ARTIFICIAL INTELLIGENCE IN THE WATER INDUSTRY: MYTH OR REALITY?

Mo Chalabi, CH2M Beca

ABSTRACT

Water suppliers are constantly seeking techniques to improve the quality of their services and reduce operational costs. This is traditionally done through ensuring that the water infrastructure is maintained regularly by performing routine maintenance and responding to faults within the infrastructure. Operators would usually discover a fault, analyse the data and respond accordingly, in a "reactive" manner. Although these techniques work, they mainly rely on human intervention which can sometimes be inefficient, slow and potentially costly. In addition, water suppliers generally follow a reactive approach to energy consumption, where potential savings are lost due to not taking external factors into account.

Artificial intelligence (AI) systems utilise real-time and historic data to optimise operations in terms of cost and quality, potentially making the water industry more "proactive". Current AI systems can significantly reduce the response time to unusual events or faults, predict faults and learn how to respond in future occurrences. These systems can also produce maintenance schedules and assign their priority based on predicted outcomes.

Another major benefit is the ability to optimise operations to save on energy costs based on various sources of data such as weather reports and electricity suppliers' charges, creating the ability to estimate the optimal balance between reducing energy consumption and maintaining the required storage volumes for peak demand. Such systems would significantly reduce the cost of operations whilst ensuring a sufficient reserve.

This paper will investigate existing applications of AI in the water industry, specifically those that are scalable, practical and require minimum change to existing systems. In addition, real-world examples of AI implementation in the water industry will be discussed with the emphasis on lessons learnt. Finally, this paper will showcase platforms data collected by SCADA and historian systems to be converted into "smart data", enabling water suppliers to utilise the latest AI technology at relatively small costs.

KEYWORDS

Artificial intelligence (AI), operational costs, energy saving, predictive maintenance, SCADA.

PRESENTER PROFILE

Mo Chalabi (Lead Author)

After graduating from the University of Canterbury with a Bachelor of Mechatronics Engineering, Mo has been working in Beca's industrial team. Mo specialises in emerging technology solutions, 360 imaging, SCADA and communications design and software development. In 2016, Mo has been awarded IPENZ Ray Meyer Medal for Excellence in Student Design as part of a successful team project. Recently, Mo was part of Beca's runner up team at the Internet of Things Cup in Melbourne.

Andrew Duffy (Support Author & Mentor)

Andrew is a chartered engineer with 15 years' experience in electrical, instrumentation and controls system engineering. He specialises in the Water and Wastewater Industry from small pump stations to large treatment plants. He has worked in various industries in New Zealand, Australia and the UK.

1 INTRODUCTION

Water utilities are continuously faced with growing demand for water resources as the world population grows. In addition, they have to overcome aging infrastructures and lack of funding for operations and maintenance while maintaining the quality of their service (Corporation, 2015).

Although the majority of the existing water infrastructure is monitored in real-time, producing large amounts of data, it still requires human intervention to understand and analyse the acquired data. Relying on operators' intuition and experience to process, analyse and forecast operational changes could lead to errors and delays, which can result in damaged equipment, lower efficiency and eventually compromised quality of service.

When it comes to maintenance, water utilities mostly rely on either reactive or preventive strategies. Reactive maintenance is only performed after equipment has failed, and therefore it is classified as failure-based strategy. On the other hand, preventative maintenance is performed regularly based on time or usage of an equipment even if the equipment is not close to failure.

Water utilities often only rely on internal acquired data for operation and maintenance purposes, leaving a huge amount of data sources that can improve the decision making process. External data sources such as weather, water consumption forecasts based on population, energy rates and infrastructure health would optimise both the operations and maintenance of water treatment plants. But as aforementioned, such massive amount of data cannot possibly be analysed by humans which creates the need for other systems such as artificial intelligence.

Artificial intelligence (AI) is defined as the mimicking of humans cognitive functions by a machine, such as problem solving, learning and reasoning. A subcategory of AI is machine learning (ML), which is a field of computer science that allows machines to learn to perform certain tasks without being explicitly programmed (Munoz, n.d.). There are different levels of AI, some of which will be discussed.

This paper will investigate existing applications of AI in the water industry, specifically those that are scalable, practical and require no change to existing systems. In addition, real-world examples of AI implementation in the water industry will be discussed with the emphasis on lessons learnt. Finally, this paper will showcase the process of converting data collected by SCADA and historian systems into "smart data", enabling water suppliers to utilise the latest AI technology at relatively small costs.

2 DISCUSSION

2.1 WHAT DOES ARTIFICIAL INTELLIGENCE MEAN?

Although scientists have been contemplating the thought of machines being able to think on their own for a long time, the term "Artificial Intelligence" was officially introduced in early 1950's, when the first conference on AI was held by John McCarthy (Smith, n.d.). AI can be broken intro two segments, Artificial General Intelligence (AGI), and Artificial Narrow Intelligence (ANI). AGI ultimate goal is to create a machine to imitate humans' actions sufficiently that a judge would not be able to tell the difference between an actual human and the machine, also known as the Turing test (Smith, n.d.). For various reasons such as the lack of computational hardware and funding allocated to AI, no machine has successfully passed the Turing test yet.

Artificial Narrow Intelligence (ANI) is defined as the ability to perform certain pre-defined tasks efficiently while learning and adapting. Examples of ANI include learning algorithms that are being used around the globe in applications such as Apple's intelligent assistant Siri and Tesla's self-driving cars. Self-driving cars have been a far-fetched reality for a few decades; however, companies such as Google and Tesla are making that a reality with almost 3 million miles being self-driven (Waymo, 2017). Driving in itself is a complicated act, for a self-driving car to function safely it would have to sense the surrounding, deal with unexpected encounters, understand and analyse human intervention, and offer reliability and security. These functions are comparable to the operation of a wastewater treatment plants, as shown in Table 1. This goes to show that AI techniques can be applied to a wide range of industries and applications.

Table 1: Comparison between self-driving car's functionality and a waste water treatment plant.

Situation	Self-Driving Car	Waste Water Treatment Plant
Sensing of surroundings	Bikers, road marking, other cars	Raw wastewater intake is measured for Flow, DO, Nitrate, pH, Ammonia and COD.
Dealing with unexpected encounters	Sudden road situations, pedestrians, bikers and cars suddenly appearing in front of the car	Unexpected dumping of waste, floods and storm water
Understanding and acting upon human intervention	Drivers' intervention	Operators' intervention
Offering reliability and security	Secure against cyber-hacks, offers redundancy	Secure against hackers, vandals, human errors

In the 1980s, business applications of AI were trialled, but did not succeed. But due to the reduction of hardware prices, the increase of computing power, and the vast amount of learning data available nowadays, AI has become business-ready. In addition, the rise of AI platforms has made it easy for business to utilise the power of AI without bearing the costs of hiring specialist staff. There has also been an increase in awareness of current AI research and its capabilities across industries. Another major motivator for AI implementation is the rise of open-source, web-based AI platforms such as Google Tensorflow, Amazon Artificial Intelligence Services and IBM Watson (Economist, 2016). Such platforms allow developers to create their instance of an AI algorithm, train the algorithm with relevant data and then eventually create their own applications.

As aforementioned, AI's advantage lies in its ability to analyse, learn and infer from large volumes of data. Therefore, it is essential access to large amounts of data sets for any AI system to function effectively. This data could come from local sources such as supervisory control and data acquisition (SCADA) and historian, or from external sources such as the internet. For a waste water treatment plant, a relatively large amount of data is being collected by instruments from both influent and effluent streams. For the last three decades, SCADA systems have been implemented in water and wastewater plants to facilitate an interface between plant operators and the machinery. Such systems acquire data from across the plant and the network, creating and storing large data sets of time-stamped readings. Although this data has the potential to support improving operations and decision making process for those plants, it rarely gets used. This is due to the inherent design of historian infrastructure which has been optimised for write access with limited search or analytics functionality. But such systems provide rich environment for AI, where AI applications can be used as an add-on to extract the data from historians, process it and make sense of it. In addition, many AI platforms take into consideration external factors such as weather condition, energy tariff as well as any other 3rd party relevant information. An ideal model for an AI platform implemented within a waste water treatment plant is shown in Figure 1. Two major aspects that AI would optimise have been identified to be operational costs savings and predictive maintenance.

Figure 1: A schematic diagram of an AI platform implementation within a typical waste water treatment plant. (Diagram based on Baxter, 2001)

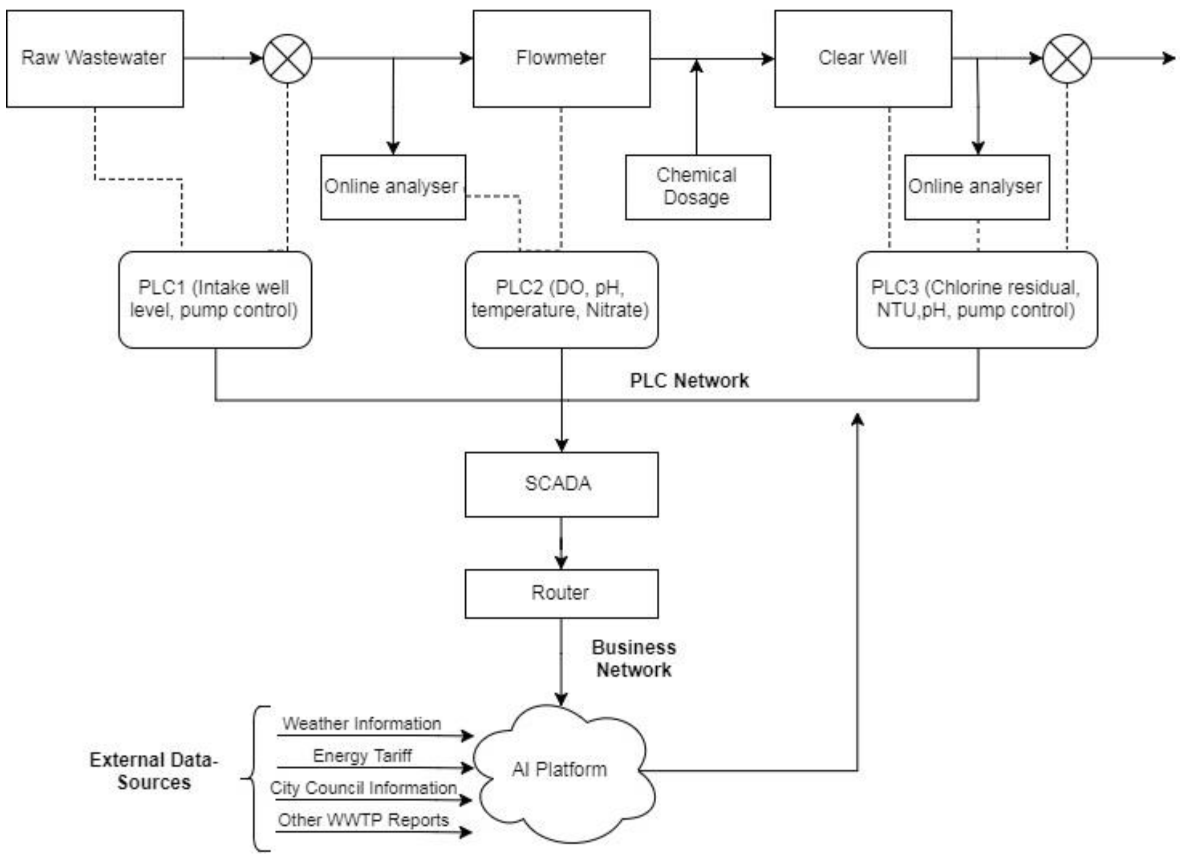

2.2 OPTIMISING OPERATIONS

In a survey done by the Economist, 49% of business leaders believe that AI will have moderate impact on their organisation in the next five years, while 20% expect little or no impact at all (Economist, 2016). There are many reasons why managers and executives, especially non-technical, might perceive AI as a concept rather than a tool. Building a viable business case requires understanding of the technology and its applications and limitations is a challenging task. In addition, the technology itself has not matured yet which makes it hard to invest large amounts of capital especially when there are no standards or guidelines around AI functionality. Therefore, implementing AI into businesses must be done on smaller scale at first, with clear and measurable milestones and on applications that allow for immediate optimisation. Two of the most suited applications of AI in the water industry are energy consumption and predictive maintenance.

2.2.1 AI APPLICATIONS IN POWER CONSUMPTION

In the water industry, energy is required throughout the processes of water production and distribution. Energy was not regarded as a place for operational cost savings; however, with the rise of energy prices, it became a main concern to water utilities regardless of their size or location. To put this in perspective, water utilities consume around 3% of all electrical energy consumption in the United States and United Kingdom. Pumping water takes between 90% and 95% of the electricity purchased (Bunn & Reynolds, 2009). The slightest improvement in pump efficiency will still achieve a huge amount of savings and reduce carbon emissions.

Pump reconditioning can be very effective in reducing operational costs, but operators and managers often lack the supporting evidence to back instituting preventative maintenance programs to optimise pumps efficiencies. However, there are companies such as AEMS Ltd. that specialise in pump energy management. In a study carried out by AEMS, a pump which has been in operation since 1963, with over 100,000 hours without any major maintenance had its efficiency raise from 70% to 82%. The process cost $20,000, but it saved $26,000 of operational cost per year, giving a 9 months' pay back period (Bunn & Reynolds, 2009).

Previous work has been done in terms of analysing individual pumps and matching pump characteristics to specific requirements for the duties of the pumps. The fundamental flaw behind those approaches is assuming that each pump runs at a single pressure and flow operating point. As water treatment plants become more complex and interconnected, pumps seldom operate in isolation. The use of variable speed drives (or inverters) is also becoming more common. So in order to capture all those interdependencies, a dynamic system is required to capture the performance of all the pumps and schedule their operations to optimise energy savings.

There are few methods that have been developed to understand and coordinate pump scheduling in order to optimise their operations. The majority of such system leverage the time-of-use tariff technique to minimise overall energy costs related to pumping. Water suppliers have the ability to offset the bulk of the pumping to off-peak hours where energy cost is lower. Existing solution are mostly based around reducing the associated cost of pumping for a small set of pumps. Additional techniques such as reducing the frequency of pump start and stop to minimise pressure surges or optimise the treatment process to avoid peak-hour charges exist, but have not been implemented widely.

A platform that has been successful in using machine learning is Aquadapt. This platform collects and stores existing plant data along with virtually generated data based on a calibrated model to build large data sets of mixed real and virtual data. This data is then used by an optimisation algorithm to highlight useful operational information such as inefficient pumps. This information is hugely beneficial for maintenance and replacement optimisation. But the platform does not only rely on operators’ response, it encompasses an automatic pump scheduler that minimise energy costs while taking into account the most suitable pumps to operate, time of day and predicted load. Implementing the Aquadapt platform has achieved energy savings of 10-20% for East Bay MUD and Eastern Municipal Water District in California, Washington Suburban in Maryland and WaterOne in Kansas (Bunn & Reynolds, 2009).

2.2.2 PREDICTIVE MAINTENANCE

There are two well-known strategies for maintenance, reactive and preventative. Reactive maintenance has the advantage of lower short-term savings, but it can lead to higher cost of repair and longer down-time in the long run. On the other hand, preventative maintenance achieves higher reliability but it has the disadvantage of being more costly if maintenance is performed without being needed (Levitt, 2011). A strategy that would achieve both reliability and reduced costs is preventative maintenance. Preventative maintenance is defined as condition-based approach where machines are only maintained when they actually require it prior to failure, as shown in Figure 2. It also enables operators to plan for failures ahead of time, therefore reducing down-time and its associated cost.

Figure 2: A diagram explaining predictive maintenance and how it compares to regular maintenance. (Diagram based on Levitt, 2011)

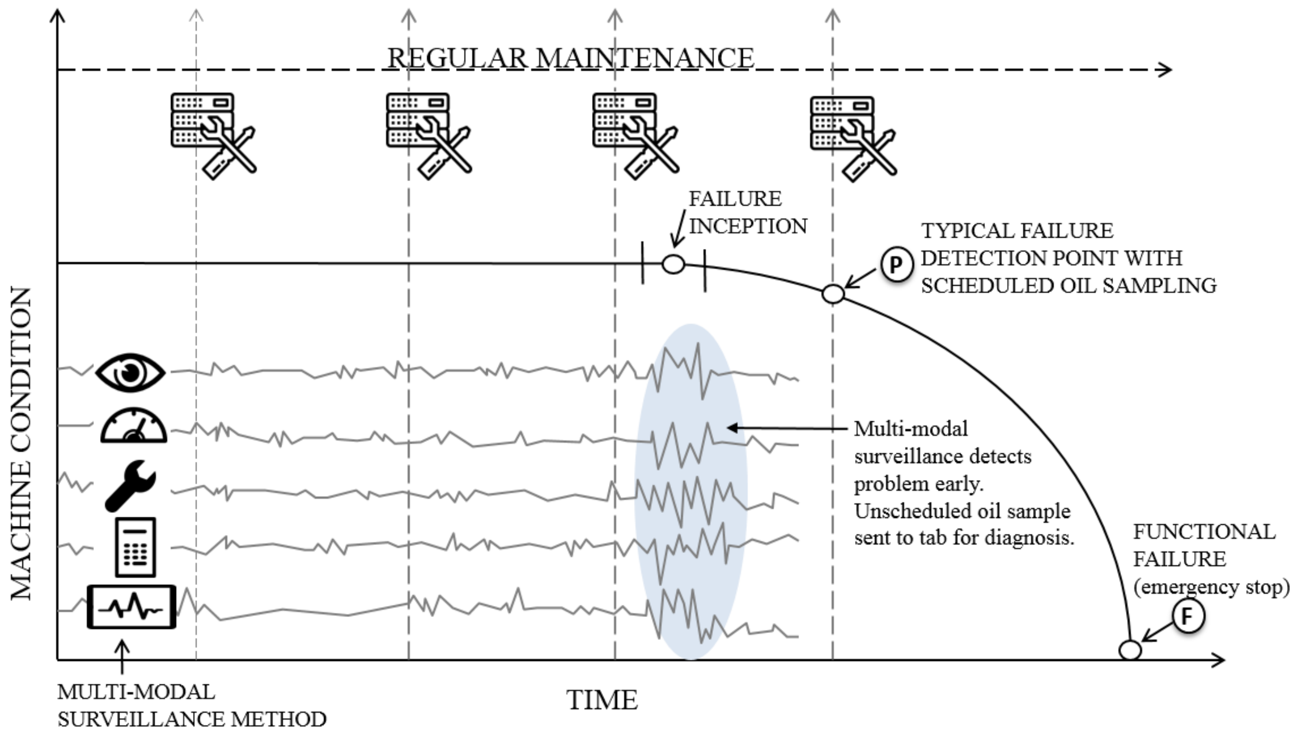

In order to efficiently implement predictive maintenance and anticipate equipment failure, a large amount of historical data sets is required along with equipment logs. Equipment logs typically contain thousands of events entries based on their pre-set time resolution. These events contain error codes, time of occurrence, event category and unstructured text messages. Typically an experienced operator would manually scan the data and find patterns and anomalies. This process is both time consuming and inefficient, as it would take a significant amount of time to go through the large amount of data being collected for one machine.

Advanced predictive maintenance platforms utilise already existing data for predicting equipment failures without the need for human intervention. Such platforms are typically divided into three segments, data acquisition, analysis and knowledge management and maintenance dashboard. The data acquisition segment deals with extracting data from existing SCADA and historian systems and processing it. The second segment is dedicated to modelling data and comparing equipment state with previously recorded states to detect any patterns that may cause failure. The third segment focuses on displaying meaningful data to operators through dashboards. An example of a predictive maintenance packages is IBM's Predictive Maintenance and Quality tool (PMQ), which is based on IBM's AI platform IBM Watson. PMQ helps organizations predict the timing and reason behind failures, as well as identifying poorly performing assets (IBM, n.d.). The PMQ tool provides operators with a simple to use dashboard that contains a health score for each machine, which makes it simple to identify poorly performing equipment and maintain them. The District of Columbia Water and Sewer Authority (DC Water) recently implemented IBM PMQ to optimise the maintenance process of their aging infrastructure and increase the reliability of their system. The platform offered near real-time information on potential issues and occurrences, based on a variety of information such as location, time, weather data and historical data. The results were promising with 36% reduction in customer complaints due to less asset downtime, as well as almost doubling the number of emergency investigation dispatched with 10 minutes. In addition, regulatory compliance reports were generated by the AI platform in seconds instead of days (IBM, 2010).

2.3 THE MOVE TO AI

2.3.1 AI PLATFORMS

There are a large number of AI platforms that are designed to take unstructured historical data as input and aid with optimising the performance of water and wastewater treatment plants. Table 2 highlights some of those platforms. It is important to note that slight differences between platforms could affect the result of implementation, as different platforms are targeted at certain fields.

Table 2: Comparison between different AI platforms for different water and wastewater applications.

AI Application	**Platforms**
Predictive Maintenance	IBM PMQ, SAP, Cisco, Intel, Siemens, Microsoft, GE, ABB, Huawie
Reducing Energy Costs & Consumption	Aquadapt, Pluto AI
Forecasting Intake, finding leaks underground	Pluto AI, H2O

In order to choose a suitable platform, it must satisfy all if not the majority of the following requirements:

- A non-intrusive data acquisition system. Allows information gathering from existing SCADA and historian systems without alternating them.

- A decision support system to provide real-time management recommendations.
- Include an interoperable hydraulic modelling capability to simulate the network behaviour, assess recommendations and analyse measured and simulated information.
- An events detection system to detect patterns, anomalies and to alert operators and managers about events whilst recommending solutions in real-time.
- Simple to use dashboards for KPI tracking. The dashboard would ideally include water resource, physical, operational or economic indicators.

3 RECOMMENDATIONS

The following points could aid with achieving a successful outcome when implementing AI into water and wastewater treatment plants:

- Identify an area within your treatment plant that requires immediate optimisation.
- It is very important to understand that no optimisation can be done without monitoring, so make sure your existing infrastructure is monitored sufficiently.
- Set a reasonable return on investment and other relevant KPIs to measure the success of implementation.
- Utilise existing SCADA and historian data and run a trial leveraging free AI platforms.
- Perform a technology assessment to identify the most suited platform. This could be done through partnership with vendor agnostic consultants.
- Run a small scale pilot study (bearing in mind the fail fast technique).
- Partnership with the chosen vendor for the pilot study.
- If successful, consider expanding the same platform to service other areas within your plant.
- Explore the use of Internet of Things (IoT) sensors in providing more data into the system.

4 CONCLUSIONS

Artificial intelligence has many use cases in almost every industry nowadays. Most of the existing water infrastructure is monitored, and the collected data often goes into a SCADA or historian system where it does not get used. Artificial intelligence has big potential at making sense of the acquired data, turning it into informative tools for operators and decision makers. This paper defined artificial intelligence, investigated existing applications of AI in the water industry, specifically reducing power costs and predictive maintenance. In addition, real-world examples of AI implementation in the water industry were discussed with the emphasis on the platforms used and lessons learnt. Finally, this paper will showcase the process of converting data collected by SCADA and historian systems into "smart data", enabling water suppliers to utilise the latest AI technology at relatively small costs.

ACKNOWLEDGEMENTS

Thank you to all Beca staff who volunteered their time to provide me with content for this paper, especially my mentor Andrew Duffy who provided support and guidance along the way.

REFERENCES

Bunn, S. & Reynolds, L., 2009. *The energy-efficiency benefits of pump-scheduling optimization for potable water supplies.,* s.l.: BM Journal of Research and Development.

Corporation, I., 2015. *IBM Intelligent Water.* New York: IBM Corporation .

Economist, T., 2016. *Artificial Intelligence in the Real World (The business case takes shape).* s.l.: The Economist Intelligence Unit Limited .

IBM, 2010. *DC Water creates a smarter utility service with IBM,* New York: IBM Corporation.

IBM, n.d. *Predictive Maintenance.* [Online]
Available at: https://www-01.ibm.com/software/analytics/solutions/operational-analytics/predictive-maintenance/products.html
[Accessed 1 July 2017].

Levitt, J., 2011. Predective Maintenance. In: *Complete Guide to Preventive and Predictive Maintenance .* s.l.:Industrial Press Inc., pp. 145-154.

Mounce, S., 2013. *A comparative study of artificial neural network architectures for time series prediction of water distribution system flow data.* Exter, The Society for the Study of Artificial Intelligence and the Simulation of Behaviour.

Munoz, A., n.d. *Machine Learning and Optimization.* s.l., s.n.

Smith, C., n.d. *The History of Artificial Intelligence.* Seattle: University of Washington.

Waymo, 2017. *On the Road.* [Online]
Available at: https://waymo.com/ontheroad/
[Accessed 5 July 2017].