporal dynamics are then meticulously captured using the BiLSTM block, effectively modeling complex sequential dependencies inherent within water distribution systems. The BiLSTM modeling temporal sequences is shown in Equation (2), where h_t is the combined hidden state at time t , and x_t is the input feature vector at time t . An attention mechanism refines the sequence representation, with weights computed in Equation (3), where v , W_h , and b_h are learnable parameters of the class-specific attention layer [2]. Crucially, the incorporation of a class-specific attention mechanism significantly improves model sensitivity by amplifying features critical to leakage events, effectively countering the inherent imbalance in leakage versus non-leakage event frequencies.

$$y = \text{ReLU}(\text{BatchNorm}(x + \text{Conv1D}(\text{BatchNorm}(\text{Conv1D}(x)))))) \quad (1)$$

$$h_t = \overrightarrow{\text{LSTM}}(x_t, h_{t-1}) \oplus \overleftarrow{\text{LSTM}}(x_t, h_{t+1}) \quad (2)$$

$$\alpha_t = \frac{\exp(e_t)}{\sum_i \exp(e_i)}, \quad e_t = v^T \tanh(W_h h_t + b_h) \quad (3)$$

The model outputs, including future water demands and leakage probability, are generated concurrently by two separate MLP branches of the neural network, leveraging shared feature representations to enhance predictive efficiency and accuracy. The flow forecasting output uses adaptive Huber loss [5] (Equation 4), where y_i is the true flow, \hat{y}_i is the predicted flow, w_i are instance weights, and δ_{Huber} is the Huber loss. Leakage classification uses weighted binary cross-entropy loss (Equation 5) and where $y_i \in \{0, 1\}$ is the true label (leak or no leak), and \hat{y}_i is the predicted probability. The adaptive Huber and weighted binary cross-entropy loss functions optimize model training, addressing class imbalance and minimizing prediction errors. FLOWAID demonstrates exceptional forecasting accuracy, with a true positive rate of 96.7% and a true negative rate of 95.9% within a 12-hour forecasting horizon, underscoring its reliability and operational applicability.

$$L_{\text{flow}} = \sum_i w_i \delta_{\text{Huber}}(y_i - \hat{y}_i) \quad (4)$$

$$L_{\text{leak}} = - \sum_i w_i [y_i \log(\hat{y}_i) + (1 - y_i) \log(1 - \hat{y}_i)] \quad (5)$$

For leakage detection, the framework leverages a domain-informed variational autoencoder (VAE) [7] integrated with support vector machines (SVM) [4]. The VAE is specifically tailored to reduce

the dimensionality of the preceding 24-hour water flow time series into a compact, interpretable, two-dimensional latent variable space, which distinctly separates leakage conditions from typical operational states. This is done by maximizing the evidence lower bound (ELBO), defined in Equation (6), where z is the latent representation, x is the input sequence, θ and ϕ are decoder and encoder parameters, and D_{KL} is the Kullback-Leibler divergence. A domain-specific contrastive loss enforces class separation (Equation 7), where m is a class-separation margin, z_i, z_j are latent points, and β scales the domain loss. This domain-informed dimensionality reduction is critical, enhancing the interpretability of otherwise complex data and providing clear, visually discernible differentiation between leakage and non-leakage conditions. This custom-designed domain-specific loss function is instrumental in enforcing clear separation within the latent space.

$$\mathcal{L}(\theta, \phi; x) = \mathbb{E}_{q_\phi(z|x)}[\log p_\theta(x|z)] - D_{KL}(q_\phi(z|x)||p(z)) \quad (6)$$

$$L_{\text{domain}} = \beta \sum_{i \neq j} \max(0, m - \|z_i - z_j\|)^2 \quad (7)$$

Subsequently, the binary SVM classifier leverages this simplified latent representation for classification into leakage or non-leakage using a kernel function $K(\cdot, \cdot)$ (Equation 8), where α_i are dual coefficients, y_i are labels, and b is the bias term. The SVM achieves exceptional accuracy, exceeding 98%, validated against the extensive real-world data from the UK. This approach ensures reliable, rapid leakage detection, which is critical for timely interventions and operational efficiency.

$$f(x) = \text{sgn} \left(\sum_{i=1}^N \alpha_i y_i K(x, x_i) + b \right) \quad (8)$$

The leakage localization stage employs an integrated graph attention network (GAT) [14] and transformer model architecture [13] to predict flow and leakage probability at a specific node given information at other nodes. The spatial component uses GAT layers to dynamically encode complex spatial relationships via attention mechanisms, effectively capturing interactions dictated by topological and geographical proximities. Topological distances (via BFS) and Haversine distances explicitly quantify these spatial relationships, enhancing model capability to accurately represent infrastructure-specific dependencies. Multi-head attention mechanisms in the GAT facilitate comprehensive and diverse spatial feature extraction, ensuring robust spatial relationship modeling. GAT computes spatial dependencies as in Equation (9) with attention weights α_{ij} from Equation (10), where W is a weight matrix, and a is a learnable vector.

$$h'_i = \sigma \left(\sum_{j \in \mathcal{N}(i)} \alpha_{ij} W h_j \right) \quad (9)$$

$$\alpha_{ij} = \frac{\exp(\text{LeakyReLU}(a^T [W h_i \| W h_j]))}{\sum_{k \in \mathcal{N}(i)} \exp(\text{LeakyReLU}(a^T [W h_i \| W h_k]))} \quad (10)$$

The temporal component utilizes transformer module, which proficiently handle long-range temporal dependencies inherent in water flow data, enhancing temporal predictive capabilities. Transformer layers handle temporal structure using scaled dot-product attention (Equation 11), where Q , K , and V are query, key, and value matrices, and d_k is the dimension of the key.

$$\text{Attention}(Q, K, V) = \text{softmax} \left(\frac{QK^T}{\sqrt{d_k}} \right) V \quad (11)$$

This hybrid GAT-transformer architecture was rigorously evaluated across diverse sensor deployment configurations, systematically varying input-output pipe ratios (e.g., 10-to-5, 12-to-3, 14-to-1 configurations), achieving consistent predictive accuracy evidenced by an average coefficient of determination (R^2) exceeding 0.9. These extensive empirical evaluations underscore the model’s reliability, scalability, and accuracy across various operational scenarios.

The proposed unified framework significantly advances water distribution management practices from traditionally reactive to proactive strategies, leveraging comprehensive data-driven insights. Its deployment facilitates near-real-time actionable intelligence, enabling water utilities to optimize operational decisions effectively, allocate resources judiciously, and minimize service disruptions through targeted interventions. The explicit incorporation of weather and environmental data further bolsters predictive reliability, allowing dynamic adaptation to variable climatic conditions, essential under ongoing climate change scenarios. Interpretability is systematically embedded throughout the model via Shapley additive explanations [9] and attention mechanisms, ensuring transparency, facilitating stakeholder trust, and enabling informed, evidence-based decision-making.

Collectively, this scientifically rigorous and comprehensive framework substantially enhances the resilience, operational sustainability, and reliability of water distribution networks. By systematically addressing forecasting, detection, and localization, the proposed approach not only mitigates immediate operational inefficiencies but also strategically safeguards essential water resources, contributing critically to their long-term sustainability and security for future generations.

REFERENCES

- [1] Water leakage statistics in the uk, 2023. Ofwat and UK Water Services Reports.
- [2] D. Bahdanau, K. Cho, and Y. Bengio. Neural machine translation by jointly learning to align and translate. In *International Conference on Learning Representations (ICLR)*, 2015.
- [3] L. Breiman. Random forests. *Machine Learning*, 45(1):5–32, 2001.
- [4] C. Cortes and V. Vapnik. Support-vector networks. In *Machine Learning*, volume 20, pages 273–297, 1995.
- [5] P. J. Huber. Robust estimation of a location parameter. *Annals of Mathematical Statistics*, 35(1):73–101, 1964.
- [6] R. E. Kalman. A new approach to linear filtering and prediction problems. *Transactions of the ASME–Journal of Basic Engineering*, 82(1):35–45, 1960.
- [7] D. P. Kingma and M. Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2014.
- [8] F. T. Liu, K. M. Ting, and Z.-H. Zhou. Isolation forest. In *Proceedings of the 8th IEEE International Conference on Data Mining*, pages 413–422, 2008.
- [9] S. M. Lundberg and S.-I. Lee. A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems*, 30, 2017.

- [10] L. McMillan, J. Fayaz, and L. Varga. Flow forecasting for leakage burst prediction in water distribution systems using long short-term memory neural networks and kalman filtering. *Sustainable Cities and Society*, 99:104934, 2023.
- [11] L. McMillan, J. Fayaz, and L. Varga. Domain-informed variational neural networks and support vector machines based leakage detection framework to augment self-healing in water distribution networks. *Water Research*, 249:120983, 2024.
- [12] M. Schuster and K. K. Paliwal. Bidirectional recurrent neural networks. *IEEE Transactions on Signal Processing*, 45(11):2673–2681, 1997.
- [13] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.
- [14] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Liò, and Y. Bengio. Graph attention networks. In *International Conference on Learning Representations (ICLR)*, 2018.