*ai*WATERS: An Artificial Intelligence Framework for the Water Sector

Darshan Vekaria

Thesis submitted to the Faculty of the

Virginia Polytechnic Institute and State University

in partial fulfillment of the requirements for the degree of

Master of Science

in

Computer Science and Applications

Sunil Sinha, Co-chair

Naren Ramakrishnan, Co-chair

Anuj Karpatne

May 1, 2023

Blacksburg, Virginia

Keywords: Artificial Intelligence, Water Sector, Machine Learning, Digital Transformation.

*ai*WATERS: An Artificial Intelligence Framework for the Water Sector

Darshan Vekaria

(ABSTRACT)

The ubiquity of Artificial Intelligence (AI) and Machine Learning (ML) applications has led to their widespread adoption across diverse domains like education, self-driving cars, healthcare, and more. AI is making its way into the industry, beyond research and academia. Concurrently, the water sector is undergoing a digital transformation, driven by challenges such as water demand forecasting, wastewater treatment, asset maintenance and management, and water quality assessment. Water utilities are at different stages in their journey of digital transformation, and its decision-makers, who are non-expert stakeholders in AI applications, must understand the technology to make informed decisions. The non-expert stakeholders should know that while AI has numerous benefits to offer, there are also many challenges related to data, model development, knowledge integration, and ethical concerns that should be considered before implementing it for real-world applications. Civil engineering is a licensed profession where critical decision-making is involved. Failure of critical decisions by civil engineers may put their license at risk, and therefore trust in any decision-support technology is crucial for its acceptance in real-world applications. This research proposes a framework called *ai*WATERS (Artificial Intelligence for the Water Sector) to facilitate the successful application of AI in the water sector. Based on this framework, we conduct pilot interviews and surveys with various small, medium, and large water utilities to capture their current state of AI implementation and identify the challenges faced by them. The research findings reveal that most of the water utilities are at an early stage of

implementing AI as they face concerns regarding the blackbox nature, trustworthiness, and sustainability of AI technology in their system. The *ai*WATERS framework is intended to help the utilities navigate through these issues in their journey of digital transformation.

*ai*WATERS: An Artificial Intelligence Framework for the Water Sector

Darshan Vekaria

(GENERAL AUDIENCE ABSTRACT)

The widespread adoption of Artificial Intelligence (AI) and Machine Learning (ML) in various industries like education, self-driving cars, healthcare, and more has spurred interest in its potential application in the water sector. As the water sector undergoes a digital transformation to address challenges such as water demand forecasting, wastewater treatment, asset management, and water quality assessment, water utilities need to understand the benefits and challenges of AI technology. Automating water sector operations through AI involves high risk as it has a huge ecological, economic, and sociological impact on society. Water utilities are non-expert end users of AI and they should be aware of its challenges such as data management, model development, domain knowledge integration, and ethical concerns when implementing AI for real-world applications. To address these challenges, this research proposes a framework called *ai*WATERS (Artificial Intelligence for the Water Sector) to help water utilities successfully apply AI technology in their system. We conduct pilot interviews and surveys with small, medium, and large water utilities across the United States to capture their current AI practices and challenges. The research results led us to find that water utilities are still at an early stage of adopting AI in their system and are faced with issues such as blackbox nature of the technology, its trustworthiness for real-world application, and sustainability at the utilities. We believe that *ai*WATERS will serve as a relevant guide for water utilities and will help them overcome current AI-based challenges.

Acknowledgments

I would like to express my heartfelt gratitude to my advisor, Dr. Sunil Sinha, for his invaluable guidance and unwavering support in developing the *ai*WATERS framework. Throughout the past two years, Dr. Sinha has been an incredible mentor, always encouraging me to strive for excellence. Without his constant motivation, completing this thesis would not have been possible. I am grateful to the Water Research Foundation (WRF) for funding this research and to Jacobs Engineering for their review and inputs related to this study. I am highly grateful to the water utilities who participated in this study. I would also like to thank Dr. Anuj Karpatne, who served as my instructor for the Data Analytics course during my first semester at Virginia Tech. The knowledge and skills I gained from this course have been instrumental in laying the foundation for the development of aiWATERS. I would also like to thank Dr. Naren Ramakrishnan for being the Co-Chair of my thesis committee and providing his valuable input. I would like to express my appreciation to the graduate students in the Sustainable Water Infrastructure Management (SWIM) Center for their constant support in development of the *ai*WATERS framework. Last, but not least, I am deeply grateful to my family and friends for supporting and encouraging me throughout my graduate studies.

Contents

List of Figures **x**

List of Tables **xii**

1 Introduction **1**

1.1 Background 3

1.2 Goals and Objectives 4

1.2.1 Research Questions 5

1.2.2 Methodology 6

2 Literature and Practice Review **9**

2.1 Literature Review 9

2.1.1 AI applications in Built Subsystem 10

2.1.2 AI applications in Natural Subsystem 17

2.1.3 AI Applications in Socio-Economic Subsystem 19

2.2 Practice Review 22

2.2.1 Metropolitan Water Reclamation District (MWRD) of Greater Chicago 22

2.2.2 Clean Water Services, Oregon 25

2.2.3 Hampton Roads Sanitation District (HRSD), VA 27

2.3 Various Use Cases of AI in the Water Sector 28
2.4 Data Driven & Model Driven Models in the Water Sector 31
2.5 Limitations of AI in the Water Sector 33
2.6 Framework Validation and Verification 35
2.6.1 *ai*WATERS Framework Verification 35
2.6.2 aiWATERS Framework Validation 36

3 *ai*WATERS: The Framework **37**

3.1 Pillars of *ai*WATERS 37
3.2 *ai*WATERS Understanding and Benefits 41
3.3 *ai*WATERS Application Goals 43
3.3.1 Using AI at Different Levels of System 45
3.3.2 Various Categories of AI Deployment 46
3.3.3 Different Modes of Building AI 47
3.4 *ai*WATERS Data Readiness 48
3.4.1 Evaluating Quality and Quantity of Collected Data 49
3.4.2 Data Preprocessing 53
3.4.3 Data Augmentation and Class Imbalance 56
3.5 *ai*WATERS Knowledge Integration 59
3.5.1 Extracting Domain Knowledge for AI Applications 61

3.5.2 Knowledge based selection of input and output features for the AI Application 64
3.6 *ai*WATERS Model Development 65
3.6.1 Selecting Input and Output Parameters for the AI Model 65
3.6.2 AI Model Selection and Hyperparameter Optimization 67
3.6.3 Model Generalization and Maintenance 72
3.7 *ai*WATERS Decision Support 74
3.7.1 Ensuring Trustworthiness of AI Decision Support 76
3.7.2 Involving Human in the Loop for Decision Support 78
3.7.3 Using Effective Visualizations for Decision Support 80
3.8 *ai*WATERS Implementation 81

4 Results **87**
4.1 RQ1: Are utilities willing to develop AI solutions and use it across the system? 89
4.2 RQ2: How do water utilities of different scale compare to each other in terms of AI implementation? 91
4.3 RQ3: What aspects of the AI technology are considered as the major challenges by water utilities? 100

5 Discussion **102**

6 Future Work **106**

7 Conclusions **108**

Bibliography **109**

List of Figures

1.1 Research Methodology for *ai*WATERS 6

3.1 Seven Pillars of *ai*WATERS. 38

3.2 Architecture of *ai*WATERS. 39

3.3 Water System of Systems 44

3.4 Different levels of AI Applications 45

3.5 Time consumption for various tasks in building AI models 49

3.6 Hierarchy of Data Source Reliability 50

3.7 5Vs of Big Data 52

3.8 Data Preprocessing Steps 57

3.9 Knowledge and Data driven Models 61

3.10 Domain Knowledge Integration in AI 62

3.11 Proposed Model Selection Process for Utilities 67

3.12 AI Trustworthiness guidelines 77

3.13 Types of Human involvements in an AI System 79

4.1 Scale and Geography of Piloting Utilities 88

4.2 Current AI usage in Utilities 89

4.3 Willingness to investigate AI solutions in near future 90

4.4 Intent to use AI for each and every process in the future 90

4.5 Mode of building AI across utilities 92

4.6 Data Quantity Evaluation by Utilities 93

4.7 Data Preprocessing by Water Utilities 94

4.8 Domain Knowledge Extraction at Utilities of all scales 95

4.9 Large scale utility preference for starting point of AI models 95

4.10 Model Experimentation Preference 96

4.11 Consideration of Model Development Tradeoffs 96

4.12 Do utilities perform hyperparameter tuning and testing? 97

4.13 Frequency of Updating AI Models at Large Scale Utilities 97

4.14 Do respondents have complete knowledge regarding the AI model? 98

4.15 Trustworthiness Factors for Water Utilities 99

4.16 Do utilities involve experts from multidisciplinary fields for decision support? 99

4.17 Major AI challenges as per water utilities 100

List of Tables

3.1 AI in Academia vs Industry 42

3.2 Key Tips for AI Model Development 68

3.3 Stakeholders vs Relevant AI Usecases for Visualizations [1] 81

Chapter 1

Introduction

With the advancement in sensory technologies, understanding of long-term performance of various civil infrastructure assets and developments in the field of computation and processing, the water sector is evolving rapidly, and various advancements are taking place in the way water is managed and used. Utilities deal with several issues related to water, like conservation of water resources, optimal energy management, reducing water consumption, recovery of nutrients in water and wastewater, separation of wastewater resources, control and automation [2]. Utilities are in the pursuit of keeping up with the latest technologies to efficiently and sustainably manage their watershed. The use of these technologies has also led to generation of huge volumes of unstructured and complex asset data during their life-cycle. However, dealing with this data is not a straightforward task as the data generated in real-world is complex and contains a variety of issues like lack of standardization, metadata, and validity, among other issues.

The era of internet, big data and machine learning has led the water utilities to undergo a digital transformation. The proliferating use of AI across various industries has also found its presence in the water sector. Water utilities can leverage AI and big data analytics to make intelligent data driven decisions. A study by International Water Association (IWA) suggests that digitalization will enable this transition and AI will be the key technology to digitalize the water utilities [3]. The scope of AI for water utilities is huge. AI can enable intelligent decision making for the conventional processes taking place at the utilities (pipe failure prediction, water demand forecasting, wastewater treatment, etc). AI can also help the utilities in achieving enhanced business intelligence by providing assistance in tasks like data integration, smart visualizations, human resource management [4].

Thus, there are arguments in favor of using AI and Machine learning (ML) in the water sector to find the complex nonlinear patterns in the data and support decisions. It can be found that various other industries like healthcare, self driving cars, education, security etc have made significant progress in terms of implementing AI in their system. Relatively, water sector is still in nascent stage of adopting AI in its system. This can be attributed to the fact that many AI models developed in literature are black-box models [5] and do not offer insights into how the algorithm reached a particular result.

Civil engineering is a licensed profession and therefore, trust in any decision support model is critical for acceptance in real-world applications. License of practice for engineers can be at risk if any decision taken by AI leads to a legal issue. Therefore, it is important to identify and discuss the true capabilities of AI. This will help the decision makers realize opportunities, challenges, and risks of implementing AI for modeling different processes within a water cycle.

Most of the open source literature publicly available talks about the state of the art AI models and corresponding results achieved for water utility operations like asset performance management, water demand forecasting, wastewater treatment process, etc. However, these are academic experiments and there are very few resources which describes how a utility uses AI in their system. Water utilities also collaborate with AI vendors and academic institutes to achieve AI driven goals. While such collaborations may provide short term benefits to the utilities, in-house knowledge of AI is also important in order for the utilities to understand the shortcomings of AI technology and to optimally integrate domain knowledge in the AI models.

*ai*WATERS, the framework proposed in this work, provides a guide for the water utilities which can help them to understand key concepts, benefits, limitations, challenges of AI and how they correlate to the water sector. The framework is built based on seven pillars of AI: Understanding AI and its Benefits, Data Readiness, Knowledge Integration, Model Development, Decision Support and Implementation.

Based on this framework, we also conduct pilot interviews and survey with water utilities of different scales to capture their AI practices and assess their maturity of AI implementation. Our goal is to not only capture current AI practices of water utilities, but to also help them improve their maturity in AI through *ai*WATERS.

1.1 Background

Civil engineering is a broad field that encompasses numerous subdomains, such as structural engineering, transportation engineering, geotechnical engineering, environmental engineering, and water resources engineering. Each subdomain deals with unique challenges, but all share a common goal of designing, constructing, and maintaining infrastructure that supports the needs of society. Of these subdomains, water infrastructure engineering focuses on the management and distribution of water resources, an essential element for the survival of human beings, wildlife, and the ecosystem.

The importance of water infrastructure in civil engineering is undeniable. Water infrastructure is responsible for delivering clean water to homes, businesses, and industries, as well as for collecting and treating wastewater. The failure of water infrastructure can result in severe consequences, making it crucial to implement state-of-the-art technology to manage and maintain water systems.

The emergence of artificial intelligence (AI) has created new opportunities for water utilities to improve the efficiency and effectiveness of their operations. AI can be applied to various aspects of water infrastructure, such as leak detection, water quality monitoring, predictive maintenance, demand forecasting, and optimization of energy consumption. However, the adoption of AI in the water sector is not without challenges, such as the availability of data, the integration of domain knowledge, the development of accurate models, and the need for human in the loop.

This report proposes a theoretical framework (and not algorithm) that aims to addresses the challenges faced by water utilities in adopting AI technology. It will help those from the AI domain who seek to understand the unique challenges of applying AI to the water sector, and those from the civil engineering domain who seek to implement AI in the water infrastructure. The proposed framework aims to bridge the gap between these two domains by providing a comprehensive understanding of AI applications in the water sector.

1.2 Goals and Objectives

The goal of this report is to facilitate the successful application of artificial intelligence (AI) in the water sector.

Water infrastructure is vast as it has various components such as drinking water, wastewater, stormwater, and others. We study the AI applications across each of these components by dividing their applications in the category of Natural Subsystem, Built Subsystem and Socio-Economic Subsystem of the water infrastructure. We carry out an extensive literature and practice review to capture the current challenges and limitations of AI in the water sector. The *ai*WATERS framework proposed in this report is based on the challenges identified in the literature and practice review. We use this framework to conduct pilot interviews with various water utilities to answer our research questions which revolve around the current state of AI implementation in the utilities and finding the digital divide between, small, medium and large utilities.

1.2.1 Research Questions

The analysis based on our pilot interviews and surveys with water utilities will enable us to answer the following research questions:

- **RQ1: Are water utilities willing to develop AI across their system for decision making?**

  Through this question, we aim to find what is the degree at which utilities want to integrate AI within their system. Do they wish to apply AI at each and every process within the utility with a human in the loop?

- **RQ2: How do water utilities of different size compare to each other in terms of AI implementation?**

  We divide piloting utilities into three size: small, medium and large utilities, based on the population they serve. The level of AI implementation and the best practices followed by them may differ based on their focus area, available expertise and prior experience with AI. We try to capture any pattern of differences observed within these 3 categories of utility.

- **RQ3: What aspects of the AI technology are considered as the major challenges by water utilities?**

  For this research question, we try to find what according to the water utilities are the major challenges for implementing AI in their system. It will also help us to find which pillars of *ai*WATERS are more important for the utilities.

1.2.2 Methodology

The *ai*WATERS framework was developed through the following steps, also mentioned in Figure 1.1

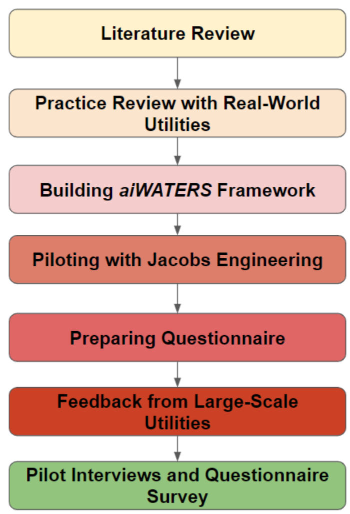

Figure 1.1: Research Methodology for *ai*WATERS

1. **Literature Review**: We conducted a comprehensive literature review by manually reviewing more than 100 articles published by water organizations, research papers from conferences and journals. The literature review captured various features of the work done by researchers, such as the domain of application in the water sector, types of AI models used and their purpose, reasoning for selecting a particular model, hyperparameter combinations tested for the model, validation process for the selected model, datasets used for training and testing, data collection and preprocessing steps followed, and programming languages/libraries used.

2. **Connecting with real-world utilities for Practice Review**: We observed that

most of the AI applications discussed in the literature are academic, and hence, to capture real-world AI practices in the water sector, we connected with some water utilities (mentioned in section 2.2) to find references of their existing AI applications.

3. **Building *ai*WATERS Fraemwork**: Our literature review and practice review helped us identify the challenges of implementing AI in the water sector. We identified seven key aspects of AI, which we call the seven pillars of AI, that can address the AI challenges to the water community. The pillars include understanding AI, its benefits and challenges as a technology, defining AI application goals, preparing AI-ready data, integrating domain knowledge in AI systems, model development, decision support and strategic planning, and implementation of AI. We named these seven pillars as our framework, *ai*WATERS (Artificial Intelligence for Water Sector).

4. **Piloting with Jacobs Engineering**: To ensure the framework's efficacy, we piloted with Jacobs Engineering during a three-day working meeting and obtained feedback from a water domain expert's perspective.

5. **Preparing the Questionnaire**: Once we finalized the framework, we prepared a questionnaire based on the 7 pillars of *ai*WATERS to capture water utilities' current practices in AI.

6. **Feedback sessions with large scale utilities**: We conducted a working meeting with HRSD and Houston water utilities to obtain feedback on the questionnaire from a utility's perspective and incorporated their suggestions. We added an extra option 'other' for each question to capture any practices that are not covered in the *ai*WATERS framework.

7. **Pilot interviews and response from Utilities**: Finally, we sent the questionnaire to various utilities to gather their current AI practices. We also conducted pilot interviews

with small, medium, and large scale utilities to understand their perspectives on the current challenges and future of AI at their utility.

The rest of this report is structured as follows: In section 2, we discuss about the various applications of AI in the water sector which are currently found in the literature. We discuss the limitations of these applications and suggest how *ai*WATERS can enable the developers in water sector to follow AI best practices and understand AI challenges in order to achieve an ideal digital transformation. In section 3, we discuss the *ai*WATERS framework which is built upon the 7 key pillars of AI. In section 4, we talk about the response received on the questionnaire survey sent to various water utilities and corresponding pilot interviews with them, to capture their AI practices and compare water utilities of small, medium and large scale with respect to their AI implementation. Section 5 talks about the inferences we make through our findings in the results section. We suggest possible future work in this research through section 6. Lastly, section 7 talks about the conclusions of this study.

Chapter 2

Literature and Practice Review

The application of Artificial Intelligence (AI) in the water sector has been rapidly growing in recent years, with many researchers and practitioners exploring its potential to improve the efficiency and effectiveness of water management. In this section, we will review both academic literature and real-world examples of AI applications in the water sector. By examining the current state of research and practice, we hope to gain insights into the strengths and limitations of AI in water management, as well as identify areas for future research and development.

2.1 Literature Review

The use of AI in the water sector is spread across the Built, Natural and Social subsystems, which makes the larger water system. Each of these sub-systems contain multiple components and elements. We define the categorization of the water system as follow:

- **Built Subsystem:** The built sub-system includes the human-made components developed for wastewater collection, stormwater capture, and treatment infrastructure. The main components of this subsystem include household/commercial/industrial, wastewater, and stormwater infrastructure.
- **Natural Subsystem:** The natural sub-system comprises all the naturally occurring

resources/ elements that exist and influence sewershed characteristics like the water, land, and climate.

- **Social Subsystem:** The social sub-system comprises the factors affecting societal aspects and having an economic impact on a sewershed and can be categorized into the community, policy, financial and economic components.

There are various AI applications associated with the component of these subsystems. The ever increasing research efforts of AI applications lacks a comprehensive study of AI on each of these subsystems. In this section, we present a summary of various applications of AI in water sector available in the literature. We also intend to highlight the limitations of these literature references to discuss how *ai*WATERS can help the water utilities in approaching their AI based solutions.

2.1.1 AI applications in Built Subsystem

Water Pipe Failure Prediction

Pipe failure analysis has been a frequently used application of AI techniques in the category of built water infrastructure. Some of the early works in this category, like survival analysis of water distribution pipes which is done by Part *et al.* [6] explored the use of proportional hazards modeling approach to obtain the survival probability of pipes. Their modeling results showed that the failure times of all the categories of iron pipes have the Weibull distribution. However, they didn't disclose the dataset used in this work.

Aslani *et al.* [7] used Boosted Regression Tree to predict pipe failure rate using the main breaks data in Tampa, Florida between the years 2015 to 2020. They identified the vulnerable areas in the city using spatial clustering analysis. The result of this analysis was translated into a categorical independent variable named hotspot level in the modeling phase.

Various parameters are used by researchers to represent pipe failure. For instance, Almheiri *et al.* [8] used risk index to represent and predict pipe failure for the city of London, Ontoario, Canada. They trained an Artificial Neural network on 2000 samples of this city. Their data contained about 70% censored data, which means that these pipes are yet to experience failure event. Using a model which is trained on only 2000 samples, with uncertainty in the nature of pipes (as these samples are yet to experinece failure) may not be ideal for translating this academic model for real world application.

A group a researchers from Iran, Shirzad *et al.* [9], compared the performance of Artificial Neural Network (ANN) and Support Vector Regression (SVR) for predicting Pipe Burst Rate in Mashhad and Mahabad districts in Iran. They found that SVR was less precise than ANN, but more robust in terms of sensitivity to input variables, allowing to capture a better system behavior. While their worked described in detail the input parameters that were tested for predicting pipe burst rate, they didn't discuss the number of samples used for training the models.

Kabir *et al.* [10] used probabilistic technique like Bayesian Model Averaging (BMA) to find failure probability of various pipes like Cast Iron, ductile iron, copper, cementitious and high-density polyethylene. They found that BMA outperformed classical regression analysis when limited pipe failure data are available. However, their data for experimentation was limited to 306 records from Greater Vernon Water, British Columbia, Canada and 199 records from City of Kelowna, Canada.

Hybrid models have also led the researchers to achieve better accuracy in their predictions. Farmani *et al.* [11] found that combining Evolutionary Polynomial Regression (EPR) with K-means clustering leads to improved prediction accuray for pipe failures when compared to compared to the non-clustered EPR approach.

Sewer Pipe Condition Assessment

Water pipes and Sewer pipes require individual analysis as there are many various factors in the environment in which both of them operates that demands an individual study of both. The content and chemicals flowing through water pipes and sewer pipes may differ significantly. This can impact the rate at which these types of pipes fail. For example, acidic content in wastewater may corrode pipes faster than non-acidic content in drinking water. Similarly, pipes located in coastal areas may corrode faster due to saltwater exposure.

Rayhana *et al.* [12] have described the traditional pipe inspection methods for sewer pipe condition assessment. The inspection of the inaccessible pipes is acheived by using artifical arms and robots which are assisted with a camera and sensor. The cameras record the internal condition of the pipes. The videos are passed to the inspection operator for survey and fault identification. Eventually, the traditional image inspection was followed by AI based condition assessment. The success of Convolutional Neural Networks (CNN) for image processing and classification tasks led researchers to explore this technique and its variants for sewer pipe condition assessment [13]. Kumar *et al.* [14] used Deep CNN (DCNN) over twelve thousand images of sewer pipelines to classify the Root intrusion, cracks and deposits. Further studies by Meijer *et al.* [15] on using DCNN for sewer pipe image classification found that the AI technique was able to outperform human operators for identifying defects in pipes.

Damaged sewer pipes are also found to be a cause of sinkholes. Kim *et al.* [16] applied logistic regression to predict sinkhole susceptibility in Seoul, South Korea. They found that length, age, elevation, burial depth, size, slope, and material of sewer pipe were the major determinants in predicting the sink hole susceptibility.

Wastewater Treatment

Wastewater treatment plants carry out various operations like effluent prediction, corrosion control, odor monitoring, effluent quality control, nutrient removal, spill event identification, removal of natural organic matter etc which can be assisted by AI techniques.

Manu *et al.* [17] used support vector machine (SVM) and adaptive neuro-fuzzy inference system (ANFIS) to predict effluent Kjeldahl Nitrogen. They used the time series data of Kavoor wastewater treatment plant (WWTP) located in Mangalore, which served a population of 440000. The data captured by them was for the period June 2014 – September 2014. However, their dataset was really small (training data had 65 data points, dated between June to August and testing data had 23 datapoints, coming from the month of September.) Thus, this application may only be considered for short term predictions.

We found that odor monitoring and control has been used in both, academia and industry. Cangialosi *et al.* [18] used Random Forest (RF) and ANN for classification and quantification of odor. Through various performance measures like R2, Root Mean Square Error (RMSE) and Normalized RMSE, then found RF outperformed ANN. However, their overall dataset was small, featuring 600 data points, which were split in a 80:20 ratio for training and testing respectively. They used data smoothing to clean the signals received from instrumental odor monitoring systems (IOMS).

A working meet with one of the officials from Metropolitan Water Reclamation District (MWRD) of Greater Chicago led us to their AI application of determing the optimal dosage of sodium hypochlorite (NaOCl) to curtail corrosion and odor by H2S. Given the amount of available data, they divided the tasks into three modules and assigned an ML model to solve each of the individual module. In Module 1, a recurrent neural network (RNN) was designed to predict wastewater characteristics. In Module 2, a random forest (RF) classifier and a support vector machine (SVM) classifier were built with the information from Module 1 along with other datasets to predict the concentrations of VFAs and H2S, respectively. Finally, in Module 3, with the information obtained from Module 2, another RF classifier was developed to predict NaOCl dosage to reduce H2S but keeping VFAs within the target range [19]. Their dataset contained 17 years of historical influential data. To predict H2S and VFA concentration, they collected online instrument data (flow, ORP (oxidation reduction potential), pH, wastewater temperature, tunnel pumping, and tunnel elevation) from SCADA system and daily precipitation data. As a part of data prepossessing, they used techniques like interpolation and oversampling for handling missing data points and class imbalance. They used Min-Max normalization for scaling their input features. They used the set of features with highest correlation with each other as the input parameters. Correlation may not be the best technique while deciding the input features and two highly correlated feature may represent the same thing and introduce redundancy in the input parameter set. All of their Machine Learning and Deep Learning analyses to predict NaOCl content for waste water plants were performed through Python, which is an open-source software, using free software packages and libraries, for example, Tensor Flow and Keras.

Water Pumping Stations

Researchers have investigated AI techniques to optimally control the flow of water at the pumping stations. Ostogin *et al.* [20] used Fuzzy logic to find the required change in control of water pumps (number of pumps to be stopped/started). They considered rate of change of level in the wet well (at interval of 5 mins), and the level in the well as the input features of their fuzzy logic. They performed this experiment on Anglian Water (UK) data. Similarly, Yagi *et al.* [21] used a combination of Fuzzy logic and genetic algorithms to determine the storm and sanitary pumping rates, but with additional input parameters like rainfall intensity, sewer water level, river water level, sewer water quality.

The latest research on pumping stations has also explored the use of deep learning techniques like CNN and LSTM (Long short term memory) to achieve water level prediction in pumping stations [22]. Zhnag *et al.* used the historical data of the pumping station project provided by the Tuancheng Lake Management Office of Beijing South-to-North Water Diversion Project for their experiment. They found that a hybrind CNN-LSTM model was able to give better results than CNN and LSTM individually. They also provided the value of hyper parameters like batch size, epochs, learning rate, dropout rate, k-fold, decay steps and decay rate for their CNN-LSTM model, making their model development reproducible by others. However, they didn't describe in detail about the feature importance and feature selection.

Accumulation of fat, oil and grease (FOG) in wastewater pumping stations is a common failure cause of failure for this infrastructure. Deep learning techniques like CNN have been explored by Moreno *et al.* [23] for estimation of FOG layer cover. The performance of their CNN architecture was found to be comparable to human classification. They used the openCV python library for camera calibration and perspective rectification of images. The CNN network was built using TensorFlow.

Green Roofs

Green roofs reduce the storm water outflows and also enables reduction in urban heat island effect, preserving the cities ecosystems and improving the urban visual amenity [24].

Tsang *et al.* [25] worked on developing a optimal irrigation strategy using machine learning, wherein they used real-time weather data to predict the soil moisture. They covered the weather data near their experimental site for the duration of August 2011 to February 2012 from the Hong Kong Observatory. They have applied ANN and Fuzzy network to predict the soil moisture content. However, there is no description about the data pre-processing methods, and model development and testing, making it difficult to reproduce this work and the models proposed in it.

One of the recent works for the Green Roof infrastructure presented by Abdalla *et al.* [26] is relatively descriptive. They compared 4 machine learning models (ANN, LSTM, KNN and M5 model tree) to simulate storm-water runoff from 16 extensive green roofs located in four Norwegian cities across different climatic zones. While they haven't described the amount of data and corresponding preprocessing techniques used for making it AI ready, they have explained their hyperparameter tuning and model testing for all 4 models in detail and found that LSTM demonstrated a better performance than other state of the art techniques.

Sludge Management

Sludge management is necessary for the utilities dealing with wastewater water treatment. Facchini *et al.* [27] used ANN with determining features as plant capacity, typology of secondary treatment of the Wastewater treatment, and the configuration of the sewage sludge treatment plant to find the optimal sewer sludge treatment process to be adopted. They are not descriptive enough about the data being used as they say that the data collected for their research was extracted from literature. The authors state that they tested their ANN model for 20 different scenarios to suggest a most effective sludge-management strategy in economic terms.

Djandja *et al.* [28] states that rising volume of sewage sludge in recent years has become a threat to the environment and that Hydrothermal carbonization (HTC) an enhance organic contaminants removal. They use feed forward neural network for predicting the nitrogen content of the hydrochar. For this work, they compiled data from 26 published papers on HTC and split it to 80:20 ratio for training and testing. They detail their neaural network model configuration as a two-layer perceptron (with one hidden layer). The limitation here is this work is limited to an academic experiment and these models may not work in real world scenario.

2.1.2 AI applications in Natural Subsystem

AI is becoming increasingly important in the management of water's natural subsystem, which includes elements such as rivers, groundwater, irrigation, and other related natural water resources. Researchers are exploring various applications of AI in Natural subsystem such as nitrogen content prediction, groundwater level and salinity prediction, surface water quality assessment and smart irrigation systems, which will be discussed in this subsection.

High concentration of nitrogen in surface water and groundwater pose a threat to humans and environment. Knoll *et al.* [29] applied AI techniques such as Regression trees, Random forest, and Boosted Regression trees for predicting groundwater nitrogen concentration for the state of Hesse in Germany. The authors discuss about how they have integrated data from various sources to collect data points for parameters like Land Use (Urban land, Arable Land, Grassland, Forest etc), Hydrology (Seepage water rate, nitrate concentration in seepage, nitrate concentration in groundwater discharge), and soil conditions (soil groups, field capacity etc). MAE, RMSE and R2 are the measures used to evaluate the performance of ML techniques and found that Random Forest achieved highest predictive performance.

Various operations of irrigation, urbanism and industrial work are associated with groundwater and scarcity of fresh groundwater has given rise to groundwater modeling [30] [31].

Recent advances in technology has led to a shift from the traditional, time consuming and expensive approaches like drilling for groundwater prediction towards digital methods like GIS (Geographic Information System), remote sensing technology and mathematical modeling. Mosavi *et al.* [32] explore the bagging and boosting methods for building ensemble methods for groundwater potential modeling. They conduct their study on the Dezekord-Kamfiruz watershed which is a part of the Fars Province in Iran. This area contains 339 groundwater resources which supply water for irrigation and drinking water purposes. The divided the groundwater resource points into a ratio of 70:30 for training and validation purposes. Recursive Feature Elimination (RFE), which is a random forest based modeling approach [33], was used in the feature selection process by the authors. The authors identified that redundant data and attributes can create problems in model performance and thus, opted for efficient feature selection process like RFE. They used Caret Package [34] in R for programming purposes. The authors found that Bagging Models (RF and Bagged CART) had a higher performance than Boosting models (AdaBoost and GamBoost).

Since ensemble learning combines multiple ML models to improve the prediction accuracy and provides room to improve an individual model's weakness [35], it has attracted more researchers in leveraging this AI techniques in their applications. Sarkar *et al.* [36] have also used ensemble modeling for groundwater potential prediction in Teesta sub-catchment, Bangladesh. The authors used RF and Random Subspace models to estimate groundwater potential, with Area Under Curve (AUC) as the performance parameter. They also found that Distance to river, slope, curvature, elevation, LULC (land use land cover) and SPI (stream power index) were the determinant factors for estimating groundwater potential.

Natural subsystem also incorporates rainfall and runoff modeling in which models are used for runoff prediction and flood forecasting. Damavandi *et al.* [37] found that deep learning techniques likes LSTM provided better performance than physical models like CaMa-Flood

[38] for predicting short term daily streamflow using current day's streamflow and climatic data for a Texas watershed.

Surface water sources like streams, lakes, rivers are highly susceptible to pollution. They resources are of great value as a lot of needs such as irrigation, hydroelectricity generation, drinking water etc are dependent on them. Thus, surface water quality assessment becomes extremely important. Shah *et al.* [39] caputred 30 years of historical water quality data of Indus River, to model total dissolved solids (TDS) and specific conductivity (EC) though techniques like ANN, LR and gene expression programming (GEP). Their model approach is explainable and reproducible as they have provided the hyperparameter values for the models they have used. Their dataset contained various parameters like Calcium, magnesium, sodium, chloride, sulphate, pH, bicarbonate, and TDS and EC. All the other seven parameters were considered as input parameters to predict TDS and EC. The authors state that all the input parameters had high correlation with the output parameter. Their dataset records (about 360 data points), were split into 70:30 ratio for training and testing respectively. The performance of GEP was the most accurate followed by ANN. While this study brings out the model development process in detail, the authors could have provided more depth into data preprocessing and feature selection.

2.1.3 AI Applications in Socio-Economic Subsystem

Cost and demand prediction, risk analysis are some commonly use cases of AI applications across various industries [40] [41] [42]. Water system consists of similar applications in form of various socio economic operations like water demand forecasting, energy cost modeling, flood risk assessment etc. Water demand forecasting is critical for anticipating changes in demand and adjusting operations accordingly. It enables the stakeholders in to make better decisions regarding water allocation, pipe capacity, water usage, pumping operations, pricing, etc [43]. Energy cost modeling helps utilities identify opportunities to improve energy efficiency, reduce costs, and reduce their carbon footprint. Flood risk assessment is essential for managing the risks associated with flooding, which can cause significant damage to water systems and infrastructure. By incorporating these elements into their operations, water utilities and other stakeholders can ensure the availability of clean water, minimize energy costs, and mitigate the risks associated with flooding.

Time series forecasting has been a widely used category of AI applications and thus, many researchers have worked in this category while exploring water demand forecasting. Guo *et al.* [44] used Gated Recurrent Unit Network (GRUN) deep learning method to make short term water demand forecasts with a 15-min time step.They compared their deep learning based apprach with artificial neural network (ANN) model (ML based) and seasonal autoregressive integrated moving average (SARIMA) model (Statistics based). Due to the stochastic and non linear nature of data, the authors found that the deep learning technique outperformed the ML and statistical model. To decide the model's input parameters based on domain knowledge, the authors select historical water demand at 15 mins time stamps as the only input, since they are making short term predictions. Population, economic development, pricing policies can be considered for medium term and long term predictions. The authors use the data of two district metered area (which are part of water distribution network from northeast of Changzhou and southwest of Changzhou) in China. The authors also provide the model configuartions for ANN and GRUN, where they achieved optimal configurations through 3 and 7 dense layers for ANN and GRUN respectively. The authors do not provide any specific details about the data preprocessing techniques used in their experiment.

Deep learning techniques like LSTM (Long Short Term Memory) have also performed well when used for short term water demand prediction. Kuhnert *et al.* [45] used LSTM model to predict water demands for next 24 hours using hourly historical Water Consumption and calendar Information as the input parameter. The authors state that their work is a research experiment and how it can be brought into real production environment is a future scope.

Water demand prediction is a globally used application of AI over across the globe. Brentan *et al.* [46] applied Support Vector Regression (SVR) to predict hourly water demand (in litre/second) for a DMA in Franca, Brazil. To overcome the challenge of identifying non linear patterns in data, the authors used a hybrid SVR + Adaptive fourier series approach. The Fourier series layer coupled with SVR output was used to better adjust the minimum and maximum demand peaks and capture parts of the time series periodicity that SVR cannot reproduce.

Water utilities also focus on optimizing their operational costs. Plant managers are interested in investing towards energy management at wastewater treatment plants due to high operational costs and considerable saving potential. Wastewater treatment plants are estimated to contribute about 1% of a country's total energy consumption [47]. Torregrossa *et al.* [48] used a database containing data of 317 wastewater treatment plants for their experiment on calculating the relation between wastewater processes and corresponding energy consumption using Neural Network and Random Forest techniques. However, in this research, the authors mention did not mention some key details regarding their experiment, for instance, the split ratio of training and testing data, model configuration and data preprocessing methods used in their ML cycle.

Zhang *et al.* [47] also used a Random Forest model to predict unit electricity consumption by using the data obtained from 2015 Urban Drainage Yearbook of China. They performed typecasting of the data attributes to convert integer data type to float, non-numerical data (like strings) to object type, and converted all the default data to NaN (not a number). Since the RF model is good at dealing with outliers, they didn't work on eliminating the outliers.

2.2 Practice Review

2.2.1 Metropolitan Water Reclamation District (MWRD) of Greater Chicago

We had setup a working meeting with the MWRD Chiacgo to capture their AI practices at their wastewater treatment plant.

MWRD treats wastewater and provides stormwater management for residents and businesses in their service area which encompasses 882.1 square miles and includes Chicago and 128 suburban communities throughout Cook County. The MWRD serves approximately 12.72 million people each day, including 5.16 million residents. This case study focuses on the implementation of AI techniques for supporting wastewater operations, specifically controlling odors and corrosion in WRRF for MWRD Greater Chicago (GC). Yang *et al.* [19] developed an ML application to find the optimal dosage of sodium hypochlorite to supress the odor and corrosion cause by H2S in their wastewater treatment plant.

There were several challenges which motivated the authors to explore AI techniques for this task. Oxygen Reduction Potential (ORP) was used earlier for identifying the dosage of sodium hypochlorite. However, the use of this method led to more overdosing. Another way adopted by them was using Hydrogen Sulphide online sensors and trying to do feedback control dosing through it. However, in this case, sensor accuracy and reliability were a concern. Thus, MWRD Chicago switched to AI in order to determine the optimal dosage of sodium hypochlorite.

Data Collection, Storage and Management: To predict H2S and VFA (Volatile Fatty Acids) concentration, they collected online instrument data regarding flow, ORP (oxidation reduction potential), pH, wastewater temperature, tunnel pumping, and tunnel elevation from SCADA system and daily precipitation data. No new sensor installation was required specifically for this project. Their project used 17 years of influential historical data. For storing their data, they used district servers. They are internal servers only accessible to operating staff. Since district servers have limited access, they have created the copy of this database in the form of Enterprise Data Management System (EDS), so that this data can be accessible to engineering and research staff as well.

MWRD Chicago stores their CSV data into the RTDS (Real Time Decision System) in order to fit it into the AI model. Since data transfer to the RTDS comes with security concerns, their team of consultants (Xylem) used a Data Splitter at the sensors (of PH, ORP, and other parameters). For the data signal coming from the device/sensor, the data splitter splits the signal and passes it on to RTDS. While this setup costs overhead in terms of finances, it will save them from transferring data from district servers which might not be fully secured. Given the high cost of maintenance, they will consider not using Data Splitter and find a secure way to transfer data from district servers directly.

Data Preprocessing: For handling missing data points and class imbalance, they use techniques like interpolation and oversampling respectively. Detection of outliers in the data was achieved by domain knowledge. Min-Max Normalization data normalization was applied to all variables before feeding them to ML models. For the sparse and imbalanced datasets, they used random oversampling, synthetic minority oversampling, and adaptive synthetic sampling for data augmentation and and nonlinear data quantization (to remove skewness and make it a multi class problem). They split their data into the 80:20 ratio for training and testing purpose respectively.

Model Selection and Hyperparameter Tuning: Their model selection was driven by the type and amount of data they had. In order to select the data attributes to train an AI model, they used a correlation map. The features showing highest correlation with each other were used and fed to the model. Domain knowledge along with the correlation mapping was also considered to select the most relevant data attributes. Given the highly diverse datasets, three subproblems were formulated, and three cascaded ML modules were developed accordingly. In Module 1, influent wastewater characteristics were predicted using historical wastewater characteristics data. In Module 2, the predicted results from Module 1 together with online sensor data were used to predict H2S and VFAs (volatile fatty acids) levels. The predicted results from Module 2 were used as input for Module 3, which predicted the NaOCl dosage. A grid search with cross-validation was used to tune the hyperparameters of each model.

Details of their AI model experimentation:

- RNN: Tested with epoch values in the range 10 - 60, batch size values 16, 32 and 64, number of hidden neurons with values 16, 32, 64.

- Random Forest: Number of trees with values ranging from 10 to 100, maximum number of features considered for splitting a node with values 1 - 5, and number of levels in decision tree with values 4 - 32.

- SVM: soft margin constant with values 1, 10, 100, and 1000, the type of kernels with values “linear polynomial” and “RBF”; and kernel parameter gamma with values 0.01, 0.001, and 0.000

For predicting wastewater characteristics, the RNN using LSTM presented the best performance due to (i) the 17-year long-term data, (ii) the learning capability from the internal LSTM structure in RNN, and (iii) the perfect matching between the available data and the model capacity. The ARIMA performed the worst due to its limited learning capability from its linear statistical model nature.

Potential limitations of this experiment: The authors mentioned that they selected the highly correlated parameters as the input to the AI models. However, this may not be the best method as two highly correlated parameters may also represent the same attribute of data.

2.2.2 Clean Water Services, Oregon

This case study is based on Utilizing Soft Sensor System for Process Control and Optimization which was submitted as a solution to The LIFT Intelligent Water Systems Challenge by Clean Water Services [49].

Clean Water Services (CWS) is a special service district that provides wastewater treatment, stormwater management, stream restoration and water resources management services to more than 620,000 residents and business in urban Washington County, Oregon.

A multidisciplinary team was formed between Clean Water Services and Princeton University to facilitate effective management of nitrification at the facilities. An AI framework for soft sensor development driven by data based decision making was developed and applied in this project.

Lu *et al.* [49] states that the existing ML lacks a framework for water and wastewater operation. Thus, they develop a framework which is applied to the Rock Creek facility to predict influent wastewater flow and its receiving river flow. Clean water service benefits through this framework as it helps them to achieve reduced chemical and energy consumption.

Objectives of the project: (i) Predict influent flows for the Rock Creek facility one day (and three days) in advance. (ii) Predict the next calendar month average Tualatin River flow. (iii) Develop a simple, easy-to-use dashboard to visualize the predictions of future influent flow and monthly average river flow and process control tools.

Data Collection: They used 11 years (2010-2020) of historical data to develop soft sensors for predicting the next day's influent wastewater flow. They conducted preliminary data analyses to understand potential problems like temporal resolution (15-min, hourly, daily data), missing data points, etc. For prediction of the next month's average river flow, they used data from two USGS river monitoring sites (Farmington and West Linn) and 40 NOAA meteorological sites based on eight ZIP areas for 81 years (10/1939 – 12/2020). Average values of available data from the 40 sites were used as preprocessed daily data, and were later converted to monthly data for all meteorological variables. For wastewater flow prediction, 11 years of daily data were first split into training (2010-2017) and testing (2018-2020) datasets. For future river flow prediction, the 81 years of monthly data were divided into training (1939-1999) and testing (2000-2020) datasets.

Data Storage: The authors used SQL servers for data storage.

Data Preprocessing: The authors used data preprocessing techniques like feature scaling, feature selection, cross-validation. They integrated data from various sources through the Clean Water Services' Power BI visualization. Conventional standard deviation and box plot based detection can overestimate the number of outliers thus, they used a distance-based detection heuristic to identify outliers. Technical outliers were identified based on domain knowledge (e.g., flow > 0), and statistical outliers were detected based on a distance-based heuristic method.All variables (regressors) were scaled to the range of [0, 1] based on their corresponding maximum and minimum values.

Model Selection:

- Prediction of Next Day's influent wastewater flow: The authors found that ETR (Extra trees regression), KRR (Kernel ridge regression), and SVR (Support vector regression) achieved good results for this task.

- Prediction of Next Month's influent wastewater flow: KNN, ETR, and KRR were the three methods that generated best candidates from the 999 tested models.

Prediction performance of all models was evaluated based on four typical metrics, adjusted R2, mean absolute percentage error (MAPE), root mean squared error (RMSE), and mean absolute error (MAE).

Tools for Decision Support: Once the model accuracy was accepted by the utility team, the dashboard was developed with Power BI as a decision support tool.

2.2.3 Hampton Roads Sanitation District (HRSD), VA

HRSD has partnered with DC Water, MWRD Colorado, University of Michigan, Ann Arbor, Northwestern University, Oak Ridge National Laboratory to work on a project titled "Crossing the Finish Line: Integration of Data-Driven Process Control for Maximization of Energy and Resource Efficiency in Advanced Water Resource Recovery Facilities". This is an ongoing project and the team has outlined the vision they want to achieve using machine learning tools. They aim to treatment goals like reduce energy and chemical inputs, maximize energy and nutrient recovery. Through this project, they want to (i) achieve maximum energy and resource efficiency, (ii) achieve performance goals; and (iii) better manage the risk of implementing new technologies in the water sector. The team intends to share the benefits of this project across the water sector through the development and demonstration of a Toolbox. Much like a cookbook, the Toolbox will illustrate how generic statistical and ML tools can be applied, given plant-specific data synthesis and blending approaches, advanced data analysis techniques and the code associated with specific examples, hardware (e.g. control network servers, DCS/SCADA), and data management software configurations.

The team is currently in the planning and scheduling phase and thus, precise technical details are not available at the time of writing this report. However, the project goals of this team clearly describe how the utilities are intending to switch towards AI to maximize benefits and operational efficieny in the water sector.

2.3 Various Use Cases of AI in the Water Sector

The current state of AI in the water industry is helping the water sector decision makers to guide the utilities towards a digital transformation. Following is a list of use cases for which AI is playing a key role in changing the water industry:

- Process Automation: AI is used to automate and optimize various processes involved in water management, such as water treatment, pump flow operations, and wastewater treatment. It can help in automating routine tasks, improving operational efficiency, and reducing human error.

- Asset Management and Maintenance: AI is being used for several tasks such as determining effluent Nitrogen, odor and corrosion control, pipe defect classification, identifying potential leaks, reduce water loss, and optimize maintenance strategies to help decision makers better manage and maintain their water and processes at the utility.

- Retaining Domain Knowledge: Domain experts with years of experience play a key role in making decisions at the utilities. Integrating AI systems with their domain knowledge helps the utilities in retaining and capturing the knowledge which makes crucial decisions, even after the experts leave the organization.

- Energy Efficiency: AI techniques can be employed to optimize energy usage in water treatment and distribution processes. By analyzing data and patterns, AI algorithms can identify opportunities for energy conservation, recommend efficient operational strategies, and enable the integration of renewable energy sources.

- Customer Experience: AI-powered smart metering systems, chatbots for capturing customer complaints, and related customer analytics are examples of how AI is used to improve customer engagement and satisfaction.

- Enhancing Traditional Methods: AI-based systems are enhancing some existing traditional methods adopted in the water sector. For instance, AI is used to optimize irrigation practices by analyzing weather data, soil moisture levels, and plant characteristics. By considering these factors, AI algorithms can recommend optimal irrigation schedules, reduce water waste, and ensure efficient water usage in agricultural and landscaping applications.

- Demand Response and Demand Management: AI helps utilities manage water demand more effectively by analyzing historical usage patterns, customer behavior, and external factors. By understanding demand fluctuations, AI can assist in demand forecasting, load balancing, and implementing demand response programs to optimize water supply and reduce costs.

- Risk Assessment to avoid emergency events: Emergency events like pipe burst can prove to be costly to the utilities in terms of resource allocation and finances. AI enabled risk assessment strategies help the utilities to obtain various analysis like risk index of pipe failure, vulnerable areas in the city prone to pipe failure, susceptibility for sink holes, etc.

- Resource Optimization: AI techniques help optimize the use of water resources by analyzing data on supply-demand patterns, weather conditions, population growth, and other factors. This enables utilities to make informed decisions about water allocation, demand management, and conservation strategies.

- Water Quality Management: AI is employed to monitor and manage water quality parameters, including pH levels, dissolved oxygen, turbidity, and contaminant detection. AI algorithms can analyze data from various sources, such as sensors, satellite imagery, and historical records, to identify water quality trends, detect anomalies, and enable proactive measures for maintaining water safety.

- Decision Intelligence to Operator: AI-based decision support systems assist water utilities and decision-makers in making informed choices. These systems utilize AI algorithms to analyze complex data, assess risks, optimize resource allocation, and recommend suitable actions. It helps in capturing patterns like pipe age vs failure count, annual variation in number of failures, pipe bursts detected vs network length covered, etc., to let the operators make intelligent decisions.

- Predictive Analytics: AI techniques have been employed to analyze historical and real-time data, enabling predictive modeling for various water-related parameters. This includes predicting pipe failures, pumping rate, water demand, water quality, flow rates, erosion rates, and other critical factors. By identifying potential issues in advance, proactive measures can be taken to prevent or mitigate them.

- Water Network Optimization: AI techniques can be applied to optimize the design and operation of water distribution networks. By analyzing network topology, hydraulic models, and real-time data, AI algorithms can optimize pump scheduling, pressure management, leak detection, and pipe network layout to improve overall network efficiency.

2.4 Data Driven & Model Driven Models in the Water Sector

Model-driven models and data-driven models are two different approaches in machine learning and artificial intelligence.

Model-driven models are built based on prior knowledge and domain expertise. These models rely on explicit rules, equations, or algorithms to represent the relationships and behavior of the system being modeled. The model is designed based on an understanding of the underlying principles and mechanisms governing the problem domain. Model-driven models typically require less data to train and can provide interpretable and explainable results. However, they may struggle with complex and nonlinear patterns that are difficult to capture with predefined rules. Examples of model-driven models include decision trees, rule-based systems, and expert systems.

Data-driven models, on the other hand, are built by learning patterns and relationships directly from data. These models do not rely on predefined rules or equations but instead use algorithms to automatically extract patterns and make predictions or classifications. Data-driven models are trained on large datasets and can discover complex patterns and nonlinear relationships that may be difficult to capture with explicit rules. However, they may require a significant amount of labeled data for training and can be less interpretable compared to model-driven models. Examples of data-driven models include neural networks, support vector machines, and random forests.

Based on our extensive literature and practice review, we found that there exist a mix of data driven, model driven and hybrid models (which has both the elements) when applying AI in the water sector. Commonly used data driven models were ANN, LSTM, Genetic Algorithms, AutoML, SARIMA, RNN, Linear Regression, Logistic Regression, KNN, K-means, PHM based survival, with ANN being the most popular technique. Model driven applications included techniques like Fuzzy Logic, SVR, Adaptive Fourier Series, SVM, Random Forest, Multivariate adaptive regression splines (MARS), Ensemble Decision Trees, Bayesian model averaging (BMA), Naive Bayes, XGBoost, Regression Trees, Naive Bayes, AdaBoost, Markov Chain Model. Tree based algorithms were found to be the most commonly used model driven models. Adaptive neuro-fuzzy inference system (ANFIS), Boosted Regression Trees, Gradient-Boosted Tree, Dynamic Bayesian Network, Gradient Boosting Machine (GBM) were the Hybrid techniques found in literature, with ANFIS being the most popular. Some of the works explored both Model Driven and Data Driven techniques. Most of the publications found Data Driven techniques outperforming Model Driven techniques in terms of accuracy. In total, 46 publications explored Data Driven models, 39 explored Model Driven models and 14 explored hybrid techniques.

Since several literature references used Data Driven models (many of which also outperformed model driven models), indicates a high emphasis on utilizing available data to make predictions or classifications rather than relying solely on predefined models. Moreover, the prevalence of data driven models highlights the significance of data in the water sector. It indicates that practitioners in the field recognize the value of leveraging data to gain insights, make accurate predictions, and address water-related challenges effectively. The presence of hybrid models indicates that the water sector is increasingly recognizing the benefits of combining both model driven and data driven approaches in the water sector. Hybrid models can leverage the strengths of both approaches to enhance accuracy, robustness, and interpretability. This suggests a growing interest in developing more comprehensive and adaptive solutions that integrate prior knowledge and available data.

2.5 Limitations of AI in the Water Sector

The extensive literature review presented in the previous subsection detailed various applications of AI in the water sector and the AI techniques implemented for various usecases. We have reviewed more than 70 papers to capture the AI practices taking place over the years in the water sector.

There are several limitations observed in AI practices studied across literature and practice which remains unaddressed. One of the major concerns is reproducibility of models. We find that many researchers mention about the techniques and AI models which they used in their application but do not mention the model and experimentation configurations in details which provided them the optimal output. It makes reproducing the outcomes of a research difficult for other researchers in this domain. We also found that most of the research is inclined towards sequence prediction and regression tasks. However, this can be attributed to the fact that most of the data in water sector received from sensors and instruments is sequential data.

The scale at which these AI experiments are achieved is an another limitation. We saw that most of the papers were experimenting with a considerably small dataset (whereas, in practice review we saw the utilities working on years of historical data). The models developed in literature are academic models and can't be scaled to real world scenario, as they have been tested on very small dataset, which may not be representative of all real world scenarios. Models in real world settings are dense and complex are they are trained upon huge volume of data. The AI models proposed in literature were relatively simpler and thus, their implementation can't be translated to real world settings. The authors also did not explain any steps taken to overcome underfitting and overfitting scenarios faced while developing the AI solution. Generalization of AI models is very important and thus, providing details about what model settings work and do not work shall provide insights to other researchers.

Another drawback from a high level view is the absence of a standard dataset. Most studies acquired data from a particular region or utility, depending on their needs. If researchers can come up with globally accepted dataset (for example Iris, ImageNet datasets) that can act as a benchmark for research in AI for water sector, it can enable quicker development of state-of-the-art models. As the researchers use one-off datasets, it becomes difficult for researchers to build upon the existing models and take the state of the art methods to next stage [50].

Adding to that, many researchers also did not provide insights regarding their technical decisions. For instance, ANN was a commonly used algorithm in most of the research works. However, many publications lacked reasoning behind the choice of one ML technique over the other. Researchers did not discuss what characteristics of a given model makes it more desirable than others for their given use case.

Through practice review, we saw that the water utilities are collaborating with academic institutes and consultants as they don't have full in house expertise. Developing in house expertise is favorable for utilities as academia and consultant will not have domain knowledge as good as that of utility to effectively integrate it in the AI applications.

While the ML applications in water sector are centered around optimization and modeling, the majority of these efforts are constrained within academic boundaries and haven't made it to the decision making tools [2]. Moreover, while we were able to capture the utility practices by piloting with them, it is difficult to find non academic ML applications in water sector within research publications, which makes it difficult to capture the best practices followed in the industry. Real world AI applications in water sector are an intersection of civil engineering and computer science, which requires an interdisciplinary effort from domain experts of various fields. The experts from these domains need to come together to work collaboratively to ensure successfully application of AI in the water sector, which has not been emphasized by the academic publications.

2.6 Framework Validation and Verification

2.6.1 *ai*WATERS Framework Verification

To identify the appropriate building blocks for holistic AI implementation in the water sector, we started by carrying out a literature review of current AI practices in the water sector. We reviewed literature articles published in various conferences, journals, water research organizations and publicly available documentation of water utilities. To understand the best practices followed in the industry and real world, we conducted a practice review which comprised of publications and literature articles discussing AI practices in real world watershed and sewer shed system. We pilot interviewed 3 large utilities to find their AI applications and best practices. Based on the literature and practice review, we analyzed the progress and limitations of AI in the water sector. This analysis and review helped us in laying the foundations of the building blocks of AI. Later, the building blocks were iteratively revised based on feedback from Jacobs Engineering, a large consulting firm, providing technical and scientific services for government and private sector projects. Jacobs has years of real-world experience in guiding hundreds of water utilities. We also held an in-person two day brain storming session with Jacobs for the revision of building blocks. We also collected inputs from the Intelligent Water Infrastructure Network (iWIN) Committee, which comprises of technical domain experts from Oak Ridge National Lab (ORNL), Jacobs, Arcadis, DC Water, City of Houston, MWRD Chicago, Clean Water Services and Virginia Tech. Involving domain experts of water sector from various organizations helped us to verify that the building blocks proposed in this framework are complete and didn't miss out on capturing any other details.

To find current AI practices of large, medium and small utilities, we created a questionnaire. The questions were prepared based on the building blocks of aiWATERS. We did a dry run of the initial version of the questionnaire with two large utilities. We conducted a feedback meeting with them to capture their input on improving the questionnaire and making it easy to comprehend for experts in the water sector. The suggestions and approval of the large utilities verified our questionnaire to be effective for capturing AI practices at water utilities.

2.6.2 aiWATERS Framework Validation

We sent the questionnaire to 100+ utilities across United States, Canada and UK. We also pilot interviewed some of those utilities through video meetings to find out about their current AI practices and future plans. The questionnaire captured that some of the utilities are already applying AI in-house, and collectively follow similar steps as proposed in the aiWATERS framework. We captured the responses of all the participating utilities and have analysed them in the Results and Discussion sections. The findings obtained from their response helped us to validate our framework.

Chapter 3

*ai*WATERS: The Framework

Many utilities aim to pilot AI in their systems to augment their capabilities through scalable and learnable decision support systems. However, the criticality of the services provided by the water sector warrants caution to ensure that appropriate safeguards are in place to mitigate any potential risks associated with AI.

There are various challenges and struggles can be associated with deploying AI in a system. A survey conducted by International Data Corporation (IDC) [51] revealed that more than 50% of 2000 technology buyers struggle with various factors like lack of skilled personnel, lack of ML operations tools and techniques, lack of adequate volume and quality of data, and trust and governance issues associated with AI. Considering various challenges found in literature review and reports, we aim to provide a framework called *ai*WATERS: AI for water sector, that provides the foundation, necessary steps, and path that utilities can follow to effectively apply AI into their systems.

3.1 Pillars of *ai*WATERS

We provide a framework to help utilities successfully navigate the issues associated with the application of AI in the water sector. The framework considers various issues associated with AI applications and is built on seven pillars of AI: Understanding Benefits and challenges, Application Goals, Data Readiness, Knowledge Integration, Model Development, Decision

Support, and Implementation, as shown in Figure 3.1.

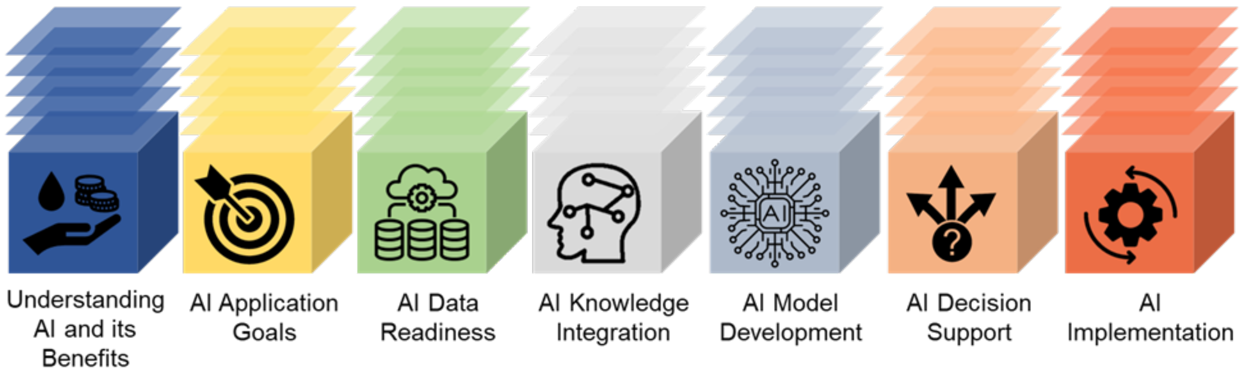

Figure 3.1: Seven Pillars of *ai*WATERS.

This framework contains recommendations based on best practices identified during an extensive review of academic literature of AI in the water sector and other domains, discussions with real-world practitioners from water utilities, and a review of industry best practices regarding AI applications.

The seven pillars presented in this framework also explore the major research domains of AI across industries. The whole life cycle of AI can be envisioned into 3 parts; the left part which involves what can be considered an input to the AI model - Data and knowledge; the middle part which process and computes the input, i.e. AI model development; and the right part which is the output/decision of the AI model. We define these as AI Left Domain Pillars, AI Middle Domain Pillar and AI right domain pillar as shown in Figure 3.2.

Before carrying out data collection and model development to make AI driven decisions, we encourage the utilities to understand the AI technology, its benefits and challenges and have defined goals for their application. We categorize understating AI and defining application goals as the AI Top domain pillars. The complete structure of *ai*WATERS Framework is shown in Figure 3.2.

Based on our comprehensive literature review, it is evident that a significant amount of

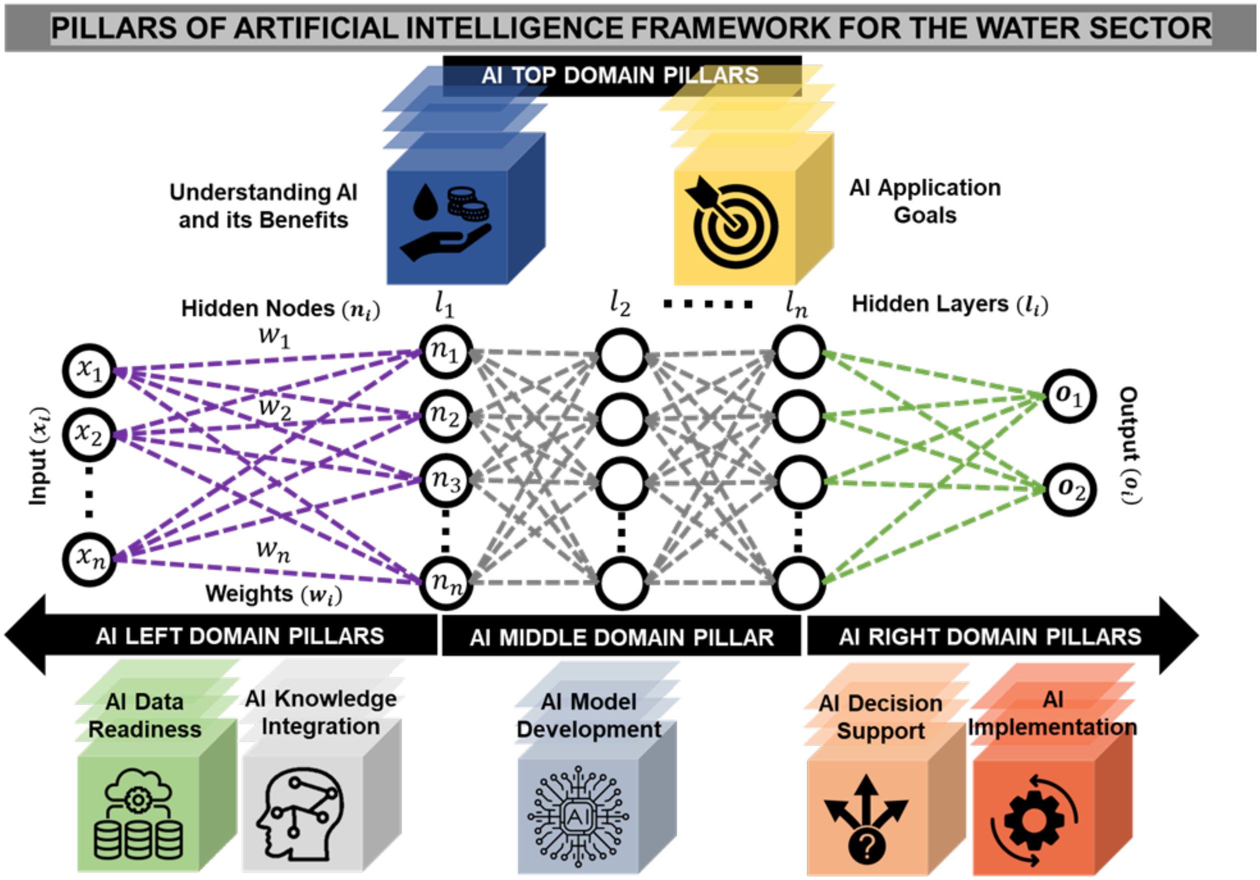

Figure 3.2: Architecture of *ai*WATERS.

attention has been devoted to the middle domain of AI, which involves the development of advanced algorithms and their optimization for publication and real-world applications. However, it is equally essential to address the left and right domains of AI, namely the data and knowledge domain and the decision-making domain, respectively. To tackle the issues related to data input, researchers are exploring cutting-edge technologies such as synthetic data generation and knowledge-guided machine learning approaches. Ensuring the trustworthiness of AI output has been a contentious issue for several years, as incorrect output generated by AI can pose significant risks to society. Therefore, researchers are actively developing frameworks to ensure the reliability and trustworthiness of AI output.

Therefore, the suggestions presented this section will also aim to discuss recent research advancement across AI left, middle and right domain pillars. Moreover, It is essential to recognize that each water utility function is unique, with its own specific challenges and issues that must be addressed to ensure sustainable and resilient service delivery. This framework is designed to offer a generalized perspective that can be applied across different types of utilities and benefit all stakeholders.

Following is a brief about the seven pillars of this framework:

1. **AI Understanding, Benefits, and Challenges:** The first pillar on understanding AI and its benefits focuses on briefly explaining the different foundational concepts of AI, difference between AI in academia and industry, benefits and challenges.

2. **AI Application Goals:** The second pillar describes how AI should be applied and discusses topics like the application at different levels of the studied system, how to deploy AI in various categories, and the different modes of building AI

3. **AI Data Readiness:** The third pillar describes how to evaluate the quality and quantity of collected data and how to preprocess data to make it ready for AI

4. **AI Knowledge Integration:** The fourth pillar explains how the knowledge in the minds of the utility experts can be integrated into AI modeling frameworks to build more robust models

5. **AI Model Development:** The fifth pillar discusses how to develop accurate and reliable AI models.

6. **AI Decision Support:** The sixth pillar explains methods to improve the trustworthiness of AI models and how to develop "human-in-the-loop" models.

7. **AI Implementation:** The seventh pillar on AI implementation explains methods for ensuring successful AI application in the real-world and continuous improvement of models.

The following subsections will discuss each of the seven pillars in detail.

3.2 *ai*WATERS Understanding and Benefits

Ever since the boom of AI technology, terms like AI, machine learning and deep learning have become buzzwords and it raises questions about whether these studies attempt to exploit the keywords, using the term 'deep learning' to take advantage of the current scientific zeitgeist.

Therefore, it essential for the water utilities to understand AI as a technology and its corresponding benefits and challenges. Artificial Intelligence (AI) is the broader field of computer science that deals with the development of intelligent machines that can perform tasks that typically require human intelligence. AI encompasses several subfields, including Machine Learning (ML) and Deep Learning (DL). Machine Learning (ML) refers to the process of enabling machines to learn from data without being explicitly programmed. Deep Learning (DL) is a specific subset of ML that is based on artificial neural networks that are designed to mimic the structure and function of the human brain. DL is particularly suited for processing complex data types, such as images and speech, and has been used in many applications, including computer vision, speech recognition, and natural language processing [52].

Understanding the differences between AI, ML, and DL is important for the water industry because it allows water professionals to identify the appropriate techniques and tools for solving different types of problems. By understanding the differences between AI, ML, and DL, water professionals can determine which approach is most suitable for a given task, taking into account factors such as the nature of the data, the desired outcome, and the available resources.

Table 3.1: AI in Academia vs Industry

Key Factors	**Research/Academia**	**Industry**
Requirements	State of the art model performance (which is optimal as recommended by existing literature) on benchmark (standard) dataset. Major focus is to meet publication standards of conference and journals	Different stakeholders have different requirements. Models should be trustworthy as licensed professionals are responsible for the model decisions. It should adhere to the legal requirements
Computational Priority	Fast Training, high throughput	Fast inference, low latency
Data	Limited (Clean, well formatted and limited data)	Large datasets, which can be noisy and unstructured
Fairness	Often not a focus, as model is not yet used for real world applications. The major focus is often the accuracy and output of the model.	Must be considered, since the model must be used for real world applications
AI Interpretability	Often not focused much. ML models are built using programming libraries and there is lack of reasoning behind using a particular model, its structure, parameters and performance compared to other techniques, as the major goal is to generate low error output.	Must be considered. Adequate experimentation and knowledge of the ML models is necessary to justify the decisions taken
Continuous Model Learning	Not a major focus. The models may not be used in real world after publication and thus addressing how to keep the model updated over time is not of much interest to the researchers	Real world settings are dynamic with changing data and knowledge which drives the model. Timely updates to model is required

AI implementation in the industry, especially water sector, is in its early stages. There are many academic literature references available that can be referred by the utilities to gain deeper insights into how AI can be implemented in the water sector. However, utilities need to aware about various differences that exist in AI implementation within academia and industry [53], as shown in Table 3.1.

The current state of Artificial Intelligence is not comparable to super intelligence [54] and thus, stakeholders from the water sector should be aware of it's limitations like data acquisition (data quality and availability are necessities for AI capabilities), limited knowledge (water sector needs AI domain experts to integrate AI in their system), computing power and resources (like GPUs, funds etc), blackbox nature AI models are often not interpretable), legal concerns (accountability of AI decisions). At the same time, AI technology also offers various benefits like includes increased efficiency and productivity, improved accuracy and precision, easily identifies trends and patterns, continuous learning and improvement, drives down the performance time and enables multi-tasking and eases workload, can handle multi-dimensional and multi-variety data, and supports decision intelligence.

Owing to certain limitations of AI, the water utilities should adopt a human in the loop approach for their solutions and should not treat AI as an autonomous technology.

Therefore, the stakeholders and decision makers from the water sector should be well versed with the AI technology, it's benefits and challenges when integrating it in their system for real world water applications.

3.3 *ai*WATERS Application Goals

AI applications currently existing within the utilities are minimal and limited to a specific and isolated to process or infrastructure. The System of Systems approach encourages water utilities to move towards expanding their internal AI applications and achieve a fully AI-driven functionality for all processes taking place in the utility at a system level [55].

The water system is a highly complex and interconnected system comprising of several subsystems, including the natural, built, and social subsystems, as illustrated in Figure 3.3. These subsystems consist of numerous elements and processes, and their interrelationships contribute to the complexity of the entire system. For instance, the built subsystem contains components such as pipes and wastewater treatment plants, which involve complex processes such as pipe failure prediction and wastewater treatment. On the other hand, the natural subsystem includes components such as groundwater and river water, which have numerous associated processes like surface water quality assessment and nitrogen content prediction. Additionally, the social subsystem involves processes such as energy cost modeling and water demand forecasting. Therefore, understanding the complexities and interdependencies of the various subsystems is crucial for effective management of the water system.

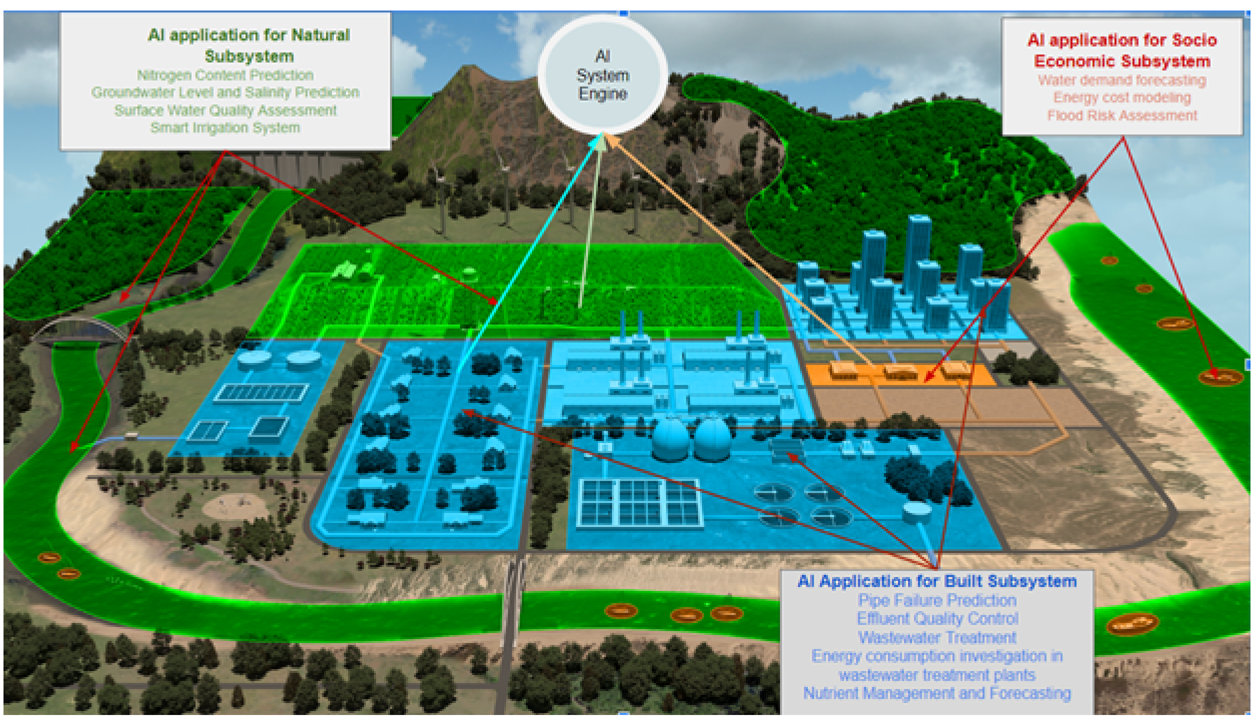

Figure 3.3: Water System of Systems

3.3.1 Using AI at Different Levels of System

Water utilities worldwide are at different stages when it comes to implementing AI within their system. The existing use of AI in water system is focused on either a specific process (for example, reducing nitrogen content from water) or an infrastructure component (for example, focused particularly on failure of water pipe infrastructure). Therefore, the current presence of AI within the utilities is limited and isolated. Figure 3.3 captures the various processes and infrastructure components at which AI/ML is currently applied within various subsystems of the water system. The implementation of AI within a utility can be categorized hierarchically as follows (as shown in Figure 3.4):

1. No AI application
2. Component and problem specific application
3. Sub-system level application
4. System level application

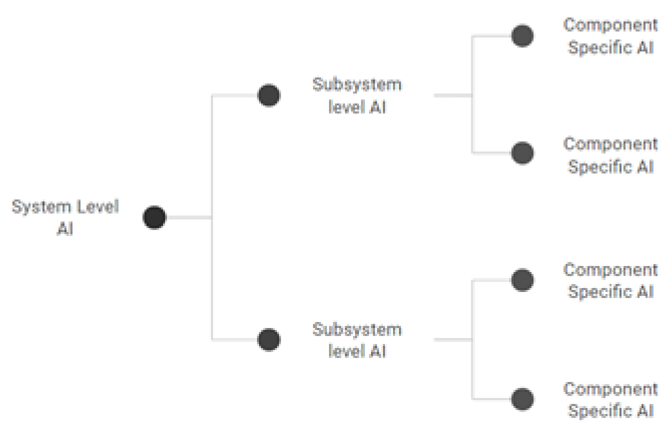

Figure 3.4: Different levels of AI Applications

In the current water utility landscape, some utilities still adhere to conventional practices where most of their operations are carried out manually. Consequently, such utilities opt for a No AI implementation approach.

Moving one step above are the utilities who apply AI techniques on specific components or problem within their system, such as utilizing AI to predict water pipe failure in their built sub-system or forecasting monthly energy costs. However, these utilities still rely on manual efforts to govern the rest of their processes and assets.

In the sub-system level application, AI is deployed at various components of the sub-system like built, natural, or socioeconomic sub-system. Here the aim is to use AI for multiple applications. Further, this implementation may also be extended to achieve a hybrid implementation so that isolated AI applications can interact with each other and the output from one application can be used to benefit another AI process within the system.

The highest level in the AI implementation hierarchy is the system-level application, where AI governs all processes in the system. In this case, AI operates all infrastructure and assets of the sub-systems, while a human is present to monitor the flow of operations. If AI is deployed at the system level, it can learn the system behavior, which can improve the operators' decision-making process. Deploying AI across the system can bring several benefits to utilities, including efficient and optimal energy usage of devices, reduced manual effort, increased automation and advanced technology, and improved maintenance of infrastructure components and elements.

3.3.2 Various Categories of AI Deployment

From our understanding, AI can be deployed in various categories within the system, which can be defined as per the degree of automation associated with the AI output. The three categories are listed as follows:

- **Manual Output:** Calculations from model used for manual actions and decision making. No automation involved.
- **Semi Automated Output:** Functioning of the system components is partially operated by the AI output. For example, a pump operator controlling the flow of water through pipes based on real time water demand prediction.
- **Fully Automated Output:** When human is observing the the process and all components are automated and operated by AI.

Each of these categories involves a different level of human interaction and engagement. In manual output, a high degree of manual interaction is involved as the human reads the AI output and accordingly carries out further processes that involve manual efforts. No automation is associated with the results given by AI. In semi-automated output, the machines are operated and controlled by the AI-based output to some extent. Lastly, a fully automated, AI-based deployment is envisioned through which the human observes the functioning of the machines that will be fully governed by AI-based decisions and outputs. The role of the human in this case would be to take an action in case the AI-based decisions operating the machines seems to be off track and doubtful.

3.3.3 Different Modes of Building AI

Three majors modes of building AI applications within the utility are:

- **In-house AI applications**: The AI applications and models are developed within the utility.

- **Vendor Based AI applications**: Utilities collaboarte with vendors and consultants to integrate AI in their system.

- **Hybrid AI Application:** Utilities build in-house AI applications with the guidance of vendors and consultants.

Utilities can build in-house models or can use the vendor-based model. Each option has its own set of advantages and disadvantages. In-house models are developed by water domain experts, which can be integrated with domain knowledge, and thus can provide effective learning for the given problem. In addition, in-house models can incorporate utility-specific scenarios and conditions accurately. However, building in-house models requires a significant investment of time and resources, as they also need an in-house AI domain expert who can collaboratively work with water domain experts. Vendor-based models may have advanced and latest techniques that can be used as they are created by AI professionals. However, they may not be able to integrate the utility-specific scenarios and conditions as accurately as the in-house models. Utilities can also follow a hybrid approach where they build in-house models under the guidance of collaborating consultants and AI vendors.

3.4 *ai*WATERS Data Readiness

"Garbage In, Garbage Out" is a popular saying in the field of computing that emphasizes the importance of high-quality input data. The quality of output from a machine learning or AI system depends on the quality of data fed into the system. If the input data is flawed or biased, the output produced by the system will also be inaccurate and biased [56]. As the water sector begins to transform itself through AI and Digital Transformation, creating this "AI-ready" data will be a key determinant of success. A lot of research is focused on improving state of the art AI models, however, creating AI ready data is equally important [57]. A data-enrichment platform for data scientists, CrowdFlower, conducted a survey of 80 data scientists and found that data processing and cleaning accounts for almost 60% time of their daily tasks as shown in Figure 3.5 [58]. Therefore, data readiness is a challenging aspect of AI applications across various sectors. Utilities should be well versed with the tasks involved in cleaning, organizing, and processing data to optimize the huge investment of time in making the data AI ready.

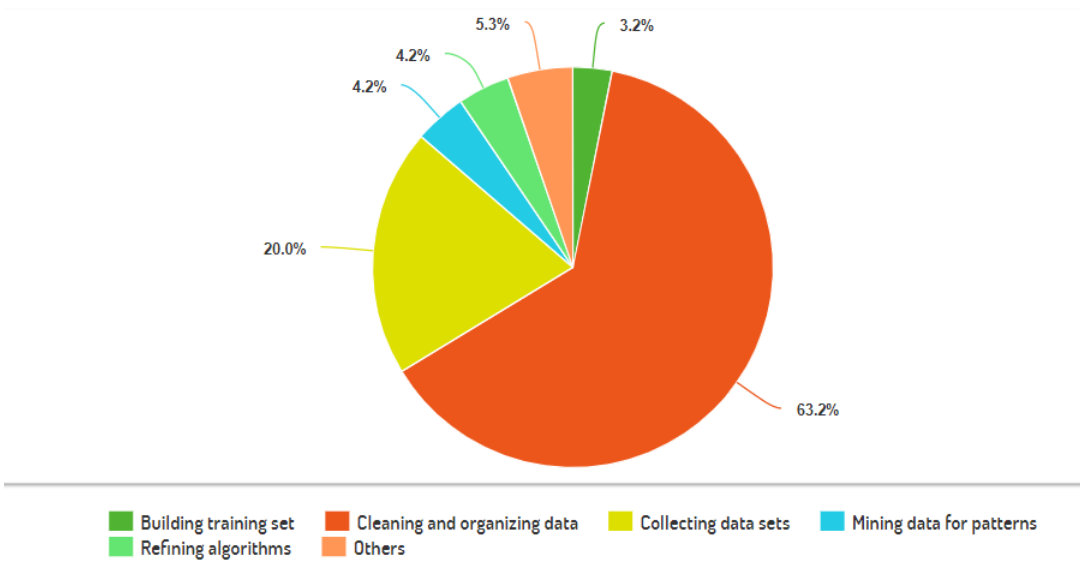

Figure 3.5: Time consumption for various tasks in building AI models

3.4.1 Evaluating Quality and Quantity of Collected Data

The quality of collected data can be determined by several factors such as data source, data integrity, data timeliness, and data relevance [59].

Data Source: Quality of data can highly correlate to the source from which data are collected. Utilities should evaluate the reliability of the data before using it for AI applications. Data reliability can be ranked based on the source from which data have been collected.

Utilities can rank the reliability of sources in the following order (highly reliable to least reliable): ground sampling and LIMS (Laboratory Information Management System) > utility database (SCADA, EAM, GIS) > derived indirectly (static pressure calculation, C Factor estimation) > educated guess (high confidence, moderate confidence, low confidence) as shown in Figure 3.6

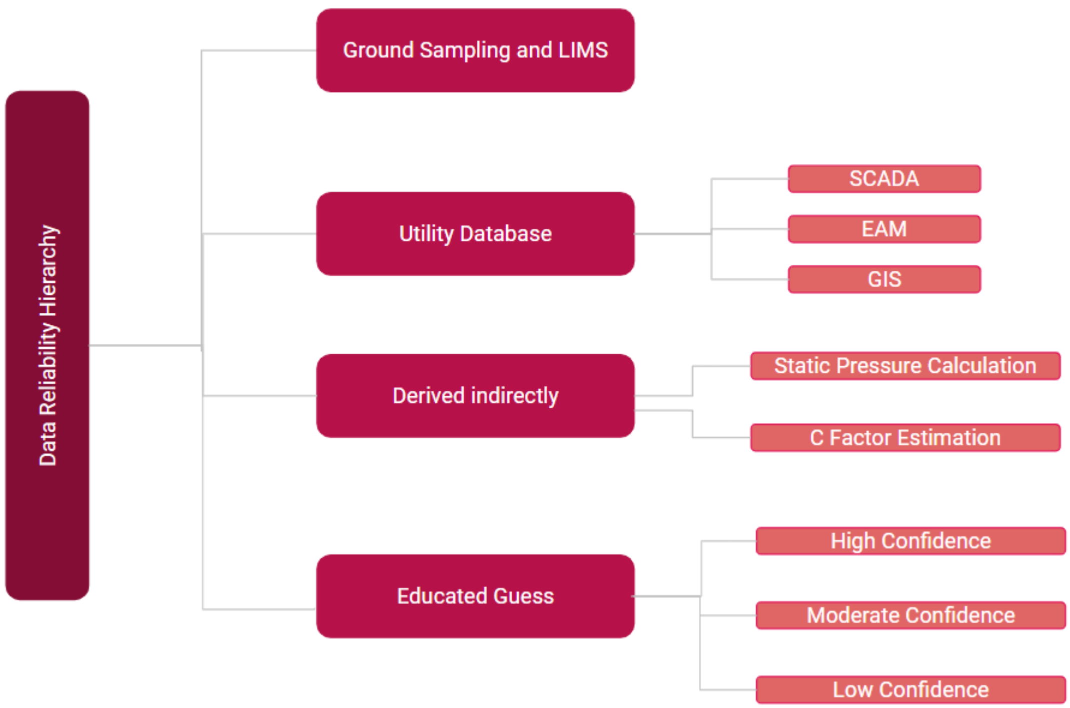

Figure 3.6: Hierarchy of Data Source Reliability

Data Integrity: Data integrity is the accuracy, completeness, and consistency of the data as maintained over the time and across formats. For example, a water utility may collect data on water quality parameters such as pH, temperature, and dissolved oxygen levels to ensure the water is safe for consumption. Questions to determine data integrity in this case could include: Are the instruments used to measure these parameters calibrated and accurate?, Are the sampling techniques appropriate for capturing the spatial and temporal variations in water quality?, and Is the data consistent across different laboratories and measurement methods?

Data timeliness: Data timeliness refers to how up to date the collected data are. For instance, a water utility may collect data on water demand and supply to optimize the operation of its water treatment plant. If the data are not up to date, the utility may end up over- or under-producing water, which can result in wastage or shortages. Therefore, the utility should ensure that the data being worked on are the latest version available to make appropriate interpretations.

Data Relevance: Data relevance is also critical for ensuring that the analysis is valuable and actionable. For example, a water utility may collect data on customer complaints about water quality to identify the most pressing issues and prioritize infrastructure investments. Gathering irrelevant information, such as customer complaints about billing, would be a waste of time and resources. Therefore, the utility must ensure that the data being collected are relevant to the problem statement being considered.

There is no fixed rule for how much data is enough for ML applications, as it depends on the specific problem domain, data sources, and ML algorithms used. The 5V's of big data (Volume, Variety, Velocity, Veracity and Value) are considered to be a relevant parameter when determining the the quantity of data required for an AI or ML application [60] [61] as shown in Figure 3.7.

Following is a description of the 5V's of big data for the water sector:

Volume: Volume refers to the size of the collected data. Having enough samples, features, and amount of data will provide the ML model with better scope to learn the patterns in the data.

Velocity: Velocity refers to the rate at which data are collected at the utility. Real-time monitoring systems can provide high-velocity data that can help utilities respond quickly to

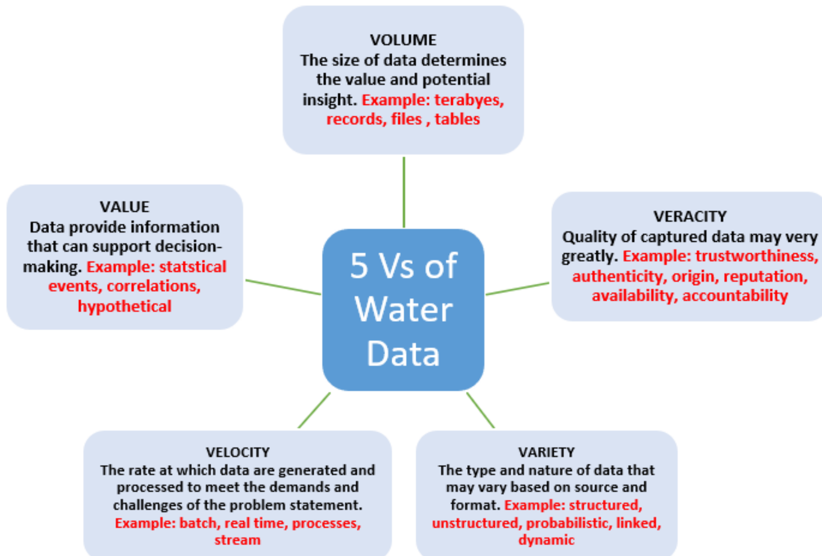

Figure 3.7: 5Vs of Big Data

changing conditions.

Variety: Variety refers to different forms of data, i.e., structured data (e.g., customer billing information), semi-structured data (e.g., maintenance logs), and unstructured data (images of pipe defects). Not all of the collected data will be in the structured/tabular form. Collecting a variety of data can help utilities gain a more comprehensive view of their operations and identify patterns that may not be apparent in structured data alone.

Veracity: Veracity is the inconsistency and uncertainty observed in the data. Big Data can be variable and highly dimensional. For instance, data in bulk could create confusion whereas less data could convey incomplete information. Water data can be subject to various sources of uncertainty, such as measurement errors or missing data. Ensuring the veracity of the data is important for avoiding biased or incorrect conclusions.

Value: Value refers to the usefulness of the collected data in generating insights or supporting decision-making. For example, water quality data can be used to optimize treatment processes and ensure compliance with regulations. Collecting and analyzing data that generates actionable insights can help utilities improve their operations and better serve their customers.

3.4.2 Data Preprocessing

We have listed some of the commonly used data preprocessing techniques in this subsection which are applied across a wide range of AI applications. It should be noted that utilities may not require to follow each of these steps in the list depending on the their tasks and data.

Data extraction, compilation and typecasting: Extract data from relevant sources and check if any of the attributes requires a type conversion. For example, image classification involves matrix multiplication, which cannot be performed with string data types. In this case, columns which represent numeric values but are in the string data type should be converted to integer/long/float data type format.

Categorical data transformation: Categorical data transformation is the process of converting categorical variables (i.e., variables that take on a limited number of values, such as colors, categories, or labels) into numerical values that can be used in data analysis. Categorical data transformation is required because many machine learning algorithms can only work with numerical data, and cannot directly process categorical variables. There are various techniques for categorical data transformation, such as one-hot encoding, label encoding, and target encoding. For instance, in the process of predicting water quality, the categorical variables might include different types of water sources (e.g., river, lake, well), different water treatment technologies (e.g., reverse osmosis, chlorination), and different water quality standards (e.g., drinking water, irrigation water). By converting these categorical variables into numerical values, machine learning algorithms could be trained to predict water quality based on the combination of these variables.

Handling missing values and outliers: Datasets may contain missing values, incomplete feature information, or missing feature columns. Missing values can be handled by eliminating data objects, estimating missing values, or ignoring missing values. Outliers are data points that significantly differ from the rest of the observations in the dataset. Outliers can also be useful and may form a new class for the dataset. Domain knowledge should be used to identify which data points are outliers. The data points can also be compared with statistics like mean or median to determine if they deviate significantly from the behavior of the majority of data.

Flag inconsistent and duplicate values: Inconsistent values are those that do not truly represent the dataset to which they belong, while duplicate values add redundancy to the dataset. These values should be identified and addressed through expert domain knowledge. For example, consider a dataset that contains measurements of water quality parameters, such as pH, temperature, dissolved oxygen, and conductivity, at different monitoring stations in a river. However, upon inspecting the dataset, we notice that there are some inconsistent and duplicate values in the dataset that need to be removed. Inconsistent values might include data that fall outside of a certain range or have extreme values that are not possible for the given parameter. For example, a pH value of 15 or a temperature of -50°C might be considered inconsistent values. These values could be removed from the dataset to ensure the accuracy of the remaining data.

Dimensionality Reduction: Complex datasets may have many features that make processing the model difficult. Dimensionality reduction reduces the dimensions of the dataset, making the newly transformed data relatively simpler. This helps avoid overfitting, as more complex data may lead to overfitting. Principal component analysis (PCA) is one of the commonly used technique for dimensionality reduction.

Discretization and Binarization: Discretization and binarization are both techniques used in data preprocessing to transform continuous data into categorical data, which is often easier to analyze using certain machine learning algorithms. Discretization involves dividing continuous data into discrete intervals or categories. This can be done using various techniques, such as equal width binning (dividing the data into intervals of equal width), equal frequency binning (dividing the data into intervals with equal number of data points), or clustering (dividing the data based on cluster analysis). Discretization can help to simplify the data and reduce the noise in the data. Binarization, on the other hand, involves converting continuous data into binary data, which takes on only two values (usually 0 and 1). This can be done by setting a threshold value and assigning a value of 1 to all data points that are above the threshold and a value of 0 to all data points that are below the threshold. Binarization can be useful in cases where only the presence or absence of a certain feature is important.

Data Normalization: Data normalization is the process of scaling numeric data to a standardized range of values, typically between 0 and 1. The purpose of data normalization is to ensure that all numeric variables are on a similar scale and have a similar influence on the analysis. Data normalization can be important in many machine learning algorithms, as some algorithms may be sensitive to differences in the scales of input features. For example, in distance-based algorithms such as k-nearest neighbors or clustering, features with larger scales may dominate the distance calculations, making the contribution of other features less important. Similarly, in gradient descent-based algorithms such as linear regression or artificial neural networks, large differences in the scales of input features may make it difficult to converge to a minimum loss.

Data Distribution: To ensure the quality and accuracy of data used to train AI models, utilities must consider the type of data distribution and skewness required for their problem statements. For example, predicting the cost of data collection instruments for yearly budgeting requires different levels of accuracy depending on the cost of the instrument. An error of 10% in predicting a cheap instrument may be acceptable, but the same error for an expensive instrument could have a significant impact on operations budget. Therefore, it is crucial to check the data distribution and sample the relevant and important data points. Thus, data should be finely distributed and have required samples as per the use case. Noisy features, such as erroneous labels or input attribute errors, should also be considered. Model developers must be aware of these factors and take measures to ensure data quality.

Feature Representation: To ensure accurate output for a problem statement, it is crucial to have all relevant features present in the dataset with sufficient sample size to learn the patterns within the data. However, in some cases, the actual data may not fully represent all important features. For instance, when predicting pipe failure, cracks can be a crucial feature, but different types of cracks such as fatigue, circumferential, and longitudinal may exist. Having enough data samples of each type of crack is essential for the model to learn and analyze all crack patterns present in a pipe.

Based on our discussion with consultants, we found that Data Extraction and Visualization are widely used steps by the utilities. We divide various preprocessing steps shown in Figure 3.8 in three columns, the left, middle and right columns depict the steps that are preliminary, intermediate and advanced in terms of complexity.

3.4.3 Data Augmentation and Class Imbalance

Data augmentation is a technique commonly used in machine learning and deep learning to artificially expand the size and diversity of a training dataset. It achieves this by generating new modified versions of the original data. The primary aim of data augmentation is to

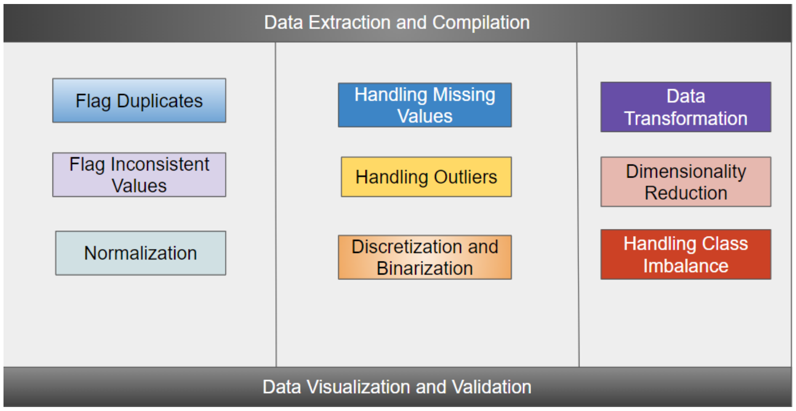

Figure 3.8: Data Preprocessing Steps

produce new training examples that retain the characteristics of existing ones while also incorporating enough variation to capture a broad range of real-world scenarios that the model may encounter.

Imbalanced data samples for a class can reduce data representativeness and introduce bias in AI models trained on such data. In such cases, developers should use class imbalance techniques to address this issue. The sensitivity of class imbalance varies with the complexity of the problem, with multiclass problems being more challenging than binary classification. Although a good model should learn to address class imbalance, deploying such a model can be challenging. Therefore, it is important to rely on class imbalance techniques to ensure model performance.

Some of the popularly used techniques to handle class imbalance are these:

- **Data level methods (resampling):** Under sampling (reducing samples of majority class) and oversampling (creating copies of minority class) are some of the popular and conventional sampling techniques.

- **Synthetic Data:** Synthetic data refers to artificially manufactured information, generated algorithmically instead of being derived from real-world events. Although synthetic data is commonly used to handle class imbalance, it is also a rapidly developing field of research in data analytics. Synthetic data offers several benefits, including the ability to reduce constraints when using sensitive or regulated data, customize data to specific conditions that may not be attainable with authentic data, and generate datasets for software testing and quality assurance. There are several methods to generate synthetic data, such as variational autoencoders, GAN models, and neural radiance fields.

Utilities can benefits in several ways by the generation of synthetic Data as it has various organizational and technical benefits, as described below:

Machine Learning Usecases:

- **Evaluation and Comparison of ML Algorithms:** As found in the literature, advanced algorithms are constantly being tested in the water industry to improve performance compared to existing approaches. However, it is important to evaluate and compare the performance of new algorithms against previous ones on realistic datasets. Synthetic data provides the advantage of generating unlimited sample sizes to explore the impact of sample size on model performance. This is particularly useful when the total amount of available data is limited.
- **Preventing Privacy Attacks on ML Models:** Utilities often handles sensitive customer data and privacy attacks pose a data security risk to ML models by exposing the training data. Training models using synthetic data can eliminate concerns about protecting the training data and minimize privacy risks. This approach enables the sharing of models within the scientific community to optimize utility and learning.

Software Testing: Utilities may use various softwares for controlling and checking the flow of water operations. Real-world data of these applications may not always create a particular flow of action on a software, and thus, synthetic data may help test the application with all edge cases.

Education and Training: Synthetic data can be used as an effective training and onboarding tool, particularly for employees who need to understand how to handle personal data. They can also use synthetic data in their hiring process, where they want the candidate to solve a problem based on utility's data characteristics without using the actual data.

Data Retention: To reduce costs associated with managing resources, infrastructure, and processes, industrial firms retain customer data for a limited time. However, some data are important and could be reused if privacy concerns are addressed. Synthetic data techniques can be used to retain data without privacy concerns, but storage costs and data relevance over time must be considered. Utilities should make informed decisions about which data to retain.

Data Sharing with Vendors: Some utilities collaborate with vendors and third-party consultants to implement AI in their operations. While collaboration can help meet business needs, sharing sensitive data with outside entities poses a risk to the organization. Synthetic data generation provides an alternative that mitigates privacy concerns. Furthermore, sharing personal data with vendors involves legal contracts and paperwork, which can be time-consuming and resource-intensive. Using synthetic data avoids these challenges, reducing overhead costs.

3.5 *ai*WATERS Knowledge Integration

Along with data, integrating domain knowledge is equally important to make effective AI applications. Science-based models, based on scientific theory, are used alongside ML models for modeling problem statements. While ML models heavily depend on data and are a black box, science-based models come with the limitation of having a large number of parameters/states and incomplete or missing physics/process knowledge. Scientific models are useful for discovering knowledge and predictive analysis. However, inferring parameters and state variables is a cumbersome process, and they are not as agile as ML models. On the other hand, ML models work well when they have a large set of representative samples but may show poor generalization performance with limited data and are not rooted to scientific knowledge.

Another problem with ML models is that they are not designed to discover new knowledge. They may be good at making predictions but, in many scientific problems, the goal is not to stop at prediction but to understand the phenomenon resulting in those predictions. Explainability or interpretability is a major goal in science. Through knowledge-guided ML, which makes use of both knowledge and data as displayed in Figure 3.9, patterns from data can be extracted without ignoring the knowledge accumulated in scientific research. Knowledge can also be induced in the ML models tuning how they make decisions at each stage of processing (for example, defining the rules for a fuzzy logic model detecting pipe failure timeline).

To effectively integrate domain knowledge in AI applications for water utilities, collaboration is required among experts from various fields, including data scientists, water experts, environment experts, and legal experts. For example, when modeling wastewater treatment to predict effluent release, environment experts must determine whether the ML model's predictions will impact natural bodies such as rivers, streams, and lakes. Similarly, legal experts must assess the legality and adherence to guidelines of the AI-based decision-making envisioned by the utility. Figure 3.10 illustrates the need for knowledge integration across multiple domains in AI for water utilities.

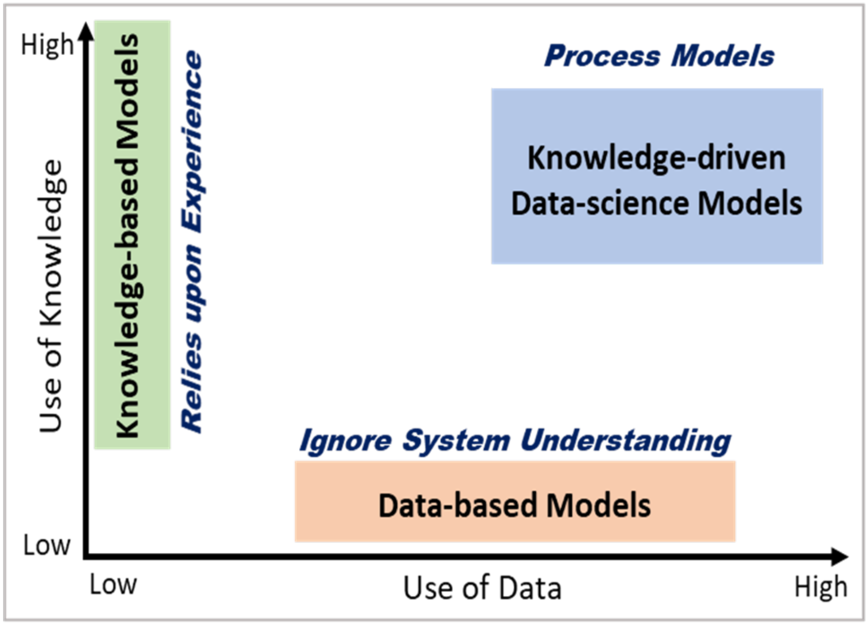

Figure 3.9: Knowledge and Data driven Models

3.5.1 Extracting Domain Knowledge for AI Applications

Domain knowledge can essentially be obtained from people, data, and experiments as described below:

Extracting Knowledge from Domain Experts: Experts with decades of experience in the utilities industry possess practical knowledge of the process and factors that influence it. Survey forms and interviews with domain experts can also provide valuable information for studying the problem statement. In addition, experts can provide structural knowledge, which is the basic problem-solving knowledge that describes the relationship between concepts and objects associated with the problem statement. Such knowledge, known as heuristic knowledge [62], is derived from domain and subject matter experts. Conferences, seminars, workshops, and industry events are also excellent resources where domain experts can be found discussing various challenges, opportunities, and research trends in the water sector.

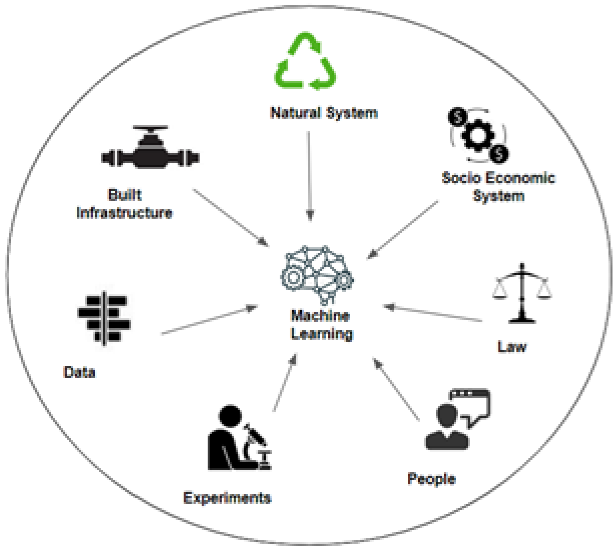

Figure 3.10: Domain Knowledge Integration in AI

Extracting Knowledge from Data: Extracting domain knowledge from data can provide valuable insights and understanding for solving complex problems in the water sector. Data sources include not only sensor data and historical records but also documents, metadata, and prior decisions and results. For instance, documents can contain expert recommendations, best practices, and insights from previous problem-solving efforts.

Literature review is an effective approach for identifying and studying relevant research papers, journals, and industry reports to gain insights into the current state of the art, challenges, and opportunities in the field. This can help to identify key concepts, techniques, and areas that require further investigation.

Recent studies suggest that AI techniques can be used to generate scientific understanding from data. AI can act as a "computational microscope" by providing information that is not yet attainable through experimental means, or as a "resource of inspiration" that expands the scope of human imagination and creativity. These dimensions of AI-assisted understanding require human refinement to develop into a full understanding. AI can also act as an "agent of understanding" by producing generalized insights from data and applying them to different tasks, as demonstrated by techniques like transfer learning. Exploring these three dimensions of AI-assisted understanding may lead to new domains of knowledge and understanding for the water sector.

Extracting Knowledge from Experiments: Experiments play a vital role in understanding the functional insights and the underlying factors that influence a given problem statement. Conducting experiments can provide procedural knowledge, which involves knowing how to perform a task and includes strategies, procedures, and rules. These experiments can take various forms, including laboratory experiments, field experiments, or scale experiments. Scale experiments are particularly useful when building and testing the actual target is not feasible due to constraints such as cost, resources, or time. For instance, to test a bridge design, it may not be viable to build a full-scale bridge at the target location, as it requires significant effort and money. Instead, a scale model can be built and tested in the laboratory to evaluate how the bridge would perform in a real-world scenario.

Another example of an experiment is to compare the effectiveness of various methods for removing contaminants from water, such as filtration, chemical treatment, or ultraviolet light exposure. The contaminant levels can be measured before and after each technique is applied to determine which method is most effective. These experiments can help to identify the best approach to solve a particular problem statement and can guide decision-making in the water sector.

3.5.2 Knowledge based selection of input and output features for the AI Application

The context of input parameters here is from the domain knowledge perspective and not the input for machine learning models. Selecting input parameters for AI applications requires a thorough understanding of the scientific process guiding the problem statement. The domain expert should analyze the relevant equations and laws to identify the parameters that govern the process. Input parameters can take various forms, such as images, variables, and numeric values. The choice of input parameters should best represent the problem statement at hand. For instance, when predicting water flow in a natural sub-system like a river, factors like location and season are crucial. In a mountainous region, precipitation levels, snow pack, and temperature would be relevant, while in coastal areas, tide levels and sea surface temperature would be more significant. Additionally, domain knowledge derived from experience, experiments, and data can help determine which set of parameters will serve as the input to define the problem statement's output.

Selecting the right output parameters is critical to the success of an AI application as it determines its usefulness and effectiveness. Defining the goals of the AI model is essential, and this involves identifying the specific task or objective that needs to be accomplished. To select the output parameters, you need to have a clear idea of the desired state of the output, which could be a specific value, a range of values, or a binary decision. The output parameters should assist you during the decision-making stage. Output parameters can be in various forms, such as a scale of 1-5 or a time series graph, depending on the problem statement. They significantly influence the decision-making process and are the defining factor of any given problem statement. Domain knowledge complements the selection of input and output parameters, and it is necessary to identify the desired state of the result for selecting the parameters that will be representative of this state, which we want to achieve with the help of the AI model.

3.6 *ai*WATERS Model Development

Designing and developing an AI/ML model is crucial to the success of an AI system. To ensure model robustness, it is important to consider factors such as input, output, model hyperparameters, training, and testing samples based on the problem statement being solved. Key considerations when selecting an AI model include problem type, data characteristics, model complexity, training and prediction time, and model interpretability. For example, understanding the type of problem to be solved (classification, regression, etc.), the data's size and structure, and the time required for training and predictions are critical in selecting an appropriate model. Model complexity and interpretability are other important considerations, as more complex models may capture intricate relationships but may lead to overfitting, while simpler models may not capture underlying relationships, leading to underfitting.

3.6.1 Selecting Input and Output Parameters for the AI Model

To optimize the performance of an AI model, selecting input parameters is a crucial process that requires experimenting with different parameters and analyzing the correlation between important attributes collected through domain knowledge.

In the practice review, we found that one of the utility used correlation heatmap to select high correlated parameters as their input for the mode. However, this may not lead to optimal model performance for all applications due to feature redundancy. For instance, consider that a utility is trying to solve the problem of predicting water demand based on various factors such as temperature, precipitation, day of the week, and holidays. They have a dataset containing these factors as input features and water demand as the output variable. Here, temperature and precipitation could be considered redundant if they both provide very similar information about the water level. For instance, if it is observed that there is a strong positive correlation between temperature and precipitation in a certain region, it is likely that they are both capturing the same environmental condition that affects water level. Thus, one of the features may be enough to represent the environmental condition affecting water level, and the other feature might not add any additional information or insights. Thus, including both features may lead to redundancy and can even decrease the performance of the ML model.

To identify and address feature redundancy, various techniques such as correlation analysis, principal component analysis (PCA), and feature selection algorithms can be employed. These techniques can help identify redundant features and select the most informative and relevant features for the model.

In addition to input selection, the output form is also important. The choice of model (prediction-based or classification-based) and the output variable (scale, class, or numeric value) may vary depending on whether a classification-based or numerical/predictive result is needed. For example, a utility may want to predict the number of years its pipes can survive (a prediction problem that requires a numeric output/result) or rate the current condition of a pipe as “good” or “bad” (a classification problem that requires identifying a class as output). Thus, the output of the AI model should vary based on the problem statement.

3.6.2 AI Model Selection and Hyperparameter Optimization

Developing and selecting the best machine learning (ML) model for a task can be a complex and tedious process due to the numerous available models and parameters that affect their learning. Additionally, these models are often referred to as black boxes, adding to the difficulty. A simplified workflow can help guide developers toward selecting the best models to solve the problem statement. We intend to guide the utilities towards ideal model development and selection by through a three step process, as depicted in Figure 3.11.

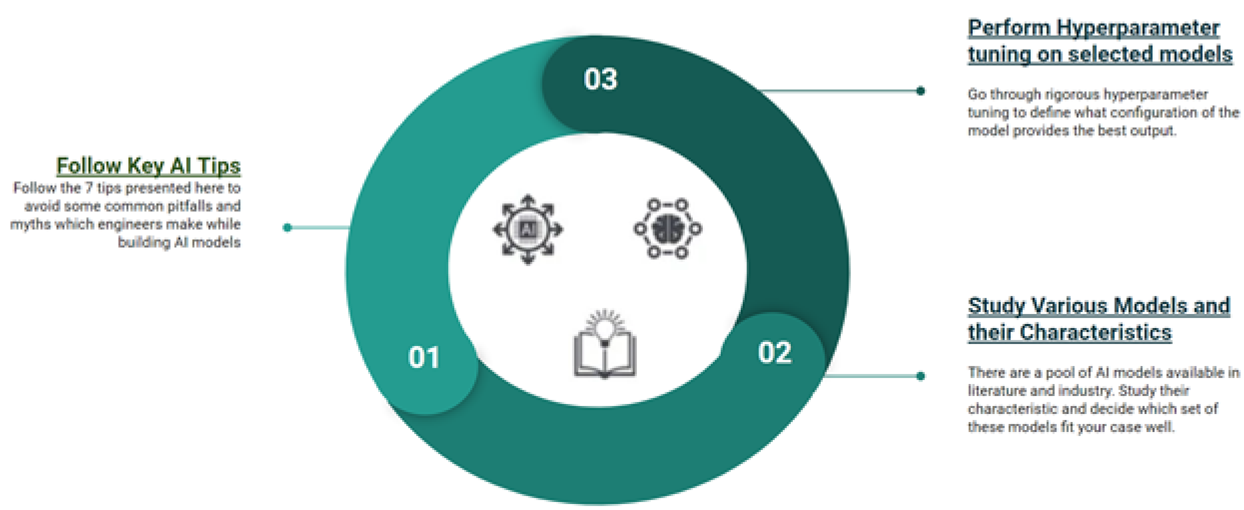

Figure 3.11: Proposed Model Selection Process for Utilities

The three step process involves following seven key tips suggested in Table 3.2 [53], studying various models and their characteristics, and performing hyperparameter tuning to fetch the optimal results from the AI model.

Table 3.2: Key Tips for AI Model Development

Key Tips	Description
1. Avoid the State-of-the-Art Trap	AI Model developers can be tempted to prioritize experimenting with newer models over relying on traditional approaches. It is important to note that researchers typically evaluate models in controlled academic settings, where their models may outperform existing models on a static and limited dataset. This does not necessarily guarantee that the same model will perform better than other models on a utility's unique dataset.
2. Start with the simplest model	The principle of ”simple is better than complex” is applicable to machine learning as well. Simplicity serves three purposes: First, simpler models are easier to deploy. Second, starting with a simple model and gradually adding complexity helps developers understand and debug the model. Third, the simplest model can serve as a baseline for comparing more complex models.
3. Avoid Human Bias in selecting Model	A developer is more excited about a particular architecture or algorithm will likely spend more time experimenting with it. If 100 experiments are being run for a particular architecture, it is not fair to run a couple of experiments for the architecture being evaluated against it.
4. Evaluate Good performance now versus good performance later	The best model now does not mean the best model a few months later. For example, if less data are available right now, a tree-based model may work better. However, if after a few months when a great deal of data are available, a neural network might perform better. A good way to estimate how model performance will change is to use a learning curve (e.g., training loss, training accuracy or validation accuracy). It can give you a sense of whether you can expect performance gain at all from more training data
5. Evaluate Trade Offs	A more complex model might deliver higher accuracy, but might require a more powerful machine, for example, a GPU instead of a CPU. Similarly, a more complex model might give better results, but the results may be less interpretable
6. Decide among In-house Vs Academic Vs Vendor Models	As discussed previously, there are various ways of building models at water utility. Models developed in-house can be potentially rich with domain knowledge. Academic models may present advanced techniques that outperform state-of-the-art models, but academic models are tested on limited data. Vendor-based models might not match the in-house domain knowledge, but they can present techniques that have been effective for many in the market. The problem should be analyzed prior to deciding which model can serve the utility best.
7. Strengthen Data x Knowledge Integration	ML models are a black box that are heavily dependent on data, whereas science-based knowledge models revolve around scientific theory. By making use of both knowledge and data, patterns can be extracted from data without ignoring the knowledge accumulated in scientific research.

Key Tips to Improve Model Development

Understanding Various Models and their Characteristics

Different models have different characteristics based on the type and nature of data and the problem statement. For instance, nearest neighbor classifiers are sensitive to noise and cannot handle missing values, while naïve Bayes classifiers are robust to noise points and irrelevant attributes. Bayesian networks and logistic regression can manage correlated and redundant attributes, whereas artificial neural networks fail to handle instances with missing values. Support vector machines can deal with irrelevant attributes. Deep learning models are effective in high-dimensional settings [63]. Therefore, based on the problem statement and data, a few models can be selected and tested through experimentation to identify the optimal model.

Following is a list of some of the widely used models and their characteritics which can enable the utilities to select the best technique for their use case:

Nearest Neighbors Classifiers

- Since they consider distance of data points/objects, they are highly susceptible to noise
- Nearest neighbor classifiers can produce decision boundaries of arbitrary shape, and thus provide more flexible model representation
- Have difficulty handling missing values
- Can handle interacting attributes (i.e the attributes which are interdependent to each other.)

Naïve Bayes Classifier

- Robust to isolated noise points because such points are not able to significantly impact the conditional probability estimates
- Can handle missing values
- Robust to irrelevant attributes
- Correlated attributes can degrade the performance

Bayesian Network

- Can easily handle the presence of correlated or redundant attributes as they don't consider conditional independence.
- Robust to the presence of noise in the training data
- Can handle missing values during training as well as testing

Logistic Regression

- Can handle irrelevant attributes
- Can handle redundant attributes
- Cannot handle data instances with missing values

Artificial Neural Networks

- Can handle irrelevant attributes
- ANN has a tendency to get stuck in local minima
- Training an ANN is a time consuming process, especially when the number of hidden nodes is large

- ANN can learn in the presence of interacting variables

Support Vector Machine (SVM)

- Can handle irrelevant attributes
- SVM can be applied to categorical data

Hyperparameter Tuning

Hyperparameter tuning is the process of selecting the optimal values for the hyperparameters of a machine learning model. Hyperparameters are parameters that are not learned from the data, but rather set by the user before the training process begins. These include things like learning rate, regularization strength, and number of hidden layers in a neural network.

Hyperparameter tuning is important because the performance of a machine learning model can vary significantly based on the hyperparameters chosen. The goal is to find the combination of hyperparameters that results in the best model performance on a validation set of data. This is typically done using a search algorithm that iteratively tests different combinations of hyperparameters.

The ideal approach to hyperparameter tuning involves a combination of automated and manual tuning. Automated methods, such as grid search, random search, and Bayesian optimization, can quickly test a large number of hyperparameter combinations. However, manual tuning by an expert with domain knowledge, prior experience and experimentation can often lead to better results by allowing for more nuanced adjustments.

To start, it's important to establish a baseline performance with default hyperparameters. Then, a search algorithm can be used to test a range of hyperparameters in a systematic manner. This can be done using techniques like grid search, which tests all combinations of hyperparameters within specified ranges, or random search, which randomly samples hyperparameters from specified distributions.

Finally, the best performing hyperparameters should be validated on a hold-out test set to ensure that the model's performance is not overfitting to the validation set. It is important to note that hyperparameter tuning should be done on a separate validation set to prevent overfitting to the test set.

3.6.3 Model Generalization and Maintenance

When the AI model is tested against unseen instances of data (i.e. the test set), one of the following three scenarios is observed [64]:

- **Underfitting:** It occurs when a model is not able to learn the underlying relationship between the input variables and the target variable, resulting in poor performance on both the training and testing datasets.
- **Overfitting:** It occurs when a machine learning model is excessively complex, capturing noise and irrelevant information in the training data that is not present in the testing data or real-world data.
- **Ideal Fit:** It is achieved when the model has both high accuracy on the training data and good generalization performance on new, unseen data.

If the utilities find their model to be overfitting or underfitting, they can resort to these approachs to overcome the same:

In case of underfitting, the developer can take the following steps:

- Increasing model complexity: Increase the complexity of the model to make it deeper, so that the model is able to capture more patterns. Hyperparameters can be tuned, such as adding more layers and number of nodes in each layer to effect higher learning from the data.

- Using a more powerful model: Switching to a more powerful model, such as a deep neural network, can help overcome underfitting.

- Feature engineering: Engineering more informative features from the raw data can help improve model performance and overcome underfitting.

Whereas, in case of Overfitting, the following approaches can be followed:

- Regularization: Regularization techniques, such as L1 and L2 regularization, can help reduce overfitting by adding a penalty term to the loss function that discourages the model from overfitting to the training data.

- Dropout: Hyperparameter tuning can also be applied in this case, like using dropout to drop a significant percentage of nodes during the training process, reducing layers and nodes in the model to decrease the aggressive memorizing nature of the model.

- Early stopping: Monitoring the performance on a validation set during training and stopping when performance on the validation set starts to degrade can help prevent overfitting.

- Ensemble methods: Ensemble methods, such as bagging and random forests, can help reduce overfitting by combining the predictions of multiple models.

To ensure the long-term performance of AI models, maintaining an ideal fit is not enough. Several factors such as data drifts and continual learning need to be addressed after deploying the model. Over time, ML models can become stale and sensitive to changes in the data associated with the problem statement. This is known as data drift, which occurs when there is a change in the statistical distribution of data from the baseline data used to build the model. Consequently, model performance can start to degrade, and continuous learning is crucial to maintain a healthy performance.

Apart from changes in the data, the conditions under which the model's decisions are applied can also change over time. Therefore, it is essential to integrate the latest domain knowledge into the models. This enables the models to make optimal decisions that take into account changing external conditions. For example, in the water sector, changes in weather patterns, availability of water sources, and water quality can all lead to changes in the statistical distribution of data. As a result, models that were once performing well may no longer be effective without continuous learning and integration of domain knowledge.

3.7 *ai*WATERS Decision Support

Machine learning systems involve not only technical aspects but also require input from various stakeholders such as business decision-makers, users, and developers. Although AI models can provide outputs without human involvement, it is crucial to exercise caution while trusting them for high-stakes real-world applications. The human-in-the-loop approach is essential to ensure the reliability of the outputs from trustworthy AI models. The decision-making process using AI technology involves several elements, such as trustworthiness of the AI, human involvement, and interpretability of results, which require careful consideration to make informed decisions. This subsection will delve into these elements in detail.

It is important to note that the impact of AI in water utilities is not limited to the processes taking place within the utility. The implementation of AI in water utilities requires a multi-disciplinary approach that involves experts from various fields. The impact of AI-based decision-making systems on water resource management is complex and far-reaching, involving social, environmental, legal, and ethical considerations.

AI systems are not perfect, and some decisions might threaten social and environmental justice. For example, access to clean and safe water is a basic human right and must not be compromised by AI-based decisions. Therefore, any AI system implemented by water utilities must consider the equitable distribution of water resources, ensuring that no social group is denied access to this essential resource.

Additionally, AI must also consider the environmental impact of its actions. Water utilities have a significant impact on the environment, and it is crucial that any AI-based decision-making system prioritizes the preservation of ecosystems and the protection of wildlife. Considerations include minimizing the risk of water pollution, protecting aquatic habitats, and preserving biodiversity.

Therefore, the integration of AI in water utilities must be seen as a cross-disciplinary effort that involves experts not limited to computer science or civil engineering, but also social scientists, environmental scientists, lawmakers, and others. This holistic approach is essential for ensuring that the impact of AI-based decision-making by water utilities is thoroughly understood and that the outcomes are socially, environmentally, legally, and ethically acceptable.

The rise of AI has also sparked discussions about its potential impact on jobs and the workforce. There has been a fear of AI replacing the humans for various tasks. While it is true that AI has the capability to automate certain tasks and processes, its impact on employment in the water sector, like any other industry, is a complex and nuanced topic. AI can bring significant benefits by improving operational efficiency, optimizing resource allocation, and enhancing decision-making processes. It can automate repetitive and mundane tasks, enabling water utilities to allocate their human resources to more strategic and higher-value activities.

It is important to note that AI is not necessarily a direct threat to jobs in the water sector. Instead, it can be seen as a tool that augments human capabilities, leading to a transformation of job roles rather than outright replacement. While certain routine tasks may be automated, AI also creates new opportunities for skilled professionals to work alongside AI systems, leveraging their insights and outputs to make informed decisions. Moreover, the water sector is inherently complex and requires domain expertise, regulatory compliance, and direct human intervention for various tasks. AI systems can provide valuable insights and support decision-making, but human oversight and interpretation remain crucial.

3.7.1 Ensuring Trustworthiness of AI Decision Support

AI systems have advanced significantly, but their increasing complexity and opaque decision-making can lead to limitations like bias, ethical concerns, and lack of transparency and accountability. Trusting the AI output is an important topic of discussion as the impact made by the AI based decisions can be of varying magnitude. For instance, the decisions of music recommendation application that uses a machine learning algorithm to suggest songs to a user will not be considered a high impact decision for the user. However, if a water treatment plant uses AI to optimize the chemical dosing, and if the model is biased towards certain input variables, it could lead to incorrect dosing and affect water quality which can have severe consequences for the environment.

We found that the European Union (EU) has developed guidelines for the design and development of trustworthy AI systems [65]. The guidelines focus on three key aspects: lawful, ethical, and robust AI systems, as illustrated in Figure 3.12. In the context of trustworthy AI, lawful means that the AI system and its decisions should comply with relevant rules and regulations. Ethical means that the AI system should align with ethical principles that prioritize human values. Robust means that the AI system should be technically strong and capable of producing optimal output, even after adhering to the principles of ethics and law.

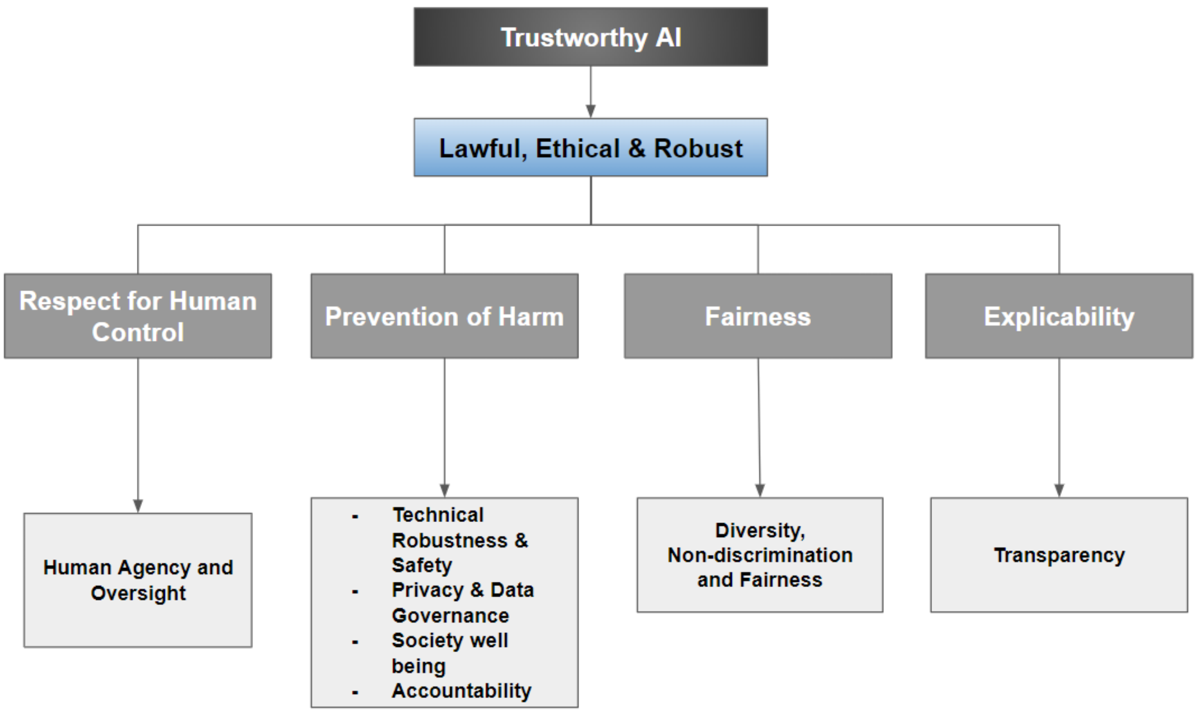

Figure 3.12: AI Trustworthiness guidelines

These three guidelines are supported by a framework that includes four ethical principles: Respect for Human Control, Prevention of Harm, Fairness, and Explicability, which are explained as follow:

Respect for Human Control means that AI systems should complement human decision-making rather than replace it. **Human agency and oversight** should be maintained, and the degree of autonomy should be based on the level of risk involved. For example, an AI system used in water management can assist operators in decision-making, but the final decision should always be made by humans.

Prevention of Harm is crucial in building trustworthy AI systems. Technical robustness and safety should be ensured to prevent errors and outside attacks. **Privacy and data governance** should also be maintained, and sensitive data should be protected from misuse. The AI system should also consider the **societal and environmental** impact of its decisions. For instance, a water utility may use AI to analyze customer data to identify who needs water conservation services. Still, the data usage should be transparent, and customer privacy should be protected.

Fairness is essential in building trustworthy AI systems. AI should treat all social groups equally and without discrimination. It should also ensure that AI services are accessible to all. For example, an AI system used to predict customers who may fail to pay their water bills should be built without bias towards any social or income group.

Explicability is necessary for building transparent AI systems. AI decisions should be interpretable and explainable, and the system should be reproducible. For example, an AI system used to detect water leaks should provide an explanation of how it arrived at a particular prediction, making it easier to identify and correct errors.

3.7.2 Involving Human in the Loop for Decision Support

The term "human in the loop" refers to the involvement of humans in various roles with AI systems, including model developers, end-users, and domain experts. This involvement can be classified into three categories: humans before the loop, humans in the loop, and humans over the loop, as shown in Figure 3.13 [65].

Humans before the loop are involved in the design phase of the AI system. They contribute to the planning, designing, and drafting of requirements for the AI system. Developers,

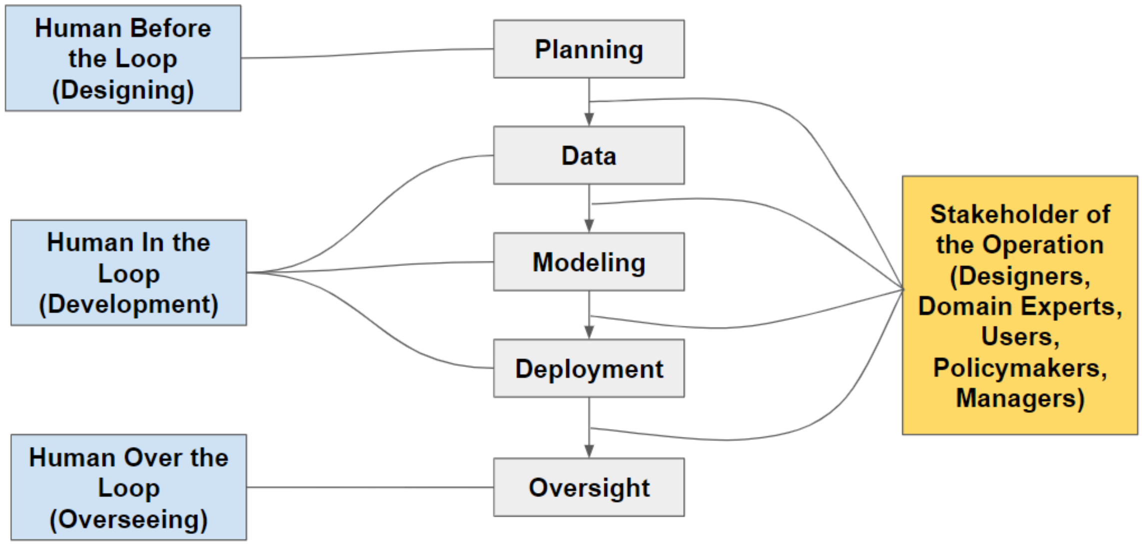

Figure 3.13: Types of Human involvements in an AI System

policy and domain experts, and end-users are "in the loop" to state their requirements.

Humans in the loop are involved in building and developing the AI system. They work on data collection, data processing, model development, knowledge integration, and continuous updating of the model. They monitor the flow and AI output at every step, which is automated by AI, and make required changes if the flow is not as expected.

Human over the loop monitors the output and decisions provided by the AI system and evaluates the system's performance. If necessary, they can override the AI-based decision if it is not suitable for real-world scenarios.

For example, in the water sector, humans before the loop could include water utility managers and engineers who plan and design the AI system. Humans in the loop could include data scientists who collect and process data to develop the AI model, as well as domain experts who provide insights into the water distribution network. Human over the loop could include regulatory authorities who monitor the performance of the AI system and evaluate its impact on the water sector.

In the application goals pillar, we defined the AI deployment level as manual, semi-automated, or fully automated. While manual and semi-automated deployment requires human involvement, humans need to be in the loop at every stage of the process to observe the flow of decisions made by AI in a fully automated scenario as well. Having humans in the loop will enable utilities to enhance trustworthiness and transparency in the system.

3.7.3 Using Effective Visualizations for Decision Support

Visualizations are a powerful tool for enhancing decision support in the water sector. They allow model developers, decision-makers, and other stakeholders associated with AI systems to draw simple and actionable conclusions. Effective visualizations can impact decision-making and judgment in the following ways [66]:

Judgement/Decision Accuracy: Well-designed visualizations can help improve problem comprehension and the quality of inference from data. As a result, better judgments and decisions can be made. For instance, visualizations of water quality data can help decision-makers understand patterns and trends that may be difficult to identify from raw data.

Response Time: Information Visualization can speed up the response time in various judgment and decision tasks since it can highlight patterns and inferences associated with the data. Effective visualization makes data interpretation simple.

Decision Confidence: Studies have shown that visual decision support systems that depict uncertainty in engineering design can lead to lower decision confidence, compared to traditional methods omitting uncertainty information.

Willingness to Act: Visualizations of KPIs (key performance indicators) can motivate managers' intention to act on information when compared to text

Empower Collaboration: Effective visualizations can make it easier for people to explain and interpret the results and decisions to officials/stakeholders across the organization, making it easier to collaborate with various entities. For instance, visualizations (instead of raw numerical data) of water demand predictions can help the water utilities communicate their conservation strategy to policymakers of the city, and can enable the non-technical end user to get insights into technical findings.

Table 3.3: Stakeholders vs Relevant AI Usecases for Visualizations [1]

Stakeholder	Usecase
ML and AI practitioners	Better understand how well the model might work for the intended use cases and track its performance over time
Model developers	Compare the model's results to other models in the same space and make decisions about training their own system
Software developers	Working on products that use the model's predictions can inform their design and implementation decisions
Policymakers	Understand how a machine learning system may fail or succeed in ways that impact people
Organizations	Make informed decisions about adopting technology that incorporates machine learning

Table 3.3 provides a summary of what various stakeholders at a utility may require from a visualizations for decision making usecases [1].

3.8 *ai*WATERS Implementation

As shown in the various pillars of AI, successful implementation of AI can be fraught with various challenges affecting the long-term sustainability of AI-based platforms and solutions.

This section addresses the various challenges in successful implementation of AI to ensure utilities can develop long-term, usable decision-support platforms.

Acceptance of AI Technology:

It is a myth to consider that AI will fully replace the human workforce. In the water sector, the more productive perspective is to view AI as a tool that can assist humans in decision-making processes. Utility managers should understand the basic characteristics of AI and start with small pilot projects to assess its potential benefits and make informed decisions before full implementation. This approach is particularly relevant in the water sector, where non-replaceable support staff with licenses at stake are necessary for working with AI. Currently, utilities lack guidance on identifying where AI-based techniques can provide the greatest benefit and estimating the financial, staffing, or workflow benefits of deploying AI-based solutions. Demonstrating and measuring benefits from AI solutions can help calculate the return on investment and support a business case for deploying advanced AI analytics solutions for decision-making.

Get Acquainted with AI and ML:

To effectively deploy AI, utilities must understand AI fundamentals and look at success stories of successful implementation from other utilities or industries. For this purpose, utilities must share knowledge within and outside their utilities from AI experts who have successfully implemented AI. Teamwork is vital to this process to enable knowledge sharing from the AI experts as well as the senior staff who may have the necessary domain knowledge that could be built into robust AI decision-support systems.

Identify the challenges you face

Not all water utility obstacles and challenges will have the same answer. Challenges can be seen as opportunities to contribute to the growing AI implementation in the water sector.

There should be some knowledge sharing platforms (like waterid.org) where utility experts could share success stories and lessons learned through case studies. Usually, vendor-based AI pilots are difficult to share due to restrictions on sharing the details of the AI platform from the vendors. However, utilities should strive to understand the complete picture before accepting or denying any solution. The case studies should focus on the technical, business, natural, social, and financial aspects of the piloting programs for better understanding and application of AI.

Define the areas that need AI Application(s)

To define the application areas of AI in the water sector, utilities should ask themselves several questions:

- Where are the potential applications of AI?
- In what areas do data and technology have a significant impact?
- Which areas are flexible enough to quickly implement innovations?
- What resources are required for AI implementation and what do we currently have?

Although utilities may identify multiple areas where AI can improve processes, it is essential to focus on top priorities. Priorities should be identified based on the areas that will result in the following:

- Increased benefits and reduced costs
- Optimized processes critical to business success
- More data-driven initiatives in the utility
- Improved interaction with stakeholders and employees.

Prioritize the main driver(s) of Value

Once the utility defines its business needs, it should identify the potential business and financial benefits of the AI project. Consider all possible AI implementations and connect each initiative with concrete returns by focusing on near-term goals and illustrating either financial or business value based on cost-benefit analysis. Do not lose sight of value drivers such as increased value for customers or improved employee productivity. Consider whether machines could handle specific time-consuming tasks more effectively than people. Avoid implementing solutions purely based on algorithm popularity. Instead, analyze how a solution fits into daily workflows, business processes, and services to boost operations over the long-term.

Strategic Planning before AI Implementation

AI is a powerful tool, but it requires responsible implementation to generate value. Developing an AI strategy is crucial for ensuring that early efforts add value and lead to future development and investment. Without a plan, AI initiatives are more likely to fail. The AI implementation plan should provide architectural and best-practice recommendations to aid domain experts, data scientists, and ML engineers in developing reliable solutions. The plan should also focus on selecting applicable technologies to improve strengths and weaknesses, with a philosophy of continuous improvement for the AI platform. Governance protocols should be established for updating, verifying, validating, and improving the user interface/user experience (UI/UX) of the AI model.

Staff the AI Team (Human-in-the-Loop AI)

To ensure successful AI implementation in a utility, a diverse team of managers and experts should work together to address potential barriers. The responsibility of developing and applying responsible AI techniques should not be solely on data scientists, but require input from water sector domain experts, social and environmental scientists, and humanities experts. Collaborative teams can be formed to advise or take action to ensure the development of fair and responsible AI. The human-in-the-loop approach shifts the focus from building an intelligent AI system to incorporating meaningful human interaction and trustworthiness.

Form a Data and Knowledge and Model Development Taskforce

AI can benefit utilities in multiple ways, but it cannot function in isolation. Instead of operating in vertical business divisions, AI can be used horizontally as an enabler across the utility. This requires a cross-functional team of experts and stakeholders to support the core AI implementation team and collaborate with other departments to identify high-impact initiatives. Combining databases and resolving disagreements are crucial steps to obtain high-quality data for successful AI implementation.

Set Up a Pilot Project(s) and start small

Before investing valuable resources into the implementation of a new technology, it is crucial to ensure its benefits. Pilot programs provide a way for utilities to experiment on a small scale before making large investments. Utilities can either rely on cloud services or construct their own AI infrastructure when starting out. Many vendors offer tools to assist in the initial implementation of AI operations. Although cloud options are initially less expensive, investing in in-house infrastructure might pay off in the long term. It is important to assess requirements such as scale and resources needed before making a choice, as benefits and drawbacks vary depending on the utility.

Identify the Gaps in the Process after AI Implementation

AI models must be retrained on a regular basis, with feedback supplied for improvement. When it comes to deciding where changes are necessary, careful examination and classification of mistakes can be of great help. Well-developed error analysis and continuous improvement protocols should allow data scientists to examine a significant number of unseen faults and get a thorough knowledge of the types of errors, their distribution, and their origins in the model.

Evaluate AI Adoption Potential

To reduce operating expenses and boost productivity, many utilities consider automating internal business operations. AI-driven business operations can be highly beneficial for some utilities, but a mismanaged AI project can be a waste of resources and discourage the utility from future AI initiatives. Thus, it puts them behind other utilities that have successfully implemented AI applications. A poor initial use case for the trial or proof-of-concept is often blamed for the failure of an AI project. However, a utility should not be deterred from using AI even if it is challenging to identify the appropriate AI application opportunity.

Keep a Check on Performance and Behavior

AI systems in the water sector are complex and difficult to evaluate. A governance approach can help. This involves defining performance and behavioral metrics, evaluating models periodically, and collecting metadata. The impact and benefits of AI implementation can also be assessed. Monitoring should extend to the training and test stages. AI teams should incorporate monitoring as an intrinsic part of the utility AI lifecycle and put their findings into practice to improve models over time.

Chapter 4

Results

In this section, we will answer the research questions mentioned in Section 1 based on our pilot interviews and questionnaire surveys with various water utilities.

Water and wastewater systems in the United States can have diverse characteristics and challenges. This diversity can arise due to differences in the number of customers served, the amount of water produced and distributed, the complexity of the infrastructure, and the level of regulatory oversight. To capture patterns regarding how the AI implementation varies based on these characteristics, we divide the utilities into three segements:

- Small Scale Utilities: They serve a population of less than 50,000 and are often located in rural areas. They tend to have simpler treatment processes and distribution systems, with fewer wells, pumps, and storage tanks. Small utilities may be operated by local governments, private companies, or cooperatives, and may face challenges in funding and maintaining their infrastructure.

- Medium Scale Utilities: They serve a population of 50,000 to 250,000 are located in suburban areas or smaller cities. They have a more extensive water and wastewater infrastructure system compared to small utilities, including multiple sources of water supply, treatment plants, and storage facilities. Medium utilities are typically operated by public agencies, such as municipal or county governments, and are subject to regulatory oversight and reporting requirements.

- Large Scale Utilities: They serve a population which is greater than 250,000 and are located in larger cities or metropolitan areas. Large utilities have complex treatment processes and distribution systems, with multiple sources of water supply, treatment plants, and storage facilities. Due to their size and complexity, large utilities may face challenges in managing their infrastructure and meeting the needs of a diverse customer base. They may also have more resources available for implementing new technologies, such as AI, to improve their operations and decision-making processes.

To ensure diversity in terms of the scale and geography of the utilities, we took response of 6 large scale utilities, 2 medium scale utilities and 2 small scale utilities, spanning across various states in the United States, as shown in the Figure 4.1a and 4.1b. However, all the participating utilities for this research were public utilities.

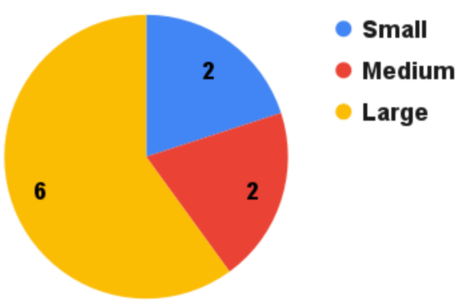

(a) Scale of Piloting Utilities

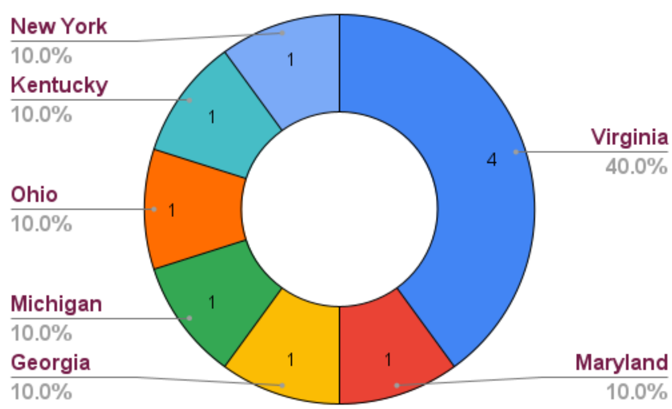

(b) Geographic Distribution of Utilities

Figure 4.1: Scale and Geography of Piloting Utilities

4.1 RQ1: Are utilities willing to develop AI solutions and use it across the system?

The goal of this research question was to understand if the utilities are currently using AI in their system and if they would continue or like to explore AI solutions in the future. Eventually, we also want to know if they envision AI to act as a system engine for their utilities, i.e., AI is applied to each and every process of the system.

As shown in Figure 4.2, only a small portion of medium and Large scale utilities are currently using AI at component level, i.e. specific to a process (like wastewater treatment, CCTV image classification, etc). Thus, it is evident that the current usage of AI in the water industry is limited and fragmented within the utilities.

AI based water meter monitoring, AI based wastewater treatment, AI based stormwater resiliency, Process Control for Maximization of Energy and Resource Efficiency, Wastewater Network Optimization and AI based CCTV image encoding were some of the use cases for which AI was used by the water utilities.

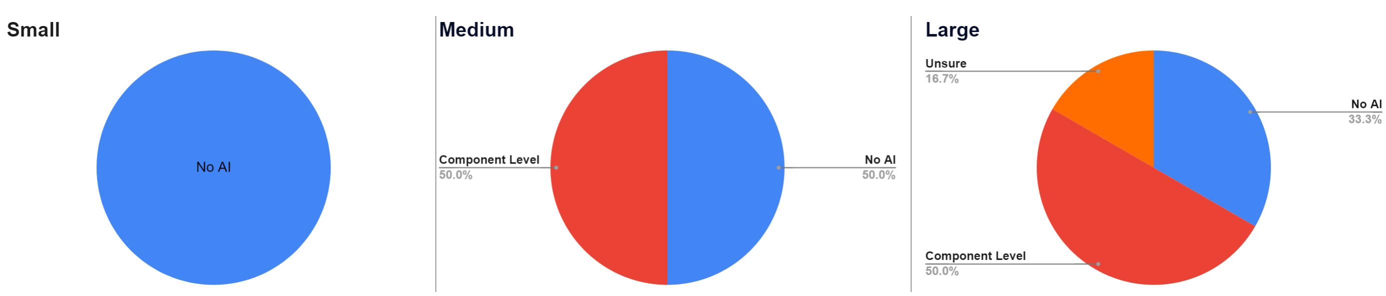

Figure 4.2: Current AI usage in Utilities

However, utilities across various scales have basic understanding of AI and its capabilities and are willing to adopt it over the next 5 to 10 years as they prepare to undergo a digital transformation, as shown in Figure 4.3.

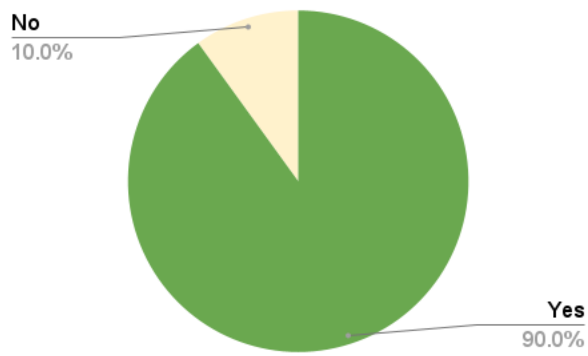

Figure 4.3: Willingness to investigate AI solutions in near future

We found that willingness to adopt AI in the future is not directly proportional to the intent of using AI for each and every process in the utility, as some utilities agreed and some disagreed on this idea while some of them were unsure if it would be a good move, as shown in Figure 4.4

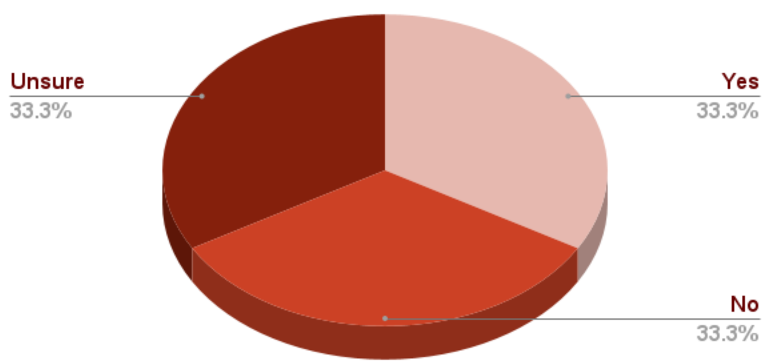

Figure 4.4: Intent to use AI for each and every process in the future

Through our pilot interviews and questionnaire responses, we found that there are various arguments supporting each possibilities. One of the large scale utility argued that using AI for each process is not a good idea as some of the processes at the utility are so simple that attempting to integrate AI overcomplicates it. The utilities in favor of implementing this idea believe that using AI across the system for every process may help in improving the decision making process at the utility. Some utilities are unsure about how this move would turn out for them as there are trustworthiness and sustainability issues associated the AI applications.

4.2 RQ2: How do water utilities of different scale compare to each other in terms of AI implementation?

We found that the utilities, irrespective of their scale, were on the same page in terms of understanding the AI technology as they were correctly able to identify its benefits and challenges. However, considering the in-house expertise and resources required to achieve AI driven goals, there exists a digital divide between the utilities of different scales when it comes to the implementation of AI.

Comparison on Building AI within the Utility:

As we can see from Figure 4.5, large scale utilities have experience working on all kinds of AI model development. This is because they have in-house expertise and also the required resources and funds to collaborate with vendors and academic institutes. Medium scale utilities are in the beginning of their AI journey and they do not have in-house expertise to create AI models. Thus, they are dependent on collaborations with vendors and academia. Small scale utilities have not developed or used any AI models yet.

Some of the utilities also responded to us about the trade offs of using in-house vs vendor based vs academic AI models. A medium scale and a large scale utility had similar opinions as they both mentioned that using in-house models requires more time but they are interpretable as the utilities are directly involved in developing them, whereas vendor based and

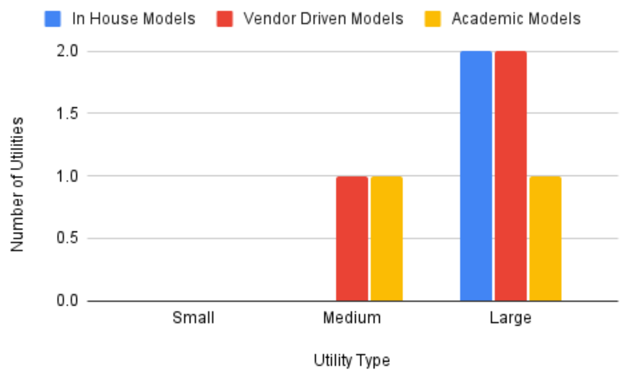

Figure 4.5: Mode of building AI across utilities

academic models are blackbox for them. Small scale utilities mentioned that since they don't have the required in-house resources, they will need to collaborate with vendors. However, that route is not sustainable for long term goals.

Comparison on Data Readiness:

We found that for almost all the utilities we piloted with considered the source reliability, data integrity, timeliness and relevance as important parameters to assess the quality of their collected data, which is in-line with our proposed framework, *ai*WATERS. However, large scale utilities follow a more holistic approach for evaluating the quantity of data as they focus on all the 5V's of Big Data, unlike small and medium scale utilities, as shown in Figure 4.6.

We found that utilities follow the very basic data preprocessing steps like removing duplicates, handling missing values, data compilation and extraction, etc, and are not involved in relatively advanced techniques like normalization, dimensionality reduction, class imbalance etc, as shown in Figure 4.7. This pattern can be attributed to a couple of facts, the first being many of the utilities are not using AI and thus, may not require steps above the basic ones, and second, that there are not many in-house AI experts helping utilities achieve the AI operations.

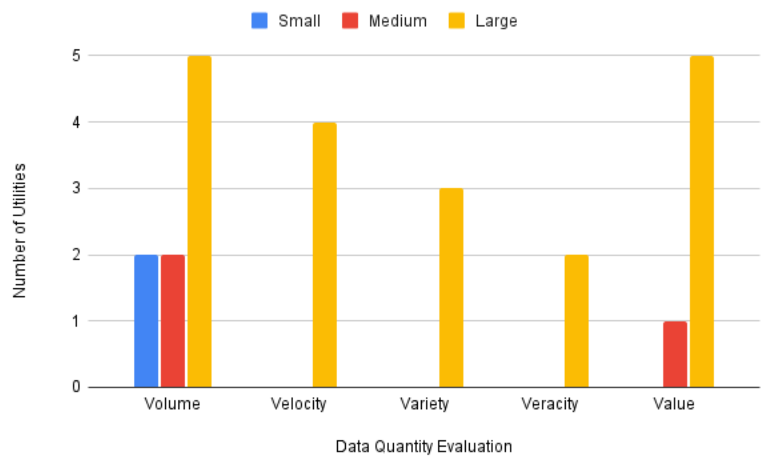

Figure 4.6: Data Quantity Evaluation by Utilities

There are no votes for handling class imbalance as none of the utilities are using techniques like sampling or synthetic data generation.

Large scale utilities mentioned that collecting the right data is highly important for their AI operations and we can say they are on the right path as large scale utilities are following all steps for quality and quantity evaluation. However, in terms of data preprocessing, utilities of all scales are not upto the mark as they do not apply some of the important data preprocessing steps.

Comparison on Domain Knowledge:

Domain experts play an important role at the utilities as we found that utilities are the most dependent on them, followed by data and experiments for extracting domain knowledge for their AI applications, as shown in Figure 4.8.

Comparison on Model Development:

Our questions regarding model development found responses mostly from large utilities, as they develop in-house models and are acquainted with the model development process.

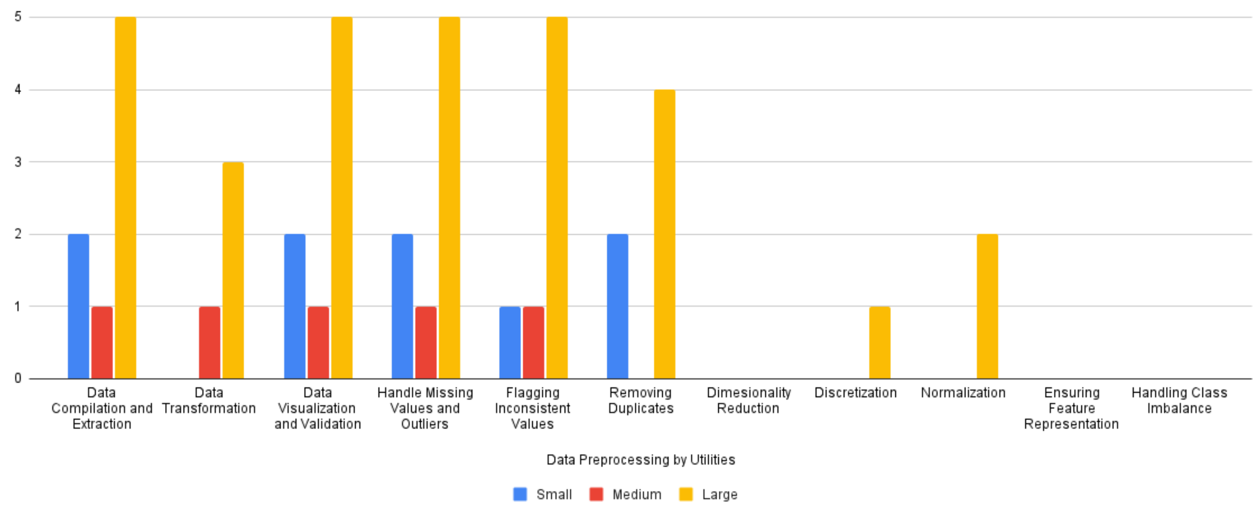

Figure 4.7: Data Preprocessing by Water Utilities

We found that all the large scale utilities **prefer using simpler techniques** (like linear regression, fuzzy logic, etc) **over state of the art techniques** (like, ANN, CNN, LSTM, etc).

Majority of the large scale utilities mentioned that they start by building most simplest model and then, work towards improving or making it more complex, as show in Figure 4.9

Along with large scale utilities, the medium scale utilities also provided their view on whether they are inclined towards experimenting more with their preferred model (like the one which performed well in the past, has good characteristics according to them, etc) or they put equal effort towards experimenting multiple potential models. We found that large scale utilities didn't have a bias towards the models and experimented equally with all potential models. One of the large scale utilities mentioned that they haven't reached this stage where they can comment on this. Whereas, for some medium scale utilities, we saw that there was a bias towards selecting a particular model, as depicted in Figure 4.10.

We also found that most of the medium and large scale utilities are aware about the trade-

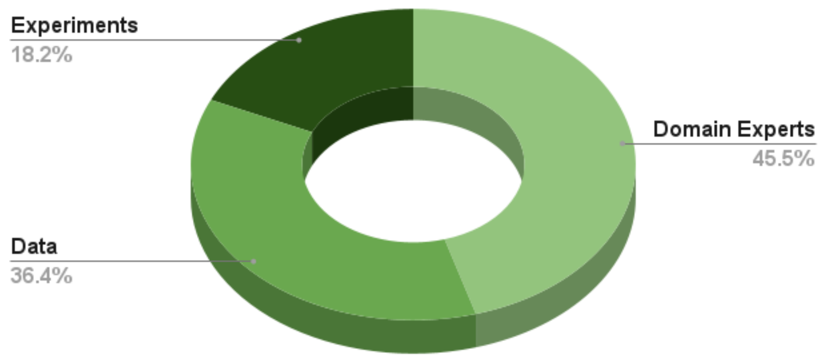

Figure 4.8: Domain Knowledge Extraction at Utilities of all scales

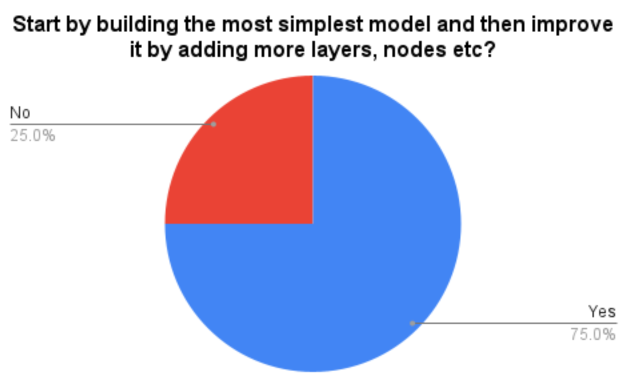

Figure 4.9: Large scale utility preference for starting point of AI models

offs like resource availability, model complexity, model interpretability since they take it in account when selecting a particular model, as shown in Figure 4.11.

One out of two medium scale utilities mentioned that they perform hyperparameter tuning and testing before selecting an AI model. However, one of the responses from the large scale utility highlighted that this information is proprietary to the vendors, as depicted in Figure 4.12. This shows that some of the large scale utilities are buying vendor based AI models for their application but not fully aware of how these models are built.

With respect to the important factor of model generalization, we found that **all the utilities**

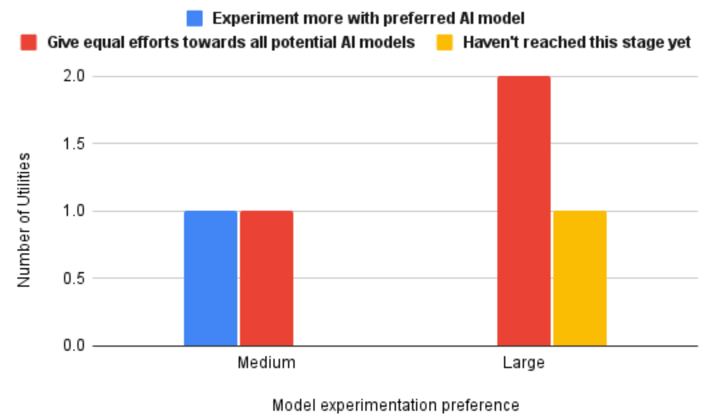

Figure 4.10: Model Experimentation Preference

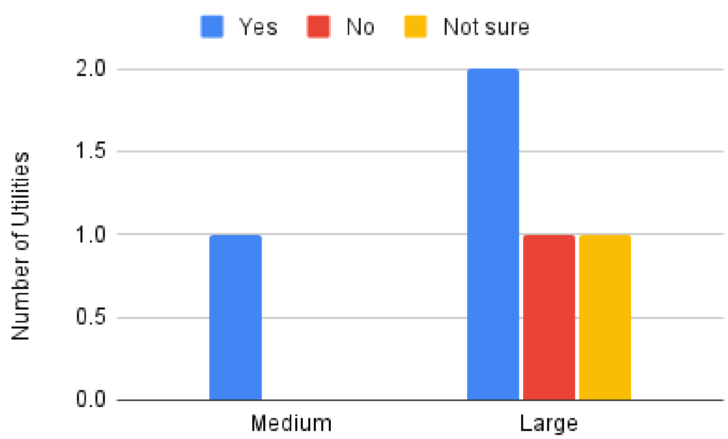

Figure 4.11: Consideration of Model Development Tradeoffs

are facing the problems of poor generalization, as their model performs well on training data but fails to perform on the unseen instances from test data.

Utilities are also not well versed with model maintenance as we got only 3 responses related to the question that asked how often do they update the model (i.e. training with new data, updating hyperparameters, replacing with new algorithm etc). Out of the 3 respondents from large scale utilities, only one of them mentioned that they update their model once every month, as shown in Figure 4.13.

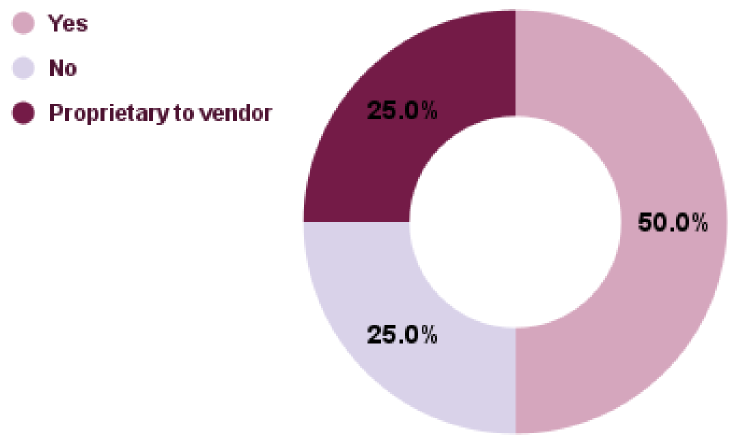

Figure 4.12: Do utilities perform hyperparameter tuning and testing?

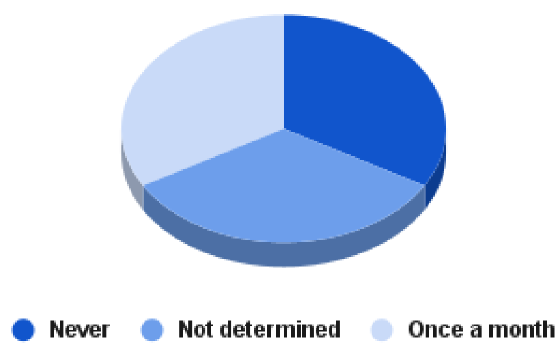

Figure 4.13: Frequency of Updating AI Models at Large Scale Utilities

Comparison on Decision Support

Since trustworthiness of AI models is a concern, we wanted to capture if the respondents had complete knowledge about the AI models being used at their utility, on factors like how the model makes decision, why a particular model was preferred over other models, advantages/disadvantages of a model. We found a diverse set of response to this question, with it being more inclined towards the respondent not versed with the AI models used at their utility, as shown in Figure 4.14.

Our respondents held the following title at their respective utility: Asset Management An-

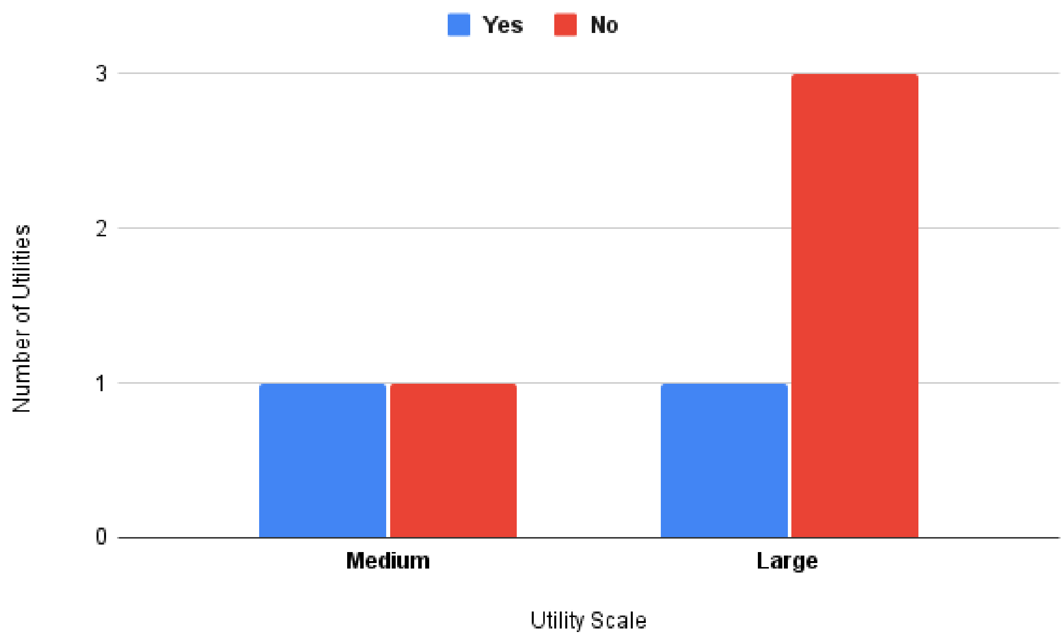

Figure 4.14: Do respondents have complete knowledge regarding the AI model?

alyst, Manager Water Planning, CMMS Administrator, Asset Manager, Field Services Director, Supervising Engineer, Strategic Initiatives & Project Delivery Director, Assistant Director of Engineering, Chief of Data and Technology Operations, Water Quality and Innovation, Process Engineer, Director of Engineering and GIS, Executive Director. As per our request to utilities about connecting with the right person for this study, these were the best point of contact to capture the AI practices at their utility. Lack of data driven roles in this list of titles represents that there was not many AI experts in the utilities, which can account for the blackbox nature of the models being used in the utilities.

To check what factors are significant to utilities when defining trustworthiness for an AI systems, we captured their views on the trustworthiness factors mentioned in the Decision support pillar of *ai*WATERS. Eight out of ten participating utilities responded to this question, and as per figure 4.15, we can say that all the factors are considered important by some or the other set of utilities.

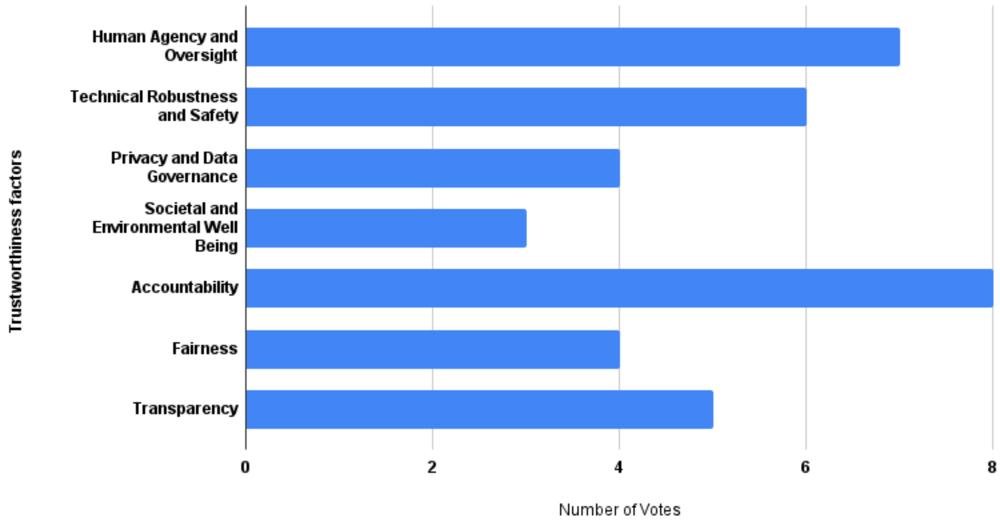

Figure 4.15: Trustworthiness Factors for Water Utilities

However, irrespective of their scale, not all utilities have involved experts from multi-disciplinary fields to ensure, social, technological, ecological and legal acceptance of their AI and technical decision support systems, as shown in Figure 4.16, who play a crucial role in ensuring the trustworthiness of AI output. Involving experts from multidisciplinary field can help the utilities assess if the AI output will be socially acceptable in the long run.

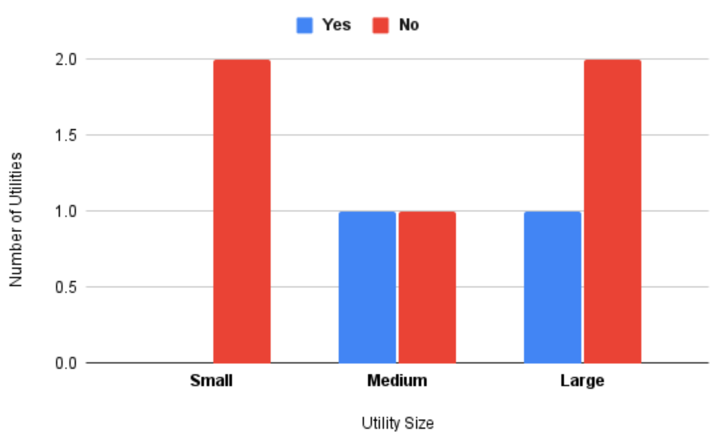

Figure 4.16: Do utilities involve experts from multidisciplinary fields for decision support?

We also found that none of the utilities currently create visualizations specific to the role of the user (as described in Table 3.3) to enhance the process of AI based decision support in their utility.

4.3 RQ3: What aspects of the AI technology are considered as the major challenges by water utilities?

We ask the utilities about what are the major challenges according to them while implementing AI in their system. As per the Figure 4.17, we find that quality data acquisition, developing AI based workforce and integration of AI in existing system are the top three key challenges for utilities.

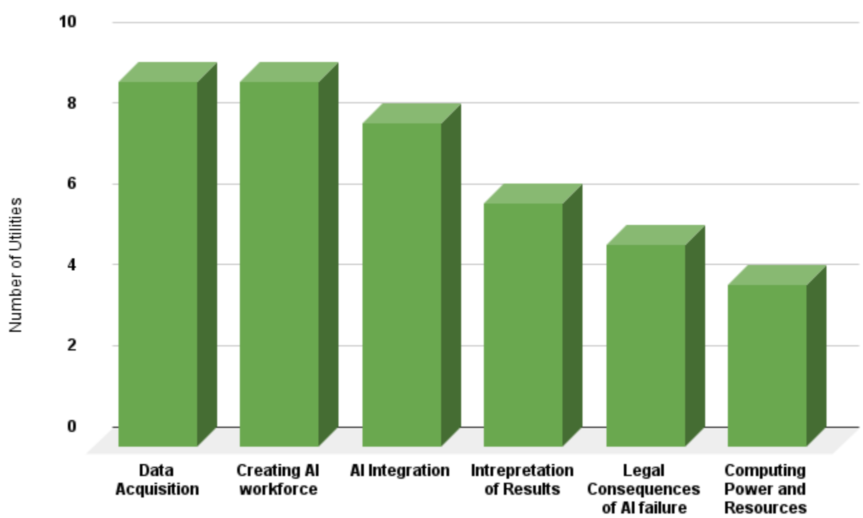

Figure 4.17: Major AI challenges as per water utilities

We also captured the perspective of problems specific to small scale and medium scale utilities based on our pilot interviews. We found that the sustainability and trustworthiness of AI are two major concerns for them. Since they are dependent on vendors, the AI solutions offered to them are blackbox, which makes it difficult for them to trust a system that has real world implications. Sustainability of the AI solutions offered to them is also difficult as they don't have in-house expertise to maintain and improve the performance of AI solutions over the time. Since the small scale utilities operate on a relatively smaller workforce, they also fear losing their job to AI technology.

Chapter 5

Discussion

The water sector is at the preliminary stages of adopting AI, as most of the utilities are not using any AI based solutions in their system. We found through RQ1 that there exist a digital divide amongst the utilities of different scales as there is no presence of AI in the small scale utilities while a few large scale utilities are already developing AI solutions. However, this does not mean that the challenges faced by any particular segment are any less significant than those of others. Even small-scale utilities have a responsibility to provide basic water needs to the communities they serve. Therefore, they need to be mindful of the responsible use of AI as a technology.

Our pilot interviews with the small scale utilities led us to find that they are not against implementing AI for their system. They have insights into what benefits and challenges AI provides for them and how AI based decision making can support their operations. Like their large-scale counterparts, the small-scale utilities are willing to upgrade to AI and use cutting edge technology for their system. As noted in RQ3, there are various challenges which they need to address for successfully integrating AI in their system. Since the small and medium scale utilities do not have AI experts in their team, they rely on consultants and AI vendors to integrate AI in their system. Although these utilities have the funds and resources to collaborate with external experts, the sustainability of AI operations within the utilities is a challenge. Consultants and vendors are not permanent employees of the utilities and can only help the small-scale utilities to setup and integrate AI in the system. The onus of sustaining and monitoring the performance and reliability of AI is on the utilities and they do not have the required experts who can keep a check on the AI system and maintain it over time. There is also a general feeling within the small utilities that AI will replace jobs that involve repetitive tasks such as data entry, customer service, and certain types of water utility services.

Another issue which is common for all utilities is the reliability of AI output. Since the utilities deal with real world water operations and provide basic water needs to people, they need to fully trust the system which they are using for services, as the stakes involved are high. Utilities haven't reached a stage where they can fully trust an AI technology instead of a human for making crucial decisions. This concern can also support the fact that it is difficult for the utilities to envision AI as a system engine (Figure 4.4 - RQ1).

Through RQ1, we found that there is a small amount of in-house expertise in the large scale utilities. Perhaps, this can be the reason behind why the large scale utilities turned out to be better than medium scale utilities in terms of data readiness (Figure 4.6 and Figure 4.7 shows large scale utilities have a more holistic approach in data quantity evaluation and preprocessing) and AI model experimentation (Figure 4.10 shows large scale utilities do not any bias or preference towards any AI model).

RQ3 led us to data acquisition as one of the major AI challenges for the utilities. At the same time, we found through RQ2 that none of the utilities are using techniques like sampling or synthetic data generation. Perhaps, if utilities start exploring synthetic data applications, it can aid in overcoming their data acquisition challenges.

One of our observations from RQ2 was that the utility representatives who communicated with us and responded to our questionnaire were civil engineering domain experts and were not designated in a AI specific role. These people may have knowledge about AI and might be leading the in-house AI operations. However, it is possible that they do not have advanced knowledge about AI techniques, which can be the reason behind why utilities follow basic data preprocessing (Figure 4.7) techniques and use simpler AI models over state of the art techniques.

We also found the presence of AI vendors in the water system, who helps the utilities achieve AI solutions. In fact, one of our respondent from the large scale utilities said that they are answering our questions based on the practice of past vendors who did pilot previously, as they don't have an in-house team. Reliance on vendors may justify why AI models are blackbox in nature for most of the utilities (Figure 4.14). Vendors may also account for the fact that many utilities are unaware about the importance and process of hyperparameter tuning, testing and timely model updation (Figure 4.12 and 4.13), as vendors usually work on contract bases and thus, do not end up maintaining the AI solutions for the utilities. The only response which said that the utility updates its AI model once a month and have basic knowledge about hyperparameters of AI models came from a large utility which uses in-house AI expertise. We also saw that utilities were not able to generalize their models. If they are relying on vendor based solutions, achieving generalization can be difficult as the vendors are not domain experts and may not customize their baseline solutions as per the utilities' needs.

Based on our discussions with medium scale utilities, we found that they are making efforts to build in-house AI expertise within their system. They are one step ahead of the small-scale utilities in terms of taking proactive measures towards developing an AI driven workforce for a digital transformation in the future. Currently, medium scale utilities have diversity and inclusion-driven workforce development programs for new hires to specialize in achieving the utility's day-to-day water operations. As the medium scale utilities prepare to enter the digital era of water, they plan to build a data and AI specialized workforce that can help them achieve in-house AI solutions for their utility. Medium scale utilities are also aware of the interdisciplinary requirements of integrating AI in their system and thus, they aim to bring together Social, Environmental, Legal and AI based domain experts in future to achieve a holistic and robust AI implementation within their utility.

Chapter 6

Future Work

The participating utilities in this research were all public utilities. In order to get a more generalized sense of current AI practices and challenges in the water sector, this study should also be extended to private and federal utilities. Since these utilities have different characteristics in terms of funding, governance, and scale of projects, piloting with these utilities can lead us to diverse inferences. Water utilities should approach the digital transformation as a multidisciplinary effort and should include computer, social and environmental, and legal domain experts to achieve a holistic and trustworthy AI implementation within their utility.

We saw that many utilities are not currently implementing AI and some are dependent on AI vendors for the same. The Strategic Implementation pillar of the *ai*WATERS framework can be extended to include a guide on how to develop in-house AI expertise for the utilities by conducting workshops for current employees and developing AI specific hiring plans and interview process. Moreover, governments and regulatory bodies could develop a regulatory framework for socially acceptable AI implementation in the water sector. This could help utilities navigate legal and ethical challenges related to AI and ensure that AI systems are transparent, fair, and accountable.

We got zero responses on the question which asked if the responding utility look up to some other utility in terms of implementing AI. It shows that utilities are not aware about how other utilities are progressing in their AI journey. Considering the digital divide, small and medium utilities can pilot and learn from the large utilities about their AI based projects and plans.

Chapter 7

Conclusions

In this study, we propose *ai*WATERS, an Artificial Intelligence framework designed to facilitate the successful application of AI within the water sector. Through an extensive literature and practice review, we identified several challenges associated with AI implementation: Understanding the technology, its benefits and challenges, Data Readiness, Knowledge Integration, Model Development, Decision Support, and Strategic planning in the water sector and proposed seven key pillars of AI revolving around these challenges. We conducted pilot interviews and administered questionnaire surveys with ten public utilities in the United States to assess their current level of AI implementation. By comparing the maturity of AI implementation between small, medium, and large utilities, we made several noteworthy findings. First, utilities in each category are at a different stage in their journey of digital transformation. Second, most of the utilities lack in-house AI expertise as they are not implementing AI or rely on AI vendors for integrating digital solutions into their systems. This reliance leads to the blackbox nature of AI for water utilities which in turn also raises trustworthiness and sustainability issues for the utilities. Thus, the *ai*WATERS framework is intended to provide comprehensive details for utilities to achieve holistic AI implementation in their system. This research concludes with recommendations for the future. It is expected that this study is extended to private and federal utilities as well to generalize our findings for the water sector.

Bibliography

[1] M. Mitchell, S. Wu, A. Zaldivar, P. Barnes, L. Vasserman, B. Hutchinson, E. Spitzer, I. D. Raji, and T. Gebru, "Model cards for model reporting," in *Proceedings of the conference on fairness, accountability, and transparency*, pp. 220–229, 2019.

[2] A. Hadjimichael, J. Comas, and L. Corominas, "Do artificial intelligence methods enhance the potential of decision support systems? a review for the urban water sector," *AI Communications. Published on*, vol. 1, 2016.

[3] V. B. Zoran Kapelan, Emma Weisbord, "Explained: Artificial intelligence solutions for the water sector." https://iwa-network.org/wp-content/uploads/2020/08/IWA_2020_Artificial_Intelligence_SCREEN.pdf. International Water Association (IWA).

[4] H. Jenny, E. G. Alonso, Y. Wang, and R. Minguez, "Using artificial intelligence for smart water management systems," 2020.

[5] A. Adadi and M. Berrada, "Peeking inside the black-box: a survey on explainable artificial intelligence (xai)," *IEEE access*, vol. 6, pp. 52138–52160, 2018.

[6] S. Park, J. Kim, A. Newland, B. Kim, and H. Jun, "Survival analysis of water distribution pipe failure data using the proportional hazards model," in *World Environmental and Water Resources Congress 2008: Ahupua'A*, pp. 1–10, 2008.

[7] B. Aslani, S. Mohebbi, and H. Axthelm, "Predictive analytics for water main breaks using spatiotemporal data," *Urban Water Journal*, vol. 18, no. 6, pp. 433–448, 2021.

[8] Z. Almheiri, M. Meguid, and T. Zayed, "Failure modeling of water distribution pipelines using meta-learning algorithms," *Water Research*, vol. 205, p. 117680, 2021.

[9] A. Shirzad, M. Tabesh, and R. Farmani, "A comparison between performance of support vector regression and artificial neural network in prediction of pipe burst rate in water distribution networks," *KSCE Journal of Civil Engineering*, vol. 18, pp. 941–948, 2014.

[10] G. Kabir, S. Tesfamariam, and R. Sadiq, "Bayesian model averaging for the prediction of water main failure for small to large canadian municipalities," *Canadian Journal of Civil Engineering*, vol. 43, no. 3, pp. 233–240, 2016.

[11] R. Farmani, K. Kakoudakis, K. Behzadian, and D. Butler, "Pipe failure prediction in water distribution systems considering static and dynamic factors," *Procedia Engineering*, vol. 186, pp. 117–126, 2017.

[12] R. Rayhana, Y. Jiao, A. Zaji, and Z. Liu, "Automated vision systems for condition assessment of sewer and water pipelines," *IEEE Transactions on Automation Science and Engineering*, vol. 18, no. 4, pp. 1861–1878, 2020.

[13] Y.-J. Cha, W. Choi, and O. Büyüköztürk, "Deep learning-based crack damage detection using convolutional neural networks," *Computer-Aided Civil and Infrastructure Engineering*, vol. 32, no. 5, pp. 361–378, 2017.

[14] S. S. Kumar, D. M. Abraham, M. R. Jahanshahi, T. Iseley, and J. Starr, "Automated defect classification in sewer closed circuit television inspections using deep convolutional neural networks," *Automation in Construction*, vol. 91, pp. 273–283, 2018.

[15] D. Meijer, L. Scholten, F. Clemens, and A. Knobbe, "A defect classification methodology for sewer image sets with convolutional neural networks," *Automation in Construction*, vol. 104, pp. 281–298, 2019.

[16] K. Kim, J. Kim, T.-Y. Kwak, and C.-K. Chung, "Logistic regression model for sinkhole susceptibility due to damaged sewer pipes," *Natural Hazards*, vol. 93, pp. 765–785, 2018.

[17] D. Manu and A. K. Thalla, "Artificial intelligence models for predicting the performance of biological wastewater treatment plant in the removal of kjeldahl nitrogen from wastewater," *Applied Water Science*, vol. 7, pp. 3783–3791, 2017.

[18] F. Cangialosi, E. Bruno, and G. De Santis, "Application of machine learning for fence-line monitoring of odor classes and concentrations at a wastewater treatment plant," *Sensors*, vol. 21, no. 14, p. 4716, 2021.

[19] F. Yang, T. B. Pluth, X. Fang, K. B. Francq, M. Jurjovec, and Y. Tang, "Advanced machine learning application for odor and corrosion control at a water resource recovery facility," *Water Environment Research*, vol. 93, no. 11, pp. 2346–2359, 2021.

[20] S. Ostojin, S. Mounce, and J. Boxall, "An artificial intelligence approach for optimizing pumping in sewer systems," *Journal of hydroinformatics*, vol. 13, no. 3, pp. 295–306, 2011.

[21] S. Yagi and S. Shiba, "Application of genetic algorithms and fuzzy control to a combined sewer pumping station," *Water Science and Technology*, vol. 39, no. 9, pp. 217–224, 1999.

[22] J. Zhang, Y. Yu, J. Yan, and J. Chen, "Data-driven parameter prediction of water pumping station," *Water*, vol. 15, no. 6, p. 1128, 2023.

[23] A. M. Moreno-Rodenas, A. Duinmeijer, and F. H. Clemens, "Deep-learning based monitoring of fog layer dynamics in wastewater pumping stations," *Water Research*, vol. 202, p. 117482, 2021.

[24] J. C. Berndtsson, "Green roof performance towards management of runoff water quantity and quality: A review," *Ecological engineering*, vol. 36, no. 4, pp. 351–360, 2010.

[25] S. Tsang and C. Y. Jim, “Applying artificial intelligence modeling to optimize green roof irrigation,” *Energy and Buildings*, vol. 127, pp. 360–369, 2016.

[26] E. M. H. Abdalla, V. Pons, V. Stovin, S. De-Ville, E. Fassman-Beck, K. Alfredsen, and T. M. Muthanna, “Evaluating different machine learning methods to simulate runoff from extensive green roofs,” *Hydrology and Earth System Sciences*, vol. 25, no. 11, pp. 5917–5935, 2021.

[27] F. Facchini, L. Ranieri, and M. Vitti, “A neural network model for decision-making with application in sewage sludge management,” *Applied Sciences*, vol. 11, no. 12, p. 5434, 2021.

[28] O. S. Djandja, P.-G. Duan, L.-X. Yin, Z.-C. Wang, and J. Duo, “A novel machine learning-based approach for prediction of nitrogen content in hydrochar from hydrothermal carbonization of sewage sludge,” *Energy*, vol. 232, p. 121010, 2021.

[29] L. Knoll, L. Breuer, and M. Bach, “Large scale prediction of groundwater nitrate concentrations from spatial data using machine learning,” *Science of the total environment*, vol. 668, pp. 1317–1327, 2019.

[30] S. A. Naghibi, M. Dolatkordestani, A. Rezaei, P. Amouzegari, M. T. Heravi, B. Kalantar, and B. Pradhan, “Application of rotation forest with decision trees as base classifier and a novel ensemble model in spatial modeling of groundwater potential,” *Environmental monitoring and assessment*, vol. 191, pp. 1–20, 2019.

[31] S. Miraki, S. H. Zanganeh, K. Chapi, V. P. Singh, A. Shirzadi, H. Shahabi, and B. T. Pham, “Mapping groundwater potential using a novel hybrid intelligence approach,” *Water resources management*, vol. 33, pp. 281–302, 2019.

[32] A. Mosavi, F. Sajedi Hosseini, B. Choubin, M. Goodarzi, A. A. Dineva, and

E. Rafiei Sardooi, "Ensemble boosting and bagging based machine learning models for groundwater potential prediction," *Water Resources Management*, vol. 35, pp. 23–37, 2021.

[33] C. Feng, M. Cui, B.-M. Hodge, and J. Zhang, "A data-driven multi-model methodology with deep feature selection for short-term wind forecasting," *Applied Energy*, vol. 190, pp. 1245–1257, 2017.

[34] M. Kuhn, "Caret: classification and regression training," *Astrophysics Source Code Library*, pp. ascl–1505, 2015.

[35] S. Talukdar, K. U. Eibek, S. Akhter, S. Ziaul, A. R. M. T. Islam, and J. Mallick, "Modeling fragmentation probability of land-use and land-cover using the bagging, random forest and random subspace in the teesta river basin, bangladesh," *Ecological indicators*, vol. 126, p. 107612, 2021.

[36] S. K. Sarkar, S. Talukdar, A. Rahman, and S. K. Roy, "Groundwater potentiality mapping using ensemble machine learning algorithms for sustainable groundwater management," *Frontiers in Engineering and Built Environment*, vol. 2, no. 1, pp. 43–54, 2022.

[37] H. G. Damavandi, R. Shah, D. Stampoulis, Y. Wei, D. Boscovic, J. Sabo, *et al.*, "Accurate prediction of streamflow using long short-term memory network: a case study in the brazos river basin in texas," *International Journal of Environmental Science and Development*, vol. 10, no. 10, pp. 294–300, 2019.

[38] "Cama-flood." http://hydro.iis.u-tokyo.ac.jp/~yamadai/cama-flood/#:~:text=Overall%20Framework-,The%20CaMa%2DFlood%20(Catchment%2Dbased%20Macro%2Dscale%20Floodplain,from%20a%20land%20surface%20model.. Online; Accessed: 3-April-2023.

[39] M. I. Shah, W. S. Alaloul, A. Alqahtani, A. Aldrees, M. A. Musarat, and M. F. Javed, "Predictive modeling approach for surface water quality: development and comparison of machine learning models," *Sustainability*, vol. 13, no. 14, p. 7515, 2021.

[40] S. Ai, A. Chakravorty, and C. Rong, "Household power demand prediction using evolutionary ensemble neural network pool with multiple network structures," *Sensors*, vol. 19, no. 3, p. 721, 2019.

[41] J. Gebhardt, H. Detmer, and A. L. Madsen, "Predicting parts demand in the automotive industry—an application of probabilistic graphical models," in *Proc. Int. Joint Conf. on Uncertainty in Artificial Intelligence (UAI'03, Acapulco, Mexico), Bayesian Modelling Applications Workshop*, 2003.

[42] U. Schmidt-Erfurth, S. M. Waldstein, S. Klimscha, A. Sadeghipour, X. Hu, B. S. Gerendas, A. Osborne, and H. Bogunović, "Prediction of individual disease conversion in early amd using artificial intelligence," *Investigative ophthalmology & visual science*, vol. 59, no. 8, pp. 3199–3208, 2018.

[43] M. Herrera, L. Torgo, J. Izquierdo, and R. Pérez-García, "Predictive models for forecasting hourly urban water demand," *Journal of hydrology*, vol. 387, no. 1-2, pp. 141–150, 2010.

[44] G. Guo, S. Liu, Y. Wu, J. Li, R. Zhou, and X. Zhu, "Short-term water demand forecast based on deep learning method," *Journal of Water Resources Planning and Management*, vol. 144, no. 12, p. 04018076, 2018.

[45] C. Kühnert, N. M. Gonuguntla, H. Krieg, D. Nowak, and J. A. Thomas, "Application of lstm networks for water demand prediction in optimal pump control," *Water*, vol. 13, no. 5, p. 644, 2021.

[46] B. M. Brentan, E. Luvizotto Jr, M. Herrera, J. Izquierdo, and R. Pérez-García, "Hybrid regression model for near real-time urban water demand forecasting," *Journal of Computational and Applied Mathematics*, vol. 309, pp. 532–541, 2017.

[47] S. Zhang, H. Wang, and A. A. Keller, "Novel machine learning-based energy consumption model of wastewater treatment plants," *ACS ES&T Water*, vol. 1, no. 12, pp. 2531–2540, 2021.

[48] D. Torregrossa, U. Leopold, F. Hernández-Sancho, and J. Hansen, "Machine learning for energy cost modelling in wastewater treatment plants," *Journal of environmental management*, vol. 223, pp. 1061–1067, 2018.

[49] "Utilizing soft sensor system for process control and optimization." https://www.waterrf.org/sites/default/files/file/2021-10/View%20Solution%201_2021%20IWS%20Challenge%20Solution_CWS_Lu.pdf. Online; Accessed: 3-April-2023.

[50] M. Sit, B. Z. Demiray, Z. Xiang, G. J. Ewing, Y. Sermet, and I. Demir, "A comprehensive review of deep learning applications in hydrology and water resources," *Water Science and Technology*, vol. 82, no. 12, pp. 2635–2670, 2020.

[51] C. Burghard, "The data dilemma and its impact on ai in healthcare and life sciences." https://blogs.idc.com/2021/06/23/the-data-dilemma-and-its-impact-on-ai-in-healthcare-and-life-sciences/. International Data Corporation (IDC).

[52] K. P. Murphy, *Machine learning: a probabilistic perspective.* MIT press, 2012.

[53] C. Huyen, *Designing Machine Learning Systems.* " O'Reilly Media, Inc.", 2022.

[54] S. Richmond, "Superintelligence: Paths, dangers, strategies. by nick bostrom. ox-

ford university press, oxford, 2014, pp. xvi+ 328. hardcover: 29.95/$18.99.*isbn* : 9780199678112," *Philosophy*, *vol*. 91, *no*. 1, *pp*. 125 − −130, 2016.

[55] A. Kossiakoff, S. M. Biemer, S. J. Seymour, and D. A. Flanigan, *Systems engineering principles and practice*. John Wiley & Sons, 2020.

[56] R. S. Geiger, K. Yu, Y. Yang, M. Dai, J. Qiu, R. Tang, and J. Huang, "Garbage in, garbage out? do machine learning application papers in social computing report where human-labeled training data comes from?," in *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, pp. 325–336, 2020.

[57] N. Sambasivan, S. Kapania, H. Highfill, D. Akrong, P. Paritosh, and L. M. Aroyo, "“everyone wants to do the model work, not the data work”: Data cascades in high-stakes ai," in *proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, pp. 1–15, 2021.

[58] "Crowdflower data science report." https://visit.figure-eight.com/rs/416-ZBE-142/images/CrowdFlower_DataScienceReport_2016.pdf. Online; Accessed: 27-Aug-2020.

[59] I. H. Witten and E. Frank, "Data mining: practical machine learning tools and techniques with java implementations," *Acm Sigmod Record*, vol. 31, no. 1, pp. 76–77, 2002.

[60] J. Warren and N. Marz, *Big Data: Principles and best practices of scalable realtime data systems*. Simon and Schuster, 2015.

[61] F. J. Ohlhorst, *Big data analytics: turning big data into big money*, vol. 65. John Wiley & Sons, 2012.

[62] M. Krenn, R. Pollice, S. Y. Guo, M. Aldeghi, A. Cervera-Lierta, P. Friederich, G. dos Passos Gomes, F. Häse, A. Jinich, A. Nigam, *et al.*, "On scientific understanding with artificial intelligence," *Nature Reviews Physics*, vol. 4, no. 12, pp. 761–769, 2022.

[63] P.-N. Tan, M. Steinbach, and V. Kumar, *Introduction to data mining.* Pearson Education India, 2016.

[64] I. Goodfellow, Y. Bengio, and A. Courville, *Deep learning.* MIT press, 2016.

[65] D. Kaur, S. Uslu, K. J. Rittichier, and A. Durresi, “Trustworthy artificial intelligence: a review,” *ACM Computing Surveys (CSUR)*, vol. 55, no. 2, pp. 1–38, 2022.

[66] K. Eberhard, “The effects of visualization on judgment and decision-making: a systematic literature review,” *Management Review Quarterly*, vol. 73, no. 1, pp. 167–214, 2023.