Case Study

AI Coagulant Dosing Predictive Model

Figure 1. *Objective Water Treatment Plant*

Task Objectives

CWT was recently engaged to develop an artificial intelligence model to predict and optimise coagulant dosing at an NSW water treatment plant. Utilising historical water quality and dosing data, the model aimed to enhance dosing accuracy, reduce chemical usage and support operator decision making on-site. Beyond high predictive performance, CWT focused on embedding the model into the operator's existing Excel logbook through a serverless cloud architecture to ensure seamless adoption.

Methodology

Model construction involved four main steps:

1. **Preparing the Data for Training** – Water quality and dosing data was analysed by CWT to identify any anomalous events. Ensuring quality training data was crucial in optimising predictive validity.

2. **Model Training** – Cleaned data was trained using the **Random Forest Regression Algorithm (RFA).** RFA results were measured through **R^2, RMSE and feature importance,** using CWT's in-house AI Training Program.

3. **Validation** – Model predictions were compared to historical dosing. Through residual analysis, CWT quantified discrepancies to confirm predictive validity.

4. **Deployment** – The model was hosted through the **Microsoft Azure** platform and deployed onto the plant's Excel workbook for the onsite operator.

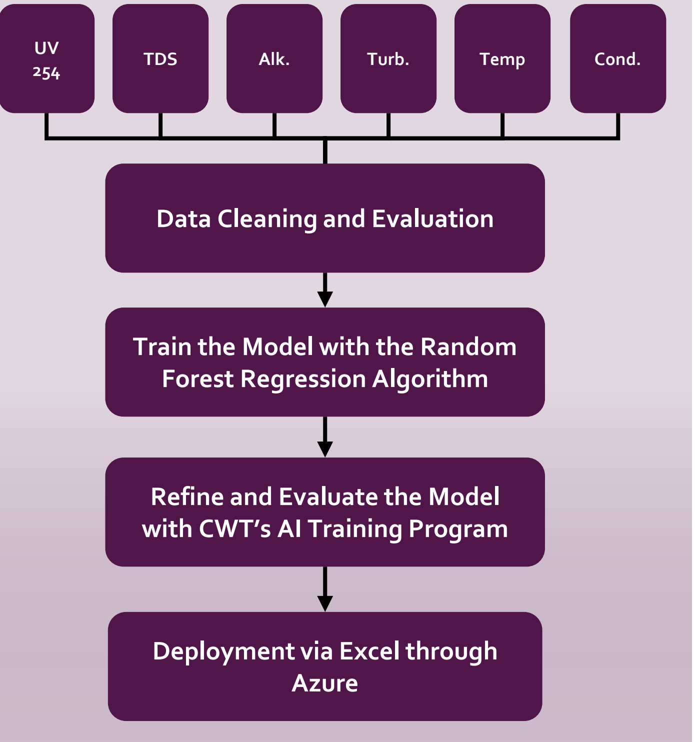

Figure 2. Model Workflow

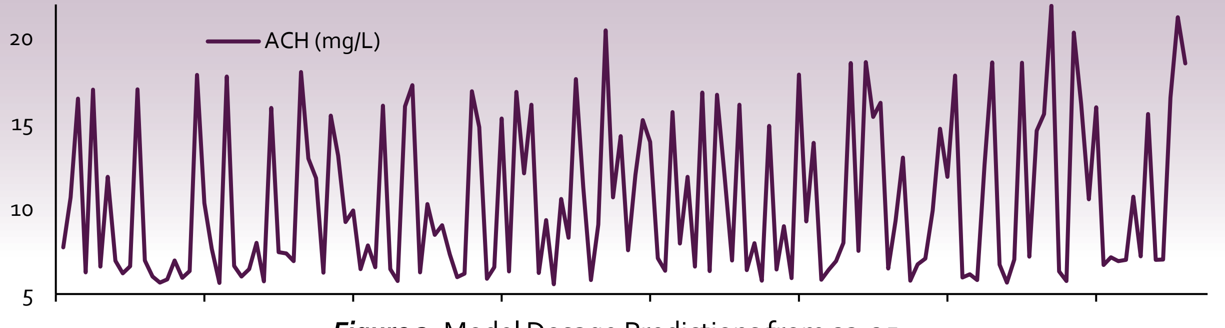

Figure 3. Model Dosage Predictions from 23-25

www.citywater.com.au
contact@citywater.com.au

AI Coagulant Dosing Predictive Model

Model Training and Evaluation

The model was evaluated through **R^2, RMSE and feature importance.** These metrics measured the model's ability to explain dosing variances, indicate average prediction error in mg/L and identify the relative influence each feature had on predictions.

R^2	RMSE (mg/L ACH)
0.8830	1.59

Table 1. R^2 and RMSE of the Plant Model

The results of the RFA showed an **R^2 of 0.883**, meaning the model explained approximately **88.3% of the variance** in historical coagulant dosing.

The **RMSE** indicated that on average, predictions deviated by about **1.59 mg/L**. Given operational dose rates typically fall in the **5–20 mg/L** range, this error was considered relatively small.

UV_{254} absorbance contributed over **60%** of the model's prediction. This is consistent with treatment expectations as it reflects primary drivers of coagulant demand, like natural organic matter.

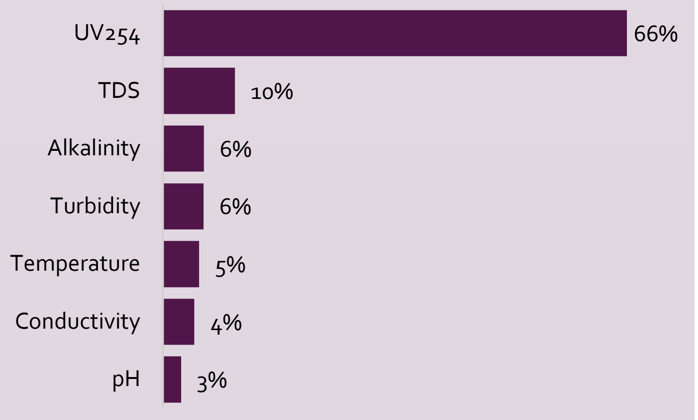

Figure 4. Relative Importance of Raw Water Quality Parameters

Model Validation - Residual Analysis

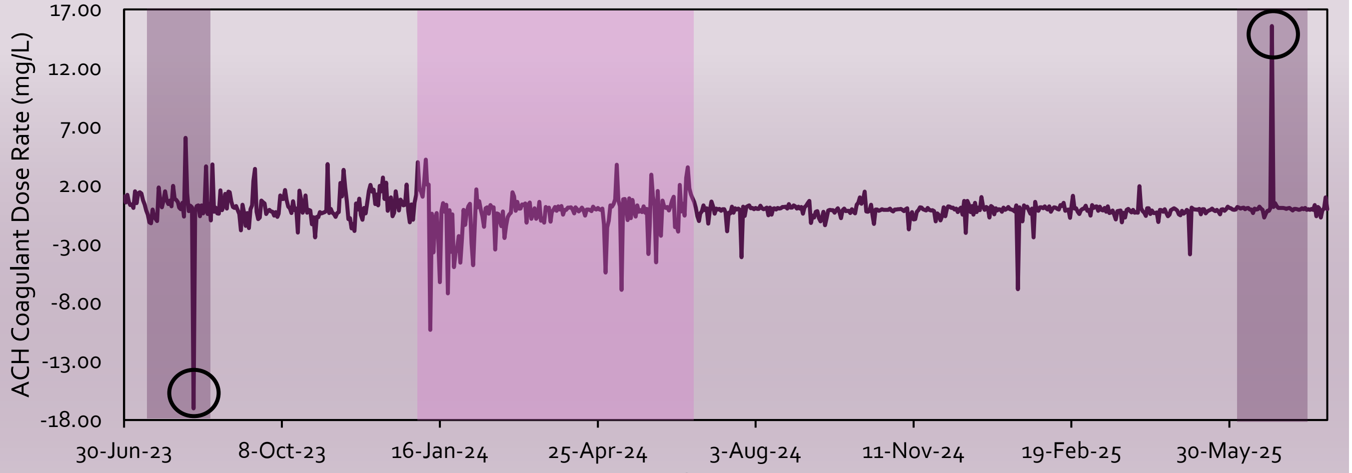

Figure 5. Dosing Residuals from June 2023 – August 2025

Residuals are the difference in predicted and historical dose rates of the plant. The model demonstrates strong performance across the dataset, with most residuals falling within operationally acceptable limits **(5–20 mg/L ACH)** and no consistent directional bias, supporting its robustness for prediction.

While small residuals indicate good model accuracy, large residuals can also be valuable by flagging potential anomalies where model predictions diverge from recorded data due to logging errors or unlogged interventions. In this case, spikes up to ±17 mg/L aligned with known inconsistencies in historical dosing records, as highlighted in Figure 5.

Case Study

WATER TREATMENT SPECIALISTS

AI Coagulant Dosing Predictive Model

Model Deployment and Cloud Hosting

The model is hosted through a **Microsoft Azure Serverless Inference Architecture (SIA)**. In this setup, computation only occurs when a request is made, which is an extremely low-cost hosting option and is well-suited given the model's high idle time between operator inputs. SIA's negligible yearly cost provides scalability for plants requiring multiple prediction models. Furthermore, using SIA supports automatic implementation of retraining and version control to ensure operators have perpetual access to the latest model.

Hosting Method	Yearly Cost (AUD)	Notes
Serverless Interface	**Negligible**	Charged only when the model is used. Cold start of less than three seconds.
Virtual Machine	~$1300-$1620 Per Model	Charged continuously even when the model is not in use.

Table 2. Breakdown of Hosting Methods

Deployment On-Site

In this deployment, the prediction model is embedded into the operator's daily Excel logbook via Power Query. When operators update their water quality parameters, Excel automatically sends the data through an API call to the SIA. The model then returns a real-time dosing prediction, displayed directly within the workbook. This approach allows operators to integrate AI data analytics into their on-site decision making, without deviating from their current working environment.

Date	
31-Jul-25	<--- **STEP 1**: Enter Date for prediction
Predicted Coag Dose	
5.77	<--- **STEP 2**: Press Ctrl + Alt + F5 to update the prediction based on the chosen date

Date	Raw pH	Raw Total Dissolved Solids	Raw Temperature	Raw Alkalinity	Raw Turbidity	Raw Conductivity	Raw UV254 nm
29-Jul-25	7.29	104	12.2	38	1.93	154	**0.0400**
30-Jul-25	7.33	104	11.7	32	1.75	154	**0.0460**
31-Jul-25	7.45	104	12.5	59	1.26	153	**0.0520**

Figure 6. Excel Interface Displaying Real-Time Dosing Predictions from the AI Model, Embedded within the Operator's Daily Logbook

Next Steps

Having achieved success in data preparation, model training, validation and deployment, CWT will be undertaking the following next steps:

Expanding Plant Development – CWT is currently working on extending the model training methodology towards a second, 2 ML/day water treatment plant in far west NSW.
On-Site Model Validation – CWT will provide support and training for operational use of the model.
Broader Prediction Parameters – CWT is investigating utilising the model to predict key parameters for application in other processes.
Integration with SCADA – CWT will explore technical requirements to link the cloud-based model to the plant's SCADA systems to integrate AI predictive data analytics into process control systems.
Deployment and Hosting Expansion – As prediction requirements shift, CWT is exploring model deployment into client-facing platforms like Power BI, and lightweight web dashboards. Deployment flexibility allow predictions to be accessed directly within existing reporting or monitoring tools.