Early Warning Detection for Water Management Systems

Dr. Meiyi Ma, Vanderbilt University

Dr. Ayan Mukhopadhyay, Vanderbilt University

1. Introduction

Efficient water service management requires timely detection of irregularities, such as valve malfunctions, pressure drops, or unexpected usage patterns. Traditional methods often rely on reactive measures, generalized thresholds and post-analysis using simulations, which can result in delays and inefficiencies.

Vision for a research-centered pilot project: We propose to develop an AI-powered anomaly detection system by leveraging historical data analysis, spatial-temporal modeling, and graph neural networks (GNNs) integrated with simulation tools to enhance the reliability and efficiency of water infrastructure management. In particular, the system will analyze temporal-spatial patterns in historical water service data and real-time sensing data, such as pressure levels, flow rates, and consumption metrics, to detect anomalies at a fine granularity. This proposal leverages cutting-edge AI techniques to enhance the resilience and efficiency of water service systems. By starting with a focused case study, we aim to develop a robust, scalable solution that can transform water infrastructure management.

Our team has been working on various projects on developing human-centered and responsible AI-powered solutions for smart cities [1-16], including safe and reliable AI-based intelligent systems, physics-informed machine learning, deep learning voice analysis and mobile / NLP / AI systems for emergency response and healthcare, integrating heterogeneous city services, etc. In this project, we will work closely with the engineers and domain experts in water service optimization to gather the data and develop the AI model.

2. Proposed Solution

We propose six objectives for this two-year project. Specifically, Objectives 1 through 4 are centered on the development and evaluation of the system, and will be carried out during the first year. Objectives 5 and 6 focus on transitioning the system to real-world deployment, incorporating real-time data inputs and developing a dashboard to visualize detection results.

Objective 1. Domain Expert-Centered Design and Historical Data Analysis.

We plan to start with a domain expert-centered design and comprehensive historical data analysis of a one-year dataset from a small pilot area. This analysis will focus on key hydraulic parameters such as water pressure, flow rates, and usage patterns to identify baseline behaviors and historical anomalies. These insights will serve as foundational inputs for model development.

To start with, we will work with a domain expert to identify the existing data, including the data being used in the simulation and additional real-time ones if exists. Data preprocessing and cleaning will be performed to ensure data integrity. This involves handling missing values and inconsistencies in sensor data, synchronizing time-series data from different sources, and filtering out erroneous readings using statistical and rule-based approaches.

Secondly, we will focus on studying the anomaly patterns from the historical data. One challenge is the sparsity of the anomaly data, i.e., there are many more normal data than anomaly data in the dataset. We plan to take two approaches targeting this problem: (1) data augmentation through statistic and simulation-based data synthesizes, and (2) anomaly property identification and knowledge injection.

Objective 2. Developing a GNN-based Deep Learning Model.

To effectively detect anomalies within the water distribution network, a Graph Neural Network (GNN)-based deep learning model will be developed to learn both spatial and temporal dependencies within the water distribution system.

Background: Graph Neural Networks (GNNs) are a class of deep learning models specifically designed to process and analyze graph-structured data. Unlike traditional neural networks that operate on Euclidean data (e.g., images or sequences), GNNs excel at modeling complex relational structures such as social networks, molecular graphs, and citation networks. In recent years, GNNs have also shown strong potential in infrastructure applications, particularly in water distribution systems (WDS)**,** where nodes represent junctions or tanks and edges represent pipes. By leveraging the inherent graph topology of WDS, GNNs enable tasks such as leak detection, sensor placement optimization, anomaly detection, and flow/pressure prediction in an end-to-end learning framework.

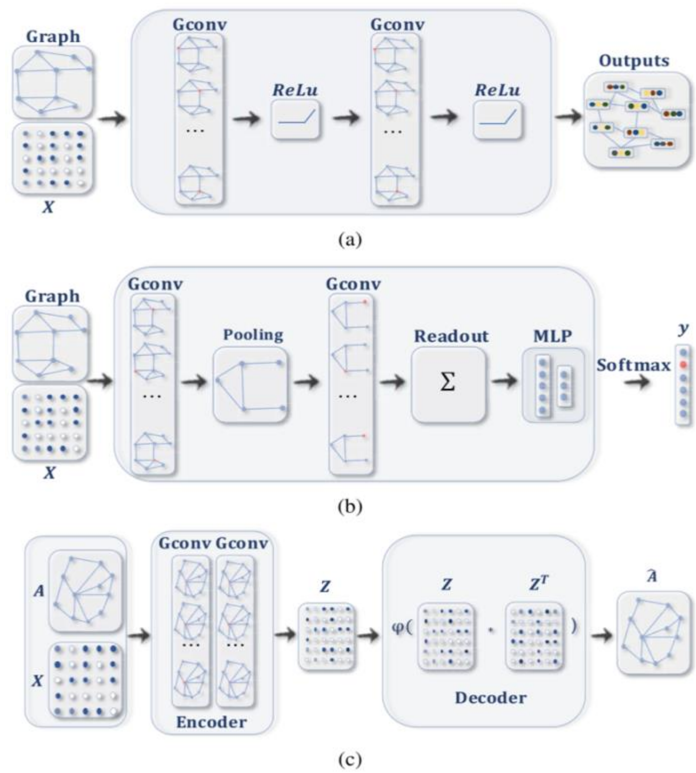

The core mechanism underlying most GNNs is information passing. In each layer, a node updates its feature representation by aggregating information from its neighbors. This process can be generalized into three steps: *message computation*: each node collects messages from its neighbors. *Aggregation*: these messages are combined (e.g., summed, averaged). *Update:* the node updates its own embedding using the aggregated message, typically through a neural network. The equation is: $x_v^{(k)} = \mathrm{UPDATE}^{(k)}\left(x_v^{(k-1)},\ \mathrm{AGGREGATE}^{(k)}\left(\left\{x_u^{(k-1)}: u \in \mathcal{N}(v)\right\}\right)\right)$, where $x_v^{(k)}$ is the node embedding at layer k, and $\mathcal{N}(v)$ is the neighborhood of node v. Through multiple message-passing layers, a node can incorporate information from multi-hop neighbors, allowing GNNs to capture both local and global graph structure. One of the distinguishing features of GNNs is they are designed for various analytic tasks depending on the target granularity.

As illustrated in Figure 1, GNNs can be designed for node-level, edge-level, and graph-level tasks. Node level tasks (figure 1b): Node-level GNNs aim to predict properties of individual nodes, such as for node classification or regression. For instance, given a person's social network graph, a GNN can infer the likely interests or attributes of their friends. In this case, the GNN aggregates information from neighboring nodes through several graph convolution layers and then applies a softmax classifier to produce the class probabilities for each node. Edge level tasks (figure 1c): These GNNs are used for link prediction and edge classification, where the goal is to infer the existence or type of connection between two nodes. For example, given two individuals in a network, the model can predict whether they are friends, colleagues, or strangers by learning a representation of each node and computing a function (e.g., dot product or MLP) over their embeddings. Graph level tasks (figure 1a): In graph-level GNNs, the entire graph is treated as a single data instance. These models are often used in chemistry, where molecules are represented as graphs and the task is to classify the molecule (e.g., toxic vs. non-toxic). After node-level convolutions, the GNN aggregates node features into a global graph representation (using techniques like sum, mean, or pooling), which is then passed to a classifier.

Proposed Plan: The first step is to construct a graph representation of the water network, where nodes will represent critical system components such as valves, junctions, tanks, and pumps, while edges will denote physical connections between these components, as shown in figure 2. Node features will be assigned based on time-series data (e.g., pressure, flow, and consumption) and connectivity attributes, while dynamic adjacency matrices will be generated to capture evolving relationships over time.

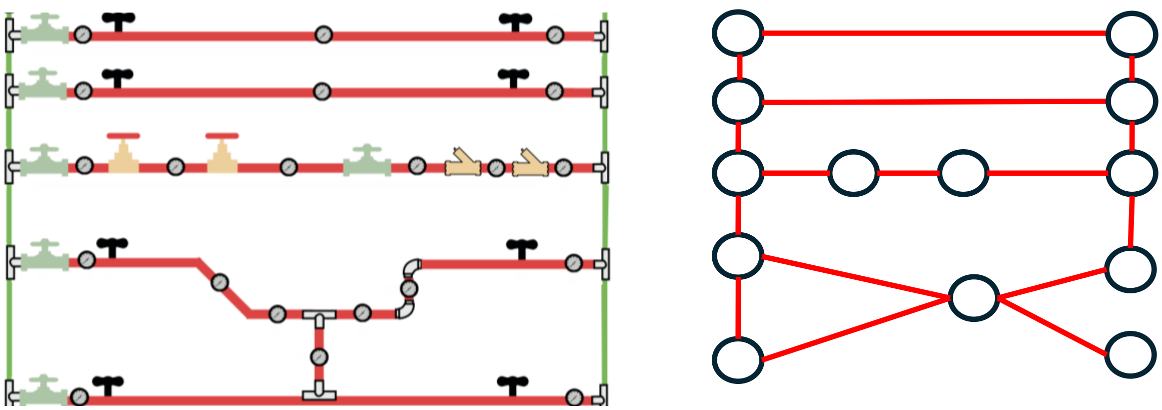

Figure 2: An example of converting a water pipeline system into a graph representation. On the left, each valve or junction is modeled as a node, while the pipes connecting them as edges.

The deep learning architecture will utilize Graph Convolutional Networks (GCNs) or Graph Attention Networks (GATs) to extract spatial relationships between nodes. Additionally, Recurrent Graph Neural Networks (RGNNs) or Temporal Graph Networks (TGNs) will be implemented to incorporate time-dependent behavior. An encoder-decoder structure will be applied to model normal operational patterns and detect deviations from expected behaviors. The anomaly detection mechanism will involve training the model in a semi-supervised or unsupervised manner to recognize deviations from learned patterns. Reconstruction-based anomaly detection techniques, such as Autoencoders and Variational Autoencoders, will be explored, alongside contrastive learning methods to enhance the model's ability to differentiate between normal and anomalous patterns. The GNN-based approach will provide an interpretable and scalable framework for anomaly detection in water distribution networks.

Objective 3. Simulation Integration and Physics Informed GNN model.

We will first carefully evaluate the GNN model proposed by Objective 2. To enhance the model's accuracy and robustness, leveraging our previous work on physics-informed machine learning, we will explore a physics-informed machine learning approach by integrating hydraulic simulations into the deep learning framework.

Specifically, we plan to build a hybrid model will by combining physics-driven models with data-driven GNN-based architectures. We will work closely with the domain expert to better understand the physics models in the current simulation and explore the best approach to incorporate it into the designed the GNN model.

The possible principles we will explore include, (1) Mass conservation: it ensures that the total inflow to a junction equals the total outflow. This is expressed as: $\sum Q_{in} = \sum Q_{out}$ Incorporating this constraint helps enforce the physical law of continuity in the GNN model, particularly at junction nodes in the water network. (2) Energy conservation along pipeline segments, typically modeled using this equation, which accounts for head loss due to friction: $h_f = f \cdot \frac{L}{D} \cdot \frac{v^2}{2g}$ This equation helps capture pressure drops due to flow resistance. (3) Hazen–Williams equation: In practice, water utilities often rely on this equation for its simplicity and effectiveness in modeling flow in pressurized systems: $h_f = 10.67 \cdot \frac{L}{C^{1.85} D^{4.87}} Q^{1.85}$. And it is often used in WDS datasets when combine with physics informed GNNs.

By embedding these physical constraints into the GNN architecture—either through the message-passing process or as part of the loss function—we aim to create a robust and interpretable model that not only fits observed data but also adheres to the fundamental laws of fluid dynamics in real-world water systems.

Objective 4. Pilot Case Study and Evaluation Plan.

Before scaling the solution to larger water networks, a controlled pilot case study will be conducted using the one-year dataset from the selected small area.

The first stage of this case study will focus on model training and validation. Data will be partitioned into training, validation, and test sets to optimize model performance. The model's effectiveness will be evaluated under different operational conditions, comparing normal and anomaly scenarios to ensure it can accurately detect deviations. Sensitivity analysis will be performed to determine key influencing factors that impact system performance.

Validation and performance benchmarking will be conducted by comparing GNN model predictions with simulation results to assess consistency. The robustness of the model will be evaluated against noise, sensor drift, and missing data scenarios, with performance metrics such as Precision, Recall, F1-score, used to measure effectiveness.

Specifically, we divide the evaluation into two phases with public datasets (as shown in Table 1) and Nashville datasets. (1) Hanoi Water Distribution Network [1]: the Hanoi network is one of the most used benchmark datasets for testing water distribution algorithms and graph-based models. It is a synthetic network developed to simulate a real urban water supply system and is often used in optimization, leak detection, and pipe failure studies. (2) L-Town Water Network [2]: The L-Town dataset is another well-known water distribution network used in research. It represents a realistic medium-sized network and is more complex than the Hanoi network. L-Town is often used in anomaly detection, simulation of sensor placement strategies, and resilience analysis.

	Hanoi [1]	L-Town [2]
Nodes	31	388
Pipes	43	429
Reservoirs	3	1
Pipe Attributes	Length, diameter, roughness	—
Available Data	Hydraulic heads, pressures	Flow, pressure, demand
Time-Series	—	Yes (multi-day simulations)

Table 1: Detailed features of the two datasets. "—" indicates that the corresponding feature is not available in the dataset.

Objective 5. Integration of Real-Time Data.

Objective 5 focuses on the critical task of integrating real-time data into large-scale water distribution networks to enhance the reliability and responsiveness of these systems. This is a complex and essential step in transitioning from theoretical modeling to practical implementation. In many engineering and scientific disciplines that deal with physical systems, one of the most persistent challenges is known as the simulation-to-reality gap. This refers to the discrepancy that often arises when a computational model, which performs exceptionally well during testing with simulated or controlled data, fails to maintain the same level of performance when applied in real-world scenarios. The root of this problem lies in the inherent unpredictability, noise, and variability of actual environmental conditions—factors that are difficult to fully capture in simulations. In the context of water distribution systems, this gap can lead to suboptimal decision-making and inefficiencies when the model is deployed in live settings. Therefore, our primary objective in this phase is to bridge that gap by refining our predictive model to handle the complexities of real-time, real-world data. This includes equipping the model to process and interpret data that may be noisy, incomplete, or subject to significant uncertainty. By developing robust algorithms and incorporating adaptive mechanisms, we aim to ensure that the model can deliver accurate and dependable predictions, even under dynamic and less predictable operating conditions.

Objective 6. Dashboard: Display Real-Time Prediction and Information.

In this stage, we aim to design and develop an interactive dashboard that presents real-time predictions and key system information generated by the GNN-based model. The purpose of the dashboard is to provide an intuitive and accessible interface for engineers and operators monitor the water distribution network, interpret model outputs, and respond promptly to anomalies or emerging issues.

The dashboard will display real-time data such as flow rates, pressure levels, and predicted events (e.g., potential leaks or abnormal patterns) across the network. Predictions will be visualized spatially over the network graph, enabling users to see which nodes or pipes are affected. In addition, the dashboard will include time-series plots, color-coded alerts, and historical trend comparisons to help contextualize the predictions. Ultimately, this tool will serve as a bridge between the GNN model's predictive capabilities and actionable insights, facilitating the practical deployment of AI-driven monitoring in operational water infrastructure.

3. Deliveries

In summary, we envision below key deliveries in this project.

- A comprehensive data analysis and integration process that combines historical and real-time sensor data, guided by domain experts to ensure the inclusion of relevant operational insights.
- The primary deliverable will be a Graph Neural Network (GNN)-based deep learning model capable of detecting anomalies within the water distribution network by learning both spatial and temporal dependencies.
- The integration of physics-based simulation into the GNN framework to build a more accurate and robust hybrid model.
- Comprehensive evaluation and pilot study.
- A Dashboard presenting real-time predictions and key system information.

4. Reference

[1] An, Ziyan, Hendrik Baier, Abhishek Dubey, Ayan Mukhopadhyay and Meiyi Ma. Enabling MCTS Explainability for Sequential Planning through Computation Tree Logic. 27th European Conference on Artificial Intelligence (ECAI), 2024.

[2] Chen, Zirong, Xutong Sun, Yuanhe Li, and Meiyi Ma. "Auto311: A Confidence-guided Automated System for Non-emergency Call." 38th Annual AAAI Conference on Artificial Intelligence (AAAI 24'). 2024 February.

[3] An, Ziyan, Taylor T. Johnson, and Meiyi Ma. "Formal Logic Enabled Personalized Federated Learning Through Property Inference." Association for the Advancement of Artificial Intelligence (AAAI 24'). 2024 February.

[4] Ma, Meiyi, Himanshu Neema, and Janos Sztipanovits. "Recovery Planning." In Autonomous Intelligent Cyber Defense Agent (AICA) A Comprehensive Guide, pp. 159-182. Cham: Springer International Publishing, 2023.

[5] Zirong Chen, Isaac Li, Haoxiang Zhang, Sarah Preum, John A. Stankovic, and Meiyi Ma. CitySpec with Shield: A Secure Intelligent Assistant for Requirement Formalization. Pervasive and Mobile Computing, 2022.11.

[6] Yiqi Zhao, Ziyan An, Xuqing Gao, Ayan Mukhopadhyay, and Meiyi Ma. Fairguard: Harness Logic-based Fairness Rules in Smart Cities. IEEE/ACM International Conference on Internet-of-Things Design and Implementation (IoTDI), 2022.

[7] Zirong Chen, Isaac Li, Haoxiang Zhang, Sarah Preum, John A. Stankovic, and Meiyi Ma. CitySpec: An Intelligent Assistant System for Requirement Specification in Smart Cities. IEEE International Conference on Smart Computing (Smartcomp), 2022: pp. 32-39.

[8] John A. Stankovic, Meiyi Ma, Sarah M. Preum, and Homa Alemzadeh. Challenges and Directions for Ambient Intelligence: A Cyber Physical Systems Perspective. IEEE International Conference on Cognitive Machine Intelligence (CogMI), 2021.11: pp. 232-241

[9] Ma, Meiyi, John A. Stankovic, and Lu Feng. "Toward formal methods for smart cities." Computer 54, no. 9 (2021): 39-48.

[10] Yukun Yuan, Meiyi Ma, Songyang Han, Desheng Zhang, Fei Miao, John A. Stankovic, and Shan Lin. DeResolver: A Decentralized Conflict Resolution Framework with Autonomous Negotiation for Smart City Services. ACM Transaction on Cyber-Physical Systems, 6, no. 4 (2022): 1-27.

[11] Ma, Meiyi, John Stankovic, Ezio Bartocci, and Lu Feng. "Predictive monitoring with logic-calibrated uncertainty for cyber-physical systems." ACM Transactions on Embedded Computing Systems (TECS) 20, no. 5s (2021): 1-25.

[12] Ma, Meiyi, Ezio Bartocci, Eli Lifland, John A. Stankovic, and Lu Feng. "A novel spatial–temporal specification-based monitoring system for smart cities." IEEE Internet of Things Journal 8, no. 15 (2021): 11793-11806.

[13] Ma, Meiyi, Ji Gao, Lu Feng, and John Stankovic. "STLnet: Signal temporal logic enforced multivariate recurrent neural networks." Advances in Neural Information Processing Systems 33 (2020): 14604-14614.

[14] Ma, Meiyi, Ezio Bartocci, Eli Lifland, John Stankovic, and Lu Feng. "SaSTL: Spatial aggregation signal temporal logic for runtime monitoring in smart cities." In 2020 ACM/IEEE 11th International Conference on Cyber-Physical Systems (ICCPS), pp. 51-62. IEEE, 2020.

[15] Ma, Meiyi, Sarah M. Preum, Mohsin Y. Ahmed, William Tärneberg, Abdeltawab Hendawi, and John A. Stankovic. "Data sets, modeling, and decision making in smart cities: A survey." ACM Transactions on Cyber-Physical Systems 4, no. 2 (2019): 1-28.

[16] Ma, Meiyi, John A. Stankovic, and Lu Feng. "Cityresolver: a decision support system for conflict resolution in smart cities." In 2018 ACM/IEEE 9th International Conference on Cyber-Physical Systems (ICCPS), pp. 55-64. IEEE, 2018.

[17] Vrachimis, S.; Kyriakou, M.; Eliades, D.; and Polycarpou, M. 2018. LeakDB: A benchmark dataset for leakage diagno- sis in water distribution networks description of benchmark. In Vol 1 of Proc., WDSA/CCWI Joint Conf. Kingston, ON, Canada: Queen's Univ.

[18] Vrachimis, S. G.; Eliades, D. G.; Taormina, R.; Ostfeld, A.; Kapelan, Z.; Liu, S.; Kyriakou, M.; Pavlou, P.; Qiu, M.; and Polycarpou, M. M. 2020. BattLeDIM: Battle of the leak- age detection and isolation methods. In CCWI/WDSA Joint Conf.

BUDGET JUSTIFICATION
Vanderbilt University
Early Warning Detection for Water Management Systems

A. SENIOR PERSONNEL

Meiyi Ma Principal Investigator

Dr. Ma will serve as Principal Investigator. They will dedicate 0.5 CAL effort months to the project. They will be responsible for leading the development of the whole project and managing the team and collaboration with the Water Service Department.

Ayan Mukhopadhyay, CO-PI

Dr. Mukhopadhyay will serve as CO-PI. They will dedicate 0.50 CAL effort months to the project. They will be responsible for leading the development of dashboard and real-time integration and collaborating on the algorithm development and evaluation.

Total Senior Personnel (A): $29,292

B. OTHER PERSONNEL

TBA 1, Research Engineer

One TBA Research Engineer will serve as a Research Engineer. They will dedicate 6.0 CAL in effort months to the project in Year 2. They will be responsible for piloting and deploying the system and developing the dashboard.

One (1) TBD Graduate Student

One graduate student will serve on this project and dedicate 12.0 CAL effort months in Year 1 and 3.0 CAL in Year 2. They will be responsible for developing and evaluating the AI-based anomaly detection algorithm.

Total Other Personnel (B): $91,210

Total Salaries and Wages (A + B): $120,502

C. FRINGE BENEFITS

The Vanderbilt University Faculty and Staff fringe benefit rate is 25.1% for full-time faculty and staff and 10.2% for temp. staff for FY25 (details can be found here). Salary and wage figures are based on yearly salaries using Vanderbilt University Human Resource categories. Salary is increased by 3.0% annually.

Total Fringe Benefits (C): $17,952

Total Salaries, Wages and Fringe Benefits (A + B + C): $138,454

D. OTHER DIRECT COSTS

Graduate Tuition

35% of VU Graduate Student academic year tuition is charged for the graduate student but is not included in the Facilities and Administrative base. The tuition charge is calculated based on graduate student FTE. The rate is $2,419 per credit hour with a minimum of $200. $15,310 is charged per full-time student in Year 1. There is a 3% inflation rate each year.

Total Graduate Tuition: $18,697

Graduate Insurance and Health Fee

Graduate Health Insurance and Student Health Fee are charged for the graduate student to the project. The Graduate Health Insurance charge is calculated based on grad student FTE. There is a $4,314 cost for annual health insurance with $441 health per semester and $130 for the summer. $5,326 is charged per full-time student in Year 1. There is a 7% inflation rate each year.

Total Graduate Insurance: $6,472

Total Other Direct Costs (G): $25,169

E. Total Direct Costs (A through G): $ 163,623

F. INDIRECT (F&A) COSTS

The current federally negotiated F&A rate with the Department of Health and Human Services for Vanderbilt University is 58.5% for FY25. The current rate agreement date is 5/15/24. F&A is charged on the Modified Total Direct Costs (MTDC). Graduate Tuition, capital equipment, participant support costs and subcontractor costs in excess of $25,000 is excluded from the MTDC base used to calculate F&A costs.

Modified Total Direct Cost Base: $ 144,926
I. Indirect (F&A) Costs: $84,782

J. Total Direct and Indirect Costs (H + I): $248,405