Artificial Intelligence for Optimising Water Management and Control in Critical Infrastructures

Development and Policy Assessment of AI-based Solutions for Water Treatment Processes in European Countries

Federico Sacile

Delft University of Technology

Artificial Intelligence for Optimising Water Management and Control in Critical Infrastructures

Development and Policy Assessment of AI-based Solutions for Water Treatment Processes in European Countries

by

Federico Sacile

6084443

Internship Supervisor |	S. Karakach
First Supervisor |	Dr. P. Stokkink
Second Supervisor |	Dr. J. Zatarain Salazar
Chair |	Dr. L. Scholten
Project Date |	November, 2025
MSc Program |	Engineering & Policy Analysis
Faculty |	Technology, Policy and Management

Acknowledgements

Working on this thesis - and this period of my life - has represented great challenges. I had to face a lot of changes all at once, such as adapting to a new city, a new language, and a new job. Yet I have always welcomed challenges, and I think that this experience has been highly valuable for my growth and development, both as a student and as a young professional.

However, it is fundamental to remember that I would have never made it alone, and I must be extremely thankful to everyone who supported me in this process.

First of all, I would like to thank my graduation committee, without whom this thesis would not have been possible. Especially, I would like to thank Lisa for making me shift trajectory and helping me turn this project into a real EPA master's thesis; Patrick, for always being available and helpful whenever I needed; and Jazmin for her precious suggestions and interview contacts, which provided great value to the policy part of this work.

In the same way, I want to thank my work supervisors, Said and Abel, and my colleagues from SUEZ who guided me through this project and helped me start my professional career as an engineer. Among them, a special thanks goes to Abir, Ibrahim, Pauline, and Chouaib, from the data science team and to the other co-interns, Imane, Pierre, and Aziz, whose friendship went far beyond professional collaboration.

I am also extremely grateful towards everyone who accepted to be interviewed as part of this research for their meaningful insights and opinions, which greatly helped me in proposing a meaningful policy at the conclusion of this work, grounded in the real-world perspectives of professionals in the sector.

Last but not least, I also want to thank my family, girlfriend, and friends, especially the ones who are here today, for all the support and affection they gave me during this project and when moving to France for the internship.

This thesis, while marking the end of my studies, is also representing the beginning of a journey I feel ready to embrace.

Thanks to all of you!

Federico Sacile
Delft, November 2025

Preface

Dear reader,

I advise you not to skip this preface and read it thoroughly, as it will be of great help to understand the rest of this work and why it has been chosen as my thesis project for my Engineering and Policy Analysis (EPA) Master's course with specialisation in Systems and Control Engineering.

I strongly believe that this research perfectly matches the goals of the EPA Master Course at TU Delft, because, like the name itself suggests, it requires both Engineering and Policy Analysis skills. Engineering methods, according to my specialisation in systems and control engineering and AI, are needed for the technical part of this project, while policy analysis frameworks and concepts are crucial to picture how this innovation can impact the sociotechnical system of critical water infrastructures.

Additionally, having always been passionate about sustainability, the environment, and water systems, I believe that the topic chosen could not be better: applying control engineering and AI methods to critical water infrastructures, perfectly matching my personal interests.

Moreover, another goal of this thesis, which coincides with the goal of EPA, is to bridge the gap between technical experts and political decision-makers. It has been decided to start from a highly technical perspective, which is the one I had in my role as a Data Scientist during my internship at SUEZ Digital Solutions.

The motivations for this choice are connected to the fact that I personally believe that high-level strategic decision-making cannot be made without a deep understanding of the lower decisional levels (tactical and operational).

That's why I chose first to dive deep into the technical aspects of AI for water treatment, learning practically how these solutions are used, developed, and implemented through the internship practice. After the internship at SUEZ Digital Solutions, where I personally analysed real operational data and built ad-hoc AI solutions to predict and control water treatment processes, I believe I gained sufficient technical knowledge to move forward and analyse higher levels of decision-making.

In the second part of my thesis, I have conducted interviews with policy makers, regulators and technical experts to understand the political and socio-technical context in which the AI solutions I developed would be implemented and try to formulate strategies to facilitate safe AI-adoption in critical infrastructure.

I think this is a similar process to what will happen (hopefully) in my work career as an EPA graduate with a technical background: I believe I will start my journey in a technical role and then move forward in the higher levels of decision-making, and this is why my thesis follows this development.

Federico Sacile
Delft, November 2025

Abstract

This thesis investigates how control systems based on Artificial Intelligence (AI) can be safely implemented to optimise the water-related processes in critical infrastructures across Europe, focusing on wastewater treatment plants (WWTPs) in Italy, France, and the Netherlands. The study addresses the overarching question: "How can AI-based solutions be technically effective, economically viable, safe, and acceptable for sustainable water management in Europe?" The research follows a dual-track approach combining engineering development and policy analysis. On the engineering side, the work develops and evaluates deep reinforcement learning (DRL) controllers for aeration optimisation within activated sludge systems using real operational data provided by SUEZ Digital Solutions. Two state-of-the-art agents, Soft Actor–Critic (SAC) and Twin-Delayed DDPG (TD3), are trained interactively on a linear model for aeration, to respect operational constraints and improve the process. The agents were trained with two different configurations: with and without a buffer of historical transitions that is used as previous knowledge. After training, the agents were benchmarked across multiple disturbance scenarios, generated from real data of energy price and inflow load. Results demonstrate significant improvements compared to baseline control, achieving lower energy consumption, stable dissolved oxygen levels, and better values of redox potential in the tank. These findings confirm the technical feasibility and scalability of DRL-based aeration control for real-world deployment. On the policy side, the research explores the institutional and governance readiness for adopting AI-based control in critical water infrastructures across Italy, France, and the Netherlands through fifteen semi-structured interviews with regulators, utility managers, and researchers. Using the Transition Model Canvas (TMC) and Multi-Level Perspective (MLP) frameworks, the analysis identifies key barriers and leverage points. In the first group, fragmented governance, infrastructural limits, and lack of AI literacy among stakeholders at all levels were identified, while the second one included regulatory sandboxes, digital-skills training, and pilot projects. Comparative insights show that France benefits from strong national coordination and incumbents (SUEZ, Veolia), Italy faces heterogeneous regional governance and uneven digitalisation, while the Netherlands provides a model of integrated and innovation-oriented regulation. By integrating both perspectives, the thesis proposes a Transition Model Canvas for AI-based wastewater infrastructure at a European level. It maps how landscape pressures (EU AI Act, Green Deal, and Urban Wastewater Treatment Directive recast) interact with regime actors and niche innovations to shape transition pathways. The work concludes with a set of policy and design recommendations for safe and responsible AI adoption, structured into short-, medium-, and long-term phases. Overall, this thesis demonstrates that AI-based control systems can substantially improve energy efficiency, regulatory compliance, and sustainability in wastewater management. Their successful adoption, however, requires coordinated regulatory frameworks, skills, and investment in digital infrastructures. The study illustrates how combining systems engineering with policy analysis can support the responsible digital transformation of Europe's critical water infrastructures.

Contents

List of Abbreviations

ABAC: Ammonia-Based Aeration Control
AI Act: Artificial Intelligence Act
AI/ML: Artificial Intelligence / Machine Learning
AOB: Ammonium-Oxidizing Bacteria
ARERA: Italian Regulatory Authority for Energy, Networks and Environment
ARG: Antimicrobial/Antibiotics Resistance Gene
ASP: Activated Sludge Process
ATT-LSTM: Attention-Enhanced Long Short-Term Memory Network
CER: Directive on the Resilience of Critical Entities
CIRSEE: Centre International de Recherche Sur l'Eau et l'Environnement
CSV: Comma-Separated Values
DDPG: Deep Deterministic Policy Gradient
DO: Dissolved Oxygen
DRL: Deep Reinforcement Learning
EC: European Commission
EIB: European Investment Bank
EPA: Engineering and Policy Analysis
EU: European Union
GPU: Graphics Processing Unit
IoT: Internet of Things
INFR: EU Infringement Procedure Code
ISO: International Organization for Standardization
ISTAT: Italian National Institute of Statistics
KPI: Key Performance Indicator
LSTM: Long Short-Term Memory Network
LSTM-ATT: Long Short-Term Memory with Attention
MBR: Membrane Bioreactor
MLP: Multi-Level Perspective
MPC: Model Predictive Control
NDA: Non-Disclosure Agreement
NH_4: Ammonium
NN: Neural Network
NOB: Nitrite-Oxidizing Bacteria

npz: NumPy Compressed Archive Format
ODE: Ordinary Differential Equation
ORP: Oxidation–Reduction Potential
OTE: Oxygen Transfer Efficiency
P / I / D: Proportional / Integral / Derivative Terms
PID: Proportional-Integral-Derivative Controller
pH: Potential of Hydrogen
ReLU: Rectified Linear Unit
RL: Reinforcement Learning
R2: Coefficient of Determination
SAC: Soft Actor–Critic
SBR: Sequencing Batch Reactor
SCADA: Supervisory Control and Data Acquisition
SDG: Sustainable Development Goal
TD3: Twin-Delayed Deep Deterministic Policy Gradient
TMC: Transition Model Canvas
UWWTD: Urban Wastewater Treatment Directive
UN: United Nations
WFD: Water Framework Directive
WHO: World Health Organization
WWT: Wastewater Treatment
WWTP: Wastewater Treatment Plant

1

Introduction

1.1. Defining the Scope and Boundaries of the Research

The objective of this thesis is to analyse, from a multidisciplinary perspective, how recent advancements in Artificial Intelligence (AI) research can contribute to optimising the management of critical water infrastructure, improving energy efficiency and compliance with sustainability goals. Given the breadth of the topic, this research narrows its technical focus on *wastewater treatment plants* (WWTPs), and more specifically, to the *aeration process* within the activated-sludge system. In this context, WWTPs are regarded as part of the broader category of *critical infrastructures*, as defined by the European Directive on the Resilience of Critical Entities (CER Directive, EU 2022/2557) [1]. These systems provide an essential public service whose disruption would have significant impacts on human health, environmental protection, and societal functioning. So, WWTPs represent a typical example of operationally critical, cyber-physical infrastructures where reliability, safety, and compliance are as important as performance and efficiency. Framing them in this way also provides a consistent link between the technical and policy dimensions of this research, since the integration of AI-based control in such systems raises domain-specific governance and regulatory challenges.

As will be discussed later (Section 1.2, Section 1.4.4), this subsystem is both representative and critical from an engineering standpoint. Aeration alone accounts for more than 50% of a WWTP's total electricity use [2], while directly influencing effluent quality and nitrogen removal efficiency [3, 4]. This combination of high energy demand, strong process interactions, and rich data availability makes aeration an ideal target for optimisation and a realistic environment in which to experiment with *data-driven control*. In particular, the thesis adopts *deep reinforcement learning* (DRL) as the main methodological frontier since it represents the latest evolution in process control and its ability to learn control policies directly from interaction data, going beyond traditional model-based or static rule-based optimisation [5, 6, 7]. This emerging paradigm aligns with current industrial research directions seeking self-adaptive control systems that can improve both energy efficiency and environmental performance in real time. The choice to focus on WWTPs is also motivated by their strategic role in achieving the Sustainable Development Goals (SDGs), particularly Goal 6 on clean water and sanitation [8], and by the pressing need to enhance energy efficiency and reduce greenhouse gas emissions from wastewater operations [9, 10]. Moreover, WWTPs are at the forefront of digitalisation within the water sector, characterised by extensive sensors, supervisory control systems (SCADA), and increasingly automated data collection, which together provide the basis for testing AI-enabled controllers [11].

From a policy perspective, this research analyses how governance and organisational frameworks influence and are influenced by the adoption of AI in critical infrastructures, with particular attention to stakeholder perspectives. For this part, the research concentrates on the *European context*, with a specific focus on Italy, France, and the Netherlands, which provide complementary institutional and technological environments. France hosts major global utilities such as SUEZ and Veolia, which are pioneers in digitalisation and AI integration in wastewater management [12]. On the other hand, Italy represents a fragmented but rapidly evolving landscape, where regional utilities increasingly participate in EU-funded innovation and modernisation programmes [13, 14]. Finally, the Netherlands is historically recognised for its world-leading expertise in water management, offering an integrated and innovation-oriented governance model exemplified by early pilot initiatives such as the *Digital Delta* and the *Water & Energy Factory* [15]. Together, these cases enable a comparative analysis of how different governance systems, regulatory pressures, and organisational cultures shape the adoption of AI-based control systems under the evolving European legislative framework [16, 1].

Although the technical case study is centred on wastewater treatment, this research also includes *interviews with experts managing other types of water-related critical infrastructures*, including dams, drinking-water networks, and desalination systems. These additional perspectives are valuable for identifying cross-sectoral challenges, such as ensuring safety, explainability, and human oversight in high-risk AI applications [17, 18], and for situating wastewater treatment within the broader transition towards resilient and digitally enabled infrastructures [19].

In summary, this thesis adopts a focused yet integrative scope that balances depth and generality. This scoped framing ensures coherence between the technical and policy analyses, allowing the study to go beyond the simple investigation of how reinforcement learning can optimise control at the process level, by also analysing its safe and effective integration into critical infrastructures within the European regulatory and institutional context.

1.2. Global Relevance of the Optimising Water Infrastructure

Water and wastewater management play a key role in sustainability and environmental protection. Globally, water scarcity, demographic growth, and urbanisation have increased the demand for efficient infrastructures and processes in the field. In 2017, the United Nations (UN) World Water Development Report (WWDR) estimated that over 80% of wastewater was being discharged untreated, posing severe risks to ecosystems and public health [20]. More recent assessments provide a clearer picture of progress: according to the UN Human Settlements Programme (UN-Habitat) and the World Health Organization (WHO), an estimated 268 billion m^3 of household wastewater was generated in 2022, of which 155 billion m^3 (58%) was safely treated, while 113 billion m^3 was still released in the environment without safe treatment [21]. Poor wastewater management also slows down progress toward a circular economy, resulting in a loss of resources [22].

In this context, WWTPs, which purify water from domestic and industrial sources before discharging it back into the environment, are classified as *critical infrastructure*. As defined in European Directive 2008/114/EC, this refers to “an asset, system or part thereof [...] which is essential for the maintenance of vital societal functions, health, safety, security, economic or social well-being of people, and the disruption or destruction of which would have a significant impact in a Member State” [19]. However, these stations are energy-consuming and emit greenhouse gas (GHG). Additionally, according to Our World in Data, wastewater treatment accounts for 1.3% of total global GHG production [23], as organic matter decomposition produces methane and nitrous oxide [24].

As a matter of fact, improving and optimising wastewater treatment is a relevant sociotechnical problem according to the UN’s 2030 Agenda for Sustainable Development [8] and it aligns with several Sustainable Development Goals (SDGs). According to the article written by Obaideen and colleagues (2022), wastewater treatment contributes effectively in achieving 11 out 17 SDGs [9]. In the specific case of this proposed thesis, great adherence can be noticed with the following:

- SDG 3 - Good Health and Wellbeing: as stressed by the health risk assessment done by Isah et al., wastewater treatment is extremely important, highly relevant to human health [8, 25, 9].
- SDG 6 - Clean Water and Sanitation: by promoting efficient water management and improving the performance of wastewater treatment systems. [26, 8, 9].
- SDG 9 - Industry, Innovation and Infrastructure: through the adoption of innovative control strategies based on Artificial Intelligence (AI) to modernise infrastructure [26, 8, 9].
- SDG 11 - Sustainable Cities and Communities: via smart wastewater management that enhances urban resilience [8, 9].
- SDG 12 - Responsible Consumption and Production: by enabling circular resource flows such as energy and nutrient recovery [8, 9].
- SDG 13 - Climate Action: by mitigating GHG emissions from energy-intensive operations like aeration and sludge handling [27, 10, 8, 9].
- SDG 14 - Life Below Water: an efficient wastewater treatment greatly reduces impacts on the aquatic environment, such as discharge of micro pollutants or eutrophication [28, 29, 8, 9].

It is clear at this point that intelligent, energy-aware optimisation of operations is essential to support climate mitigation and transition towards circular economies. With accelerating urbanisation and population growth, wastewater infrastructure is under pressure. Dynamic, AI-based control systems could offer a promising solution to handle fluctuating influent loads and stringent effluent standards.

From a technical perspective, inefficiencies in current wastewater treatment operations result in high energy use and operational costs, paired with avoidable environmental impacts. At the same time, influent variability and strict effluent standards demand adaptive, real-time optimisation. From a strategic and governance perspective, the sector faces growing pressure to integrate digital and AI-based solutions into critical infrastructure, while ensuring transparency, accountability, and broad stakeholder acceptance. This dual urgency underscores the need for innovative approaches that are not only technologically sound but also institutionally viable and aligned with long-term sustainability goals.

1.3. Relevance to the Engineering and Policy Analysis Program

This research is tightly aligned with the major objectives of the Engineering and Policy Analysis (EPA) Master's programme, as it addresses a complex sociotechnical system where infrastructure, policy, and innovation intersect. Specifically, this work tackles the Grand Challenge of achieving sustainable, climate-resilient resource management by transforming critical wastewater infrastructure.

The thesis offers a unique dual-perspective on the topic, diving deep into its different dimensions in line with the multidisciplinary nature of the EPA programme. In this context, the adoption of AI for optimising water management and control can be framed in two ways.

On the one hand, it is mandatory to consider the engineering side of the problem, which involves addressing technical feasibility, system development, cost optimisation, and environmental benefits. On the other hand, implementing AI on critical infrastructure is also a strategic challenge. It is crucial to assess if the system is ready to safely implement AI-based control for water treatment processes and whether adequate regulation exists to govern and support this innovation. And, if preparedness has not been reached yet, the goal is to understand why and which actors are driving the resistance.

This wicked problem cannot be effectively addressed by treating the technical and strategic perspectives separately. A robust engineering understanding is essential not only to develop and assess innovative solutions but also to inform sound policy decisions. At the same time, strategic insight is required to ensure that innovation is implemented safely, responsibly, and in alignment with broader societal goals. These perspectives are deeply connected: policies must be grounded in technical realities, and industrial innovation must anticipate regulatory and societal dynamics. Navigating this interdependence requires a rare combination of skills, which reflect precisely the kind of interdisciplinary expertise cultivated within the EPA programme.

1.3.1. Engineering and Technical Relevance

From a technical standpoint, this thesis builds on the engineering pillars of the EPA programme, requiring modelling methods and strong data science skills to develop AI-based control systems. The author's specialisation choices also play a decisive role: *Digital Control* from the MSc in Systems and Control contributes tools for advanced controller design, while *Planning and Decision-Making* from the MSc in Robotics offers techniques for exploring scenarios and optimising performance under uncertainty.

Additionally, from an academic perspective, this work is highly innovative, exploring a new frontier in the application of AI to critical infrastructure. Research in this domain remains in its early stages, and while some studies have applied advanced AI techniques (e.g. deep RL controllers) to wastewater treatment processes, these efforts have predominantly relied on simulations rather than real-world implementations. The technical novelty of this thesis is better addressed in the next Chapter in Section 2.2.

1.3.2. Policy Relevance

From a policy standpoint, this thesis builds on the strategic and institutional pillars of the EPA programme, drawing on courses such as *Actor and Strategy Models* and *Behaviour in Transitions*, which provide the tools to analyse innovation in complex governance environments.

In particular, this work applies the Transition Model Canvas (TMC) [30] to assess how AI-based solutions for wastewater treatment interact with regulatory regimes, incumbent practices, and landscape-level pressures. This approach is highly relevant to the EPA programme, as it combines socio-technical transition analysis with stakeholder models and interviews to assess the system preparedness and understand how to push innovation. While the detailed novelty of this contribution is better addressed in Section 2.2, it is worth mentioning that few studies have explored the governance of AI in critical infrastructure, making this work a substantive and original contribution to the policy field.

1.4. Introduction to Key Topics in Wastewater Treatment

After the previous high-level overview, the reader can understand the global relevance of optimising wastewater treatment and how this needs to be addressed from a multidisciplinary perspective. Now, before diving deeper into the rest of this work, it is mandatory to provide an introduction to key topics in wastewater treatment, to familiarise oneself with concepts that will be later taken for granted.

1.4.1. WWTPs and Wastewater Treatment Processes

WWTPs are designed to protect environmental and public health by removing contaminants before discharging effluents back into the environment. In these infrastructures wastewater undergoes a treatment process divided in four sequential stages, as described by Lakshmana Prabu and colleagues. and as shown in Figure 1.1. These are: *preliminary*, *primary*, *secondary*, and *tertiary* treatment [31]. Other classifications, however, might group preliminary operations under primary treatment, and thus commonly refer to the three main stages of wastewater treatment as primary, secondary, and tertiary.

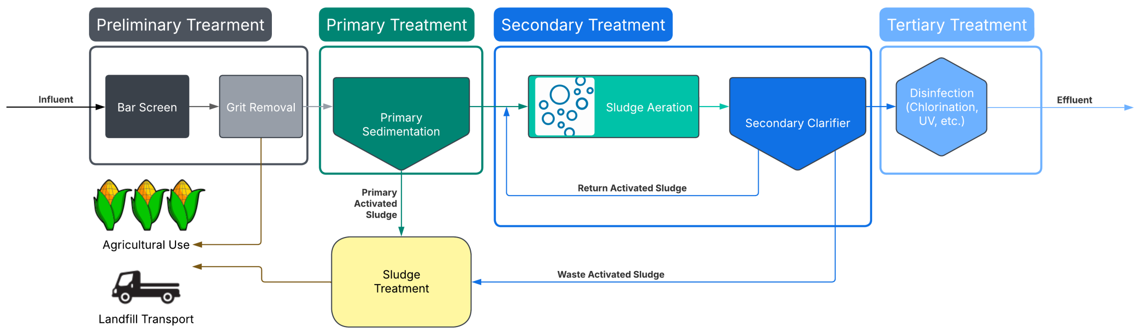

Figure 1.1: Main Stages in Wastewater Treatment (created with Lucidchart)

Preliminary treatment consists of physical operations such as screening, grit removal, and grease elimination. This dispose of large debris, abrasive particles, and oils that would interfere with downstream processes. Other preliminary operations can include flow equalisation and neutralisation [31].

Primary treatment involves sedimentation, also called primary clarification or decantation. In this phase, wastewater is temporarily held in large tanks to remove settleable solids thanks to gravity. This step can be enhanced by *chemical precipitation*, where reagents convert dissolved substances into compounds that settle with the sludge. The result is a clarified effluent suitable for biological treatment [31].

Secondary treatment involves removing soluble organic matter and smaller suspended solids by using activated sludge systems, trickling filters, aerated lagoons, bioreactors, and ponds [31]. In particular, among these, many authors mention the activated sludge process (ASP) as the most widely used technique for the biological treatment of domestic and industrial wastewaters [32, 33]. Given relevance both globally and specifically for this research, a better description of the ASP is provided in Section 1.4.3.

Finally, tertiary treatment provides advanced purification by removing remaining nutrients, heavy metals, pathogens, and residual organics. Techniques employed in this stage include chemical coagulation followed by sedimentation, filtration, activated carbon adsorption, ion exchange, and membrane processes such as reverse osmosis. Disinfection methods, including chlorination, ozonation, and ultraviolet (UV) irradiation, are also used to safeguard effluent quality before discharge or reuse [31].

1.4.2. Main Water Pollutants and Their Effects

In the matter of pollutants, extensive research is available in the literature, as shown, for example, by the work of Akpor and colleagues [34]. For the purpose of this thesis, it is sufficient to introduce the group of contaminants most relevant to the optimisation focus of this research: excess nutrients, particularly nitrogen (N) and phosphorus (P). These components, which remain only partially removed in many WWTPs, when released in aquatic habitats, trigger algal blooms, causing eutrophication, oxygen reduction, fish deaths, and long-term biodiversity loss [34]. This ecological damage directly impacts water supply, recreational use, and fisheries. Eutrophication remains an EU-wide issue, affecting more than 30% of rivers, lakes and coastal waters and 80% of EU marine waters. Reflecting this urgency, the Urban Waste Water Treatment Directive recast (2024) provisions have been updated and harmonised to further stress nutrient removal as a priority in WWTPs [35].

Given this context, it is then clear how nutrient management remains one of the main long-standing priorities in wastewater research and practice, directly motivating the focus on aeration optimisation adopted in this work, since nitrification, and hence aeration, is central to nitrogen removal. For completeness, a more detailed overview of water pollutant classes and associated treatment considerations is presented in Appendix A.

1.4.3. Activated Sludge Process

Given that the ASP is both the foundation of biological treatment and the stage involving aeration, which is the focus of the technical part of thesis, it is now discussed in greater detail.

Historical Background

The activated sludge process was first introduced by Edward Ardern and William T. Lockett, two chemists working at the Davyhulme Sewage Works in Manchester, in 1914 [36]. In their pioneering paper, they demonstrated that sewage aeration alone was insufficient for practical purification, but that reintroducing the biologically active solids from previous aerated sewage accelerated both oxidation and nitrification [36]. This material, which they named *activated sludge*, contained microbial flocs capable of sustaining rapid and continuous treatment when recycled into fresh sewage. Their findings provided the scientific basis for modern biological wastewater treatment, shifting the field from empirical filtration and chemical precipitation methods toward engineered microbial processes.

General Description

Nowadays, this crucial process is widely known and studied. As described by the Water Environment Federation, in a conventional ASP configuration, influent wastewater is combined with a microbial community and oxygen in the aeration basin, where microorganisms oxidise organic matter (ammonium and phosphorus composites) and form flocs. The mixed liquor then passes into a secondary clarifier, where solids settle to the bottom. A portion of the settled biomass is recycled as return activated sludge (RAS) to maintain process stability, while the excess is removed as waste activated sludge (WAS). A visual scheme of the process is provided below in Figure 1.2. ASP can be implemented through several reactor configurations, each with specific advantages and limitations. While these variations are not essential for the discussion here, a more complete overview is provided in Appendix A.

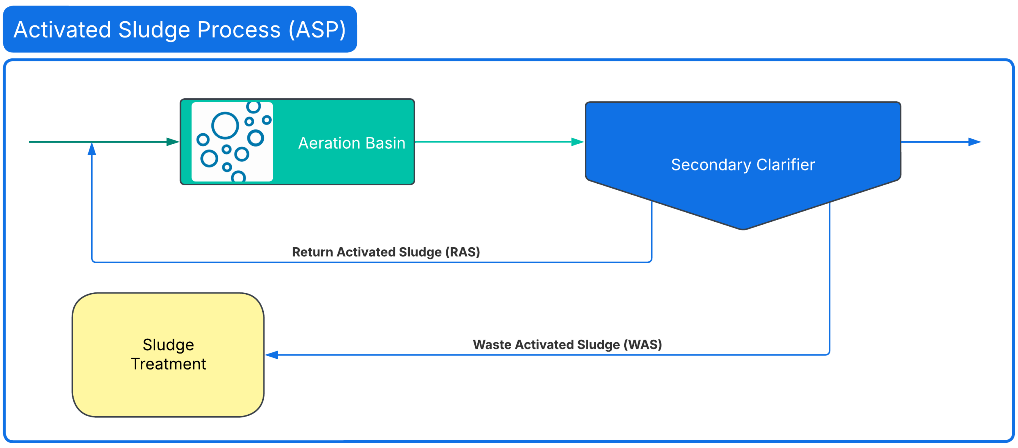

Figure 1.2: The Activated Sludge Process (created with Lucidchart)

Microbiology

From a microbiological perspective, activated sludge relies on diverse species of bacteria, protozoa, fungi, and other organisms to carry out oxidation and nutrient transformations, making it the most versatile and prevalent biological treatment process worldwide [37].

Ammonium-oxidising bacteria (AOB) such as *Nitrosomonas* convert ammonia into nitrite, while nitrite-oxidising bacteria (NOB) such as *Nitrobacter* and *Nitrospira* transform nitrite into nitrate. In this two-step nitrification, which is the first part of the nitrogen cycle, bacteria require oxygen to sustain microbial metabolism, oxidise organic matter, remove nutrients, and ensure effluent quality. Thus, this step is directly linked to the aeration regime, which is better described in Section 1.4.4.

Subsequently, under anoxic conditions (absence of dissolved oxygen but presence of nitrate), denitrifying bacteria such as *Paracoccus* and *Pseudomonas* reduce nitrate to nitrogen gas. This step completes the nitrogen cycle and is essential to prevent eutrophication in receiving waters [38, 39].

1.4.4. Aeration: Relevance and Methods

As said in Section 1.4.2, nutrient removal is one of the main priorities in wastewater research, and aeration is the precise step that impacts nitrification (Figure 1.4.3). Additionally, water aeration is particularly energy-intensive, accounting for more than 50% of the total electricity demand of a WWTP [2]. This combination led the research to focus on this process to develop one of the first industrial implementations of AI to control WWTPs.

A more precise (and mathematical) description of what happens in an aeration tank from a physical, chemical, and microbiological perspective can be found in Chapter 5, where an Ordinary Differential Equations (ODEs) model for the nitrogen cycle is developed and explained.

Among the several aeration methods currently employed in WWTPs, the most widespread is diffused aeration, where air is introduced at the bottom of tanks via fine or coarse-pore diffusers. Fine-pore diffusers provide high oxygen transfer efficiency (OTE) but are prone to fouling and scaling, which increase energy demand unless cleaning and maintenance are performed regularly [40]. In contrast, mechanical surface aerators, which agitate the water using paddles or brushes, offer robust mixing but have comparatively lower oxygen transfer efficiency and higher energy consumption [2]. In smaller or rural systems, typically in mountain regions, the so-called ”cascade” or ”fall method” can also be used. This consists of letting wastewater flow down steps or cascades to absorb atmospheric oxygen.

Recent innovations have targeted energy savings. High-speed turbo blowers can reduce aeration energy use by up to 42% compared with traditional blowers [41, 42]. Dynamic aeration techniques, where airflow is pulsed or periodically interrupted, improve oxygen transfer rates by up to 24% [43]. Finally, smart aeration control strategies, such as Ammonia-Based Aeration Control (ABAC), integrate online ammonia sensors and predictive algorithms to dynamically match oxygen supply with microbial demand, achieving 20 to 40% reductions in energy use without compromising effluent quality [44].

In conclusion, aeration is both the biological ”engine” of the ASP (sustaining the nitrogen cycle and ensuring compliance with effluent standards) and the main driver of WWTP energy consumption, dominating operational costs. The wide range of available methods reflects the continuous search for a balance between treatment efficiency and energy sustainability. This dual importance makes aeration the most relevant process to improve through AI, soft sensing, and reinforcement learning.

1.4.5. Sensors in Aeration Tanks

As data-driven innovation is done, in this case, by using sensor data, an introduction to the typical sensors in aeration tanks is a must for understanding what is done later. To track biological processes and performances in modern aeration tanks, three main measurements are monitored: dissolved oxygen (DO), oxidation–reduction potential (ORP), and ammonium (NH_4^+).

DO is an important process variable for aerobic treatment and energy use and sensors to track it are widely used in aeration tanks. Two families dominate: (i) electrochemical (polarographic/galvanic) cells, whose current scales with oxygen flux through a membrane; and (ii) optical sensors, where excited luminophores undergo oxygen-dependent fluorescence quenching (intensity or lifetime) described by Stern–Volmer relations. Optical probes have largely displaced classical cells due to lower flow dependence and maintenance burden [45, 46, 47, 48].

ORP is an aggregate redox indicator measured as the potential of an inert metal (often Platinum (Pt)) versus a reference electrode. Although non-specific, ORP tracks transitions between oxidising and reducing regimes and can mark denitrification endpoints, inform intermittent aeration, and flag process upsets when interpreted in context. Recent evaluations document both its usefulness and the pitfalls of over-interpreting single set-points across varying matrices [49, 50, 51]. In practice, ORP is paired with DO to avoid aerating through anoxic targets or to schedule aeration pulses. Due to its simplicity and low price, it is always present in aeration tanks.

Real-time NH_4^+ enables ABAC by matching aeration to nitrification demand. Three online approaches are common: ion-selective electrodes (ISEs) using ammonium-selective ionophores in polymer membranes; wet-chemistry colourimetric analysers (e.g., salicylate/indophenol); and UV/UV-Vis spectroscopic methods. ISEs are reagent-free but sensitive to pH and K^+; colourimetric analysers achieve low detection limits at the cost of reagents and maintenance; UV approaches are compact but matrix-dependent and often require site-specific calibration. Reviews and field protocols provide practical guidance for achieving stable ISE performance in wastewater and wetlands [52, 53].

Table 1.1 shows typical price bands for these sensors, to understand the type of investment that is required to move towards digital control in WWTPs.

Sensor	**Example item / source**	**Indicative price**	**Ref.**
DO (portable)	YOGAYE Portable Dissolved Oxygen Meter, Amazon IT [54]	€90–€130 (entry-level)	[55]
DO (industrial RS485/4–20 mA)	L-Com SRWQ100-DO104 Industrial DO Sensor [56]	US $1,620 (≈€1,500)	[55]
ORP (handheld)	Generic ORP Pen Tester, Amazon IT [57]	€25–€35 (entry-level)	[58]
ORP (industrial transmitter)	Omega TX-ORP Industrial Transmitter [59]	US $200–300 (typical)	[58]
NH_4^+ (ISE, educational)	Vernier Ammonium Ion-Selective Electrode (S16214ND) [60]	US $345 (≈€320)	[52]
NH_4^+ (cabinet analyzer)	Hach Amtax sc NH_4^+ Analyzer	Quote-based, ∼US $10,000–15,000	[61]

Table 1.1: Indicative prices for sensors in aeration control (checked Sep 10, 2025). Marketplace prices vary.

1.4.6. Key Points

It can be said that the ASP remains the backbone of biological wastewater treatment due to its adaptability, efficiency, and established engineering frameworks. Within this process, aeration, while central to biological activity, is also the largest energy consumer [2]. Therefore, it is clear that optimising this subsystem offers significant potential for cost reduction and environmental performance. As such, it presents a promising target for the application of AI-based control and soft sensing.

1.5. Foundational Concepts for a Broad Audience

Given the goal of this research, which is to bridge the gap between technical experts and policymakers, it has been decided to provide explanations of key concepts in both worlds, to make the work self-standing and clearly understandable by people with different backgrounds. However, these concepts can seem unnecessary for certain readers; therefore, the explanation has been located in Appendix A.

There, readers with a non-technical background can find definitions and high-level explanations of key concepts in Systems and Control Engineering that are taken for granted in this research. These are: Open- and Closed-loop control, Relay Controllers, Proportional-Integrative-Derivative (PID) Controllers, Optimisation-based controllers (from Linear Optimisation (LO) to Model Predictive Control (MPC)), and AI-Based Controllers (Neural Networks (NN), Reinforcement Learning (RL), and Deep Reinforcement Learning (DRL)).

In the same way, readers who are not familiar with social and decision sciences will find there useful information on concepts such as Decision-Making Levels (Strategic, Tactical, and Operational), Socio-Technical System, Actor, Policy (Control Engineering and Public Governance Meanings), Policy vs Politics, and the Transition Model Canvas to analyse socio-technical transitions.

1.6. Industrial Perspective: SUEZ Digital Solutions

To conclude this chapter, an introduction to the SUEZ Group and its business unit SUEZ Digital Solutions has to be made, as it is one of the main partners of this project, providing technical expertise on the topic.

Figure 1.3: Main entrance (left) and geographical location (right) of SUEZ Digital Solutions headquarters in Le Pecq, Île-de-France, France (Source: Google Maps, accessed on Sept 11, 2025).

The SUEZ Group

SUEZ is a top global water and waste management company operating in more than 40 countries and is highly relevant in the industry thanks to its emphasis on innovation and digital transformation. With over 160 years of history, SUEZ has developed strong expertise in water supply, wastewater treatment, resource recovery, and waste valorization, becoming a major player in the ecological transition and the circular economy [62]. The Group is recognized worldwide for its ability to develop solutions that improve sustainability and operational efficiency across industries.

SUEZ Digital Solutions

At the core of SUEZ's innovation strategy is *SUEZ Digital Solutions*, formerly known as SUEZ Smart Solutions (3S). This expert business unit employs over 600 data scientists, engineers, and AI specialists dedicated to developing and implementing advanced digital tools for process optimization, risk prediction, and smart infrastructure management [12]. Its solutions are widely deployed in water and WWTPs as well as solid waste management facilities, reshaping operational paradigms and supporting clients' transition toward sustainable and resilient systems.

SUEZ Digital Solutions originated from the *Centre International de Recherche sur l'Eau et l'Environnement* (CIRSEE) and operates as a subsidiary of both SUEZ and SUEZ Eau France (SEF). With over a decade of experience, the company has grown to become a market leader in smart water solutions, serving more than 600 clients globally and contributing to sustainable city development worldwide.

Mission and Vision

The mission of SUEZ Digital Solutions is to support the environmental and digital transformation of regions worldwide by integrating innovative technologies with operational expertise. Its vision is to enable rational use of resources, enhance resilience to crises, and accelerate the ecological transition through data-driven, adaptive solutions.

Data Science and Optimization Team

This thesis internship is carried out within the *Data Science and Optimization* team at SUEZ Digital Solutions. The team's focus is on applying state-of-the-art machine learning and optimization techniques to propose new solutions and improve processes in water and waste management. In the specific case of this project, the objective is to explore and validate new approaches (such as DRL) for enabling intelligent, adaptive control mechanisms that reduce energy consumption, enhance process efficiency, and ensure compliance with regulatory requirements.

2

State of the Art and Knowledge Gaps

2.1. Literature Synthesis

Following the dual track of this research, two different literature reviews have been conducted to introduce both the technical and the policy problem. This perspective on the topic is already unique in the literature, representing a good innovation point for the research. However, the final goal is to be able to guarantee innovation for both sides, even if considered independently.

2.1.1. Control and Decision-Making in Water Critical Infrastructure

The application of AI to water infrastructure has significantly expanded in recent years, driven by the need to manage increasingly complex systems under uncertainty, while balancing multiple and often conflicting objectives. Much of this work originates in water resources engineering, where strategic operational decisions, such as reservoir releases, flood control, and water allocation, have long required sophisticated optimisation frameworks.

A first generation of methods relied on evolutionary multi-objective optimisation, which enabled researchers to derive control policies that represent trade-offs between competing goals. For example, [63] examines how the choice of radial basis functions affects the performance of Evolutionary Multi-Objective Direct Policy Search (EMODPS) policies in benchmark reservoir systems. Their work highlights a central insight for water-related problems: the representation of the policy is as important as the optimisation algorithm itself, as it shapes the transparency, stability, and interpretability of decisions.

Building upon these approaches, recent research has turned toward reinforcement learning (RL), and in particular *multi-objective* RL (MORL). MORL allows learning a *set* of policies corresponding to different trade-offs, rather than collapsing objectives into a single scalar reward. This is particularly appealing in water governance domains where competing interests must be considered simultaneously. However, the adoption of MORL introduces a new challenge: how should decision-makers interpret and compare multiple learned policies? Addressing this question, [64] proposes an explainability-driven method that clusters and summarises Pareto-optimal policies based on both behavioural similarity and trade-off structure. The method provides a compact and interpretable representation of MORL solutions, an important step toward making AI-based decision support acceptable to operators and regulators.

The growing maturity of MORL methods is reflected in their application to real, large-scale, high-stakes systems. A recent contribution by [65] applies MORL to the Nile River Basin, one of the most politically and hydrologically complex water systems in the world. By modelling the basin as a multi-objective Markov decision process, the authors optimise dam releases across Ethiopia, Sudan, and Egypt while balancing hydropower production, irrigation needs, and downstream water security. The study illustrates the potential of learning-based approaches to navigate long-term strategic trade-offs under deep uncertainty.

Beyond reservoirs and river basins, AI techniques have also been adopted in other critical water infrastructures. In urban drainage and combined sewer overflow (CSO) systems, RL has been used to design real-time control strategies that coordinate gates and pumps to reduce flooding and overflow volumes compared to rule-based or Model Predictive Control (MPC) baselines [66, 67]. Deep RL has been also applied to drinking water networks, to optimise pump scheduling and pressure control [68, 69]. Across these diverse applications, AI consistently provides value by supporting anticipatory, adaptive, and multi-objective decision-making in systems that are too complex for traditional deterministic or rule-based control.

Taken together, this body of literature reveals a clear trend: AI, and in particular RL, is increasingly used to manage complex interactions, uncertainties, and trade-offs in water systems. Building on this, the next section, turns to the specific challenges, opportunities, and existing research on optimisation and smart control strategies for aeration, situating the present thesis within this emerging frontier.

2.1.2. Optimisation and Smart Control Strategies for Water Aeration

The treatment of municipal wastewaters within modern WWTPs is a complex, non-linear process which relies on several dynamic inputs [5]. As introduced in the previous Chapter, in Section 1.4.4, in these facilities, the aeration process is particularly energy-intensive, accounting for more than 50% of a WWTP's total electricity demand [2]. Therefore, optimising its control, while representing a relevant technical challenge, also guarantees huge improvements in energy consumption.

Historically, aeration has been controlled through simple PID or Relay controllers. These simple and reactive controllers often struggle with influent variability, sensor uncertainty, and the multi-objective trade-offs between energy costs and effluent compliance [70]. A more advanced control solution is represented by MPC, which chooses a control action on a "rolling horizon" by solving an optimisation problem. However, for WWTPs, unexpected behaviour of the microorganisms, temperature changes, disturbance in the influents, and volume increase effects could lead to modelling errors [71].

Nowadays, advances in digital technologies and AI have opened up new dimensions in water management, particularly in WWTP operation optimisation via data-driven control and predictive analytics [11]. The new technologies enable real-time decision-making, fault detection, and adaptive process control, therefore ensuring both operational effectiveness and environmental compliance.

One of the earliest innovations was the development of *soft sensors*, which provide estimations and virtual measurements of key parameters (e.g., ammonium or nitrate) from online data, without relying on expensive or failure-prone hardware sensors and enabling near-real-time feedback. Recent reviews highlight a shift from mechanistic models to ML-based predictors, including neural networks (NNs) and tree ensembles [72]. Benchmarking studies demonstrate that soft sensors can achieve robust, accurate, and cost-effective monitoring, supporting improved process control [73, 11].

The most recent and dynamic frontier is RL, where controllers learn through interaction with the plant model or a digital twin. RL is particularly well-suited to wastewater treatment because it can balance competing objectives in uncertain environments. Studies applying deep RL algorithms such as Twin Delayed Deep Deterministic Policy Gradient (TD3) and Proximal Policy Optimisation (PPO) to the widely used Benchmark Simulation Model No. 1 (BSM1) show promising results. For example, in 2023, Croll and colleagues report that TD3 reduced aeration and pumping energy demand by 14,3% compared to the BSM1 default control while maintaining effluent standards [5]. This outperforms advanced domain-based control strategies such as ABAC, although future work is necessary to improve robustness [5]. Similarly, but one year later, another study on TD3 has been conducted by [74], which again demonstrated significant improvements in aeration control. However, also in this case, the experiments were performed again on the BSM1 simulation [74]. Extensions have incorporated sustainability goals, with [6] demonstrating that RL can reduce greenhouse gas emissions alongside energy use. The study of [7] brings additional improvements, showing a 46% reduction in operational costs and a 12% decrease in energy consumption while maintaining compliance with effluent discharge standards thanks to the combination of RL algorithms and Bayesian Optimisation.

A systematic literature review on RL applications in water resource management that analysed a total of 40 articles in March 2025, highlighted the fact that in all the studies considered, RL agents are trained in simulated environments rather than directly on historical data and that future research should focus on bridging the gap between simulation and real-world applications [75].

2.1.3. Regulatory Landscape and Governance of AI and Water Infrastructure

As shown previously (Section 2.1.2), the main technical challenges for implementing AI systems in critical water infrastructure include bridging the model-to-reality gap, robust fail-safe mechanisms, and improving interpretability for operator trust [75]. However, it is not only the technical side that should adapt itself to the designed policies, and it should not be like this. As a matter of fact, good and efficient policies should be flexible enough to push progress and innovation safely. Therefore, this second part of the literature review explores the European regulatory background that surrounds AI for monitoring and controlling water infrastructures, how this is managed in other parts of the world, the main barriers to its adoption, and the common strategies to push innovation.

European Regulatory Landscape

Nowadays, in Europe, AI deployment in WWTPs must comply with multiple frameworks:

- AI Act (2024): Regulation (EU) 2024/1689 establishes harmonised rules for AI across the Union and classifies AI systems used in the management of *critical infrastructure* - a category that explicitly includes water and wastewater systems under the EU Critical Entities Resilience (CER) Directive - as *high-risk* [16, 1, 19]. This classification triggers strict obligations for both providers and deployers, including implementation of a risk management system, quality assurance for training/validation data, detailed technical documentation and record-keeping, transparency obligations, human oversight measures, and continuous post-market monitoring. High-risk AI must undergo a conformity assessment before being placed on the market or deployed, and utilities must ensure ongoing compliance during operation. The table below (Table 2.1) summarises the key obligations most relevant to WWTP operators and technology providers.

- GDPR, Data Governance Act, and Data Act: These EU frameworks become relevant when WWTP data contains *personal data* (e.g., household-level inflows or identifiable operator logs). GDPR then mandates principles such as purpose limitation, data minimisation, and, in some cases, Data Protection Impact Assessments. The Data Governance Act (Reg. 2022/868) and the Data Act (Reg. 2023/2854) establish mechanisms for trusted data sharing and user rights, enabling cross-utility data exchange under secure and transparent conditions. For most aggregated process data, these acts function less as barriers and more as enablers for lawful data re-use and AI model training [17, 76, 77].
- Water Framework Directive (2000) & UWTD Recast (2024): The Water Framework Directive (2000/60/EC) establishes the objective of achieving “good ecological status” for surface waters and binds Member States to manage water bodies in a catchment-wise manner under the river basin paradigm [78]. Meanwhile, the recast UWTD (Directive (EU) 2024/3019) updates the former Urban Wastewater Treatment Directive (91/271/EEC) with substantially tighter effluent requirements, expanded pollutant removal (nutrients, micropollutants), and new energy-related obligations [13]. In particular, the recast mandates that, at the sectoral (national) level, the total energy consumed by WWTPs treating ≥10 000 p.e. shall not exceed the energy produced from renewable sources by 2045, and requires periodic energy audits of treatment plants [13].

Table 2.1: AI Act obligations for high-risk AI systems relevant to WWTP control.

Obligation	Relevance for Water Utilities
Risk management system	Hazard analysis for effluent compliance, safety constraints on aeration control.
Data & data governance	Documentation of training/validation datasets (e.g., sensor history, lab data).
Technical documentation	Clear description of model architecture, objectives, and limitations.
Record-keeping	Logging of AI decisions and operational context for audits.
Transparency & Info	Operators must understand the purpose and functioning of AI.
Human oversight	Human-in-the-loop or fallback control to prevent violations.
Post-market monitoring	Continuous performance evaluation, incident reporting.

These new obligations push utilities to reconsider how to design, operate, and upgrade wastewater systems. Rather than focusing only on pollution removal, compliance now includes strict control over energy balance and GHG emissions [13]. Recent studies, such as the one by Andrea Capodaglio, specifically discuss how the 2024 UWTD recast introduces an “energy neutrality” objective and the requirement for energy audits, even though detailed methodologies for implementing them in WWTPs are still emerging [79]. As a matter of fact, this regulatory shift creates a stronger incentive for advanced process control, optimisation, and digital innovations (such as AI) that can dynamically balance pollutant removal and energy use. However, the recast does not itself prescribe how to certify or regulate adaptive controllers, nor does it create safe testing environments for novel control systems, leaving a gap between ambition and governance [13, 14].

In addition to European regulation, it is important to note that each country has its own national strategy (France’s Plan Eau, Italy’s PNRR funding, Dutch digitalisation strategy) to support and shape digital innovation, creating enabling conditions for AI pilots.

A Global Perspective: EU Strategies Compared to US and China

In contrast to Europe’s prescriptive regulatory regime, the United States and China have adopted more flexible or sectoral approaches to AI that may facilitate faster infrastructure deployment.

In the U.S., the federal “America’s AI Action Plan” explicitly aims to “remove red tape and onerous regulation” to accelerate AI innovation and accelerate infrastructure and energy projects [80].

In China, while AI is subject to oversight, governance remains oriented around strategic state planning rather than rigid ex ante regulation; China has released interim measures for generative AI services and continues to evolve its AI legal framework as it balances innovation and control [81, 82]. These differing regulatory postures may enable more rapid experimentation and scaling of AI in infrastructure sectors outside Europe.

Governance of Water Critical Infrastructure in European Countries

The governance of water critical infrastructure in Europe sits at the intersection of multilevel water management arrangements, fragmented institutional responsibilities, and the emerging regulatory regime for digital and AI-based technologies. Literature on water governance therefore offers essential insights into why the integration of AI into operational decision-making is not merely a technical challenge but fundamentally an institutional one.

A substantial body of scholarship on *adaptive* and *polycentric* water governance shows that European water management is characterised by dispersed authority, diverse actor constellations, and decision processes spanning administrative, hydrological, and political boundaries. Huitema et al. identify four principles underpinning contemporary governance models: collaboration within polycentric systems, public participation, experimentation, and management at the bioregional scale. They also highlight the persistent difficulty of putting these principles into practice due to coordination costs, uneven power dynamics, and the political sensitivity of experimentation in essential services [83].

These issues are particularly salient within the European Union, where the Water Framework Directive (WFD) institutionalises river basin governance as a core organising principle. Gupta and Pahl-Wostl situate this arrangement within the broader debate on *multilevel governance*, showing that while authority distributed across EU, national, regional, and basin scales enhances representation and accountability, it simultaneously complicates coherent policy implementation [84]. Water in the European context is framed as a human right, an ecological medium, a strategic resource, and a security issue. As a result, responsibilities are fragmented across ministries, regulators, basin authorities, municipalities, and utilities, producing what Pahl-Wostl et al. describe as persistent ”missing links” in agenda setting, implementation, and learning across scales [83].

Comparative governance analyses further demonstrate that fragmentation is not only administrative but also political. Özerol et al. show recurring difficulties in public participation, intersectoral coordination, and justice considerations across a wide range of European case studies—including transboundary rivers, drought regimes, and flood protection—and emphasise the limited availability of longitudinal evidence on how governance arrangements evolve under technological or regulatory change [85]. Taken together, this literature reveals several structural patterns that are common across European Member States, including the three countries analysed later in this thesis (Italy, France, and the Netherlands).

First, the European water governance is inherently polycentric. Authority is distributed across numerous actors: EU institutions, national ministries, regional and basin-level bodies, municipalities, and utilities. This arrangement supports representation and learning, yet introduces significant coordination challenges [83, 84]. Second, institutional fragmentation appears not as an anomaly but as a structural feature. Even robust systems such as the Dutch water boards or the French basin agencies face difficulties aligning horizontal policies across sectors such as agriculture, land use, and energy [85, 83]. These difficulties complicate the allocation of risk and responsibility in domains such as data governance or AI-based operational control.

Third, collaboration and tensions routinely coexist, as argued by Mancilla García et al. conflicts in water management arise from the competing demands of environmental protection, agriculture, industry, and urban development [86]. Participation and collaboration, while normatively desirable, may reproduce existing power asymmetries, and conflict itself can foster institutional learning. This dynamic is directly relevant for AI-based control systems, which may shift authority between operators, engineers, utilities, and regulators.

Fourth, strong EU regulatory steering shapes national priorities. The WFD, Floods Directive, Drinking Water Directive, and the recent recast of the Urban Wastewater Treatment Directive set performance and investment requirements. At the same time, horizontal legislation creates a multilayered regulatory environment whose implications for operational AI in water systems are still insufficiently understood.

Fifth, accountability for water and AI risks is often unclear in polycentric systems. Responsibilities for human oversight, transparency, documentation, and liability may fall across utilities, basin authorities, ministries, and private vendors.

Finally, European utilities display highly heterogeneous levels of digital maturity: differences in size, capacity, and resources result in uneven readiness for deploying high-risk AI systems [87, 17].

When combined with emerging scholarship on AI governance under the EU AI Act, these governance characteristics reveal important tensions. The AI Act introduces a relatively centralised regulatory framework classifying AI used in managing water infrastructure as *high-risk* [16]. This contrasts sharply with the deeply polycentric structure of water governance, where discretion, adaptation, and experimentation are traditionally distributed across actors and levels. Research on AI governance shows that while ethical principles and risk-based approaches are well developed, regulatory clarity for high-risk domains—particularly in critical infrastructure—lags behind technological innovation [88]. Industry actors have similarly raised concerns that rigid high-risk classifications may inhibit innovation by imposing extensive compliance burdens.

In sum, the integration of AI into Europe's water critical infrastructure unfolds within a governance landscape that is at once enabling and constraining: enabling because adaptive and learning-oriented approaches are already embedded in water governance frameworks; constraining because fragmentation, conflicting actor coalitions, heterogeneous organisational capacities, and strict compliance obligations create substantial barriers to experimentation, deployment, and scaling. As later sections of this chapter will show, this multi-actor, multi-level context strongly shapes how utilities interpret the AI Act, how national strategies interact with EU directives, and how innovation pathways develop across France, Italy, and the Netherlands.

2.1.4. Barriers to AI Adoption

In the current situation, even given the EU law adaptations, obstacles for adoption are still identifiable: literature shows barriers across technical, organisational, and regulatory domains.

Data quality and integration with legacy Supervisory Control and Data Acquisition (SCADA) systems remain major issues [89]. Many European utilities continue to operate with heterogeneous SCADA systems and variable levels of automation. According to [87], the adoption of AI in infrastructure sectors necessitates an interdisciplinary approach that includes technological development, workforce adaptation, and coordinated investment strategies. In the absence of alignment among digital infrastructure, capacity-building initiatives, and organisational capabilities, such adoption is likely to remain constrained to isolated pilot implementations.

Organisationally, utilities face a digital skills gap and cultural resistance, prioritising short-term compliance over innovation. Effluent violations in Europe carry financial, legal, and reputational penalties. Utilities are therefore highly risk-averse and hesitant to rely on "black-box" AI controllers. According to [90], AI adoption requires transparency and certification, while [91] highlight that perceived trustworthiness directly shapes adoption decisions.

Additionally, regulatory uncertainty (AI Act compliance costs, liability under the upcoming AI Liability Directive) can create a *regulatory chill*, discouraging deployment [17]. Utilities are often risk-averse and prioritise short-term regulatory compliance over strategic innovation. This creates what is known as an ambidexterity dilemma — the organisational challenge of simultaneously pursuing exploitation (refining and improving existing processes for efficiency and reliability) and exploration (investing in new, potentially disruptive technologies and business models). Over-emphasising exploitation can lead to organisational inertia and missed opportunities, whereas too much exploration can compromise operational stability and compliance. Managing this balance is crucial for enabling digital transformation in critical infrastructures such as wastewater treatment plants [92].

All the barriers to adoption identified are summarised in Table 2.2, for a clear and fast overview.

Table 2.2: Key barriers to AI adoption in water infrastructure.

Category	Barrier	Description
Technical / Data	Data quality	Noisy, incomplete datasets; difficult to aggregate across plants.
	Integration	Challenges connecting AI to legacy control systems.
	Validation	Need to guarantee effluent compliance under all conditions.
	Explainability	Black-box nature limits operator trust [75].
Organisational	Skills gap	Lack of in-house AI/data science capacity.
	Culture	Conservative mindset; ambidexterity dilemma.
	Procurement	Traditional processes discourage agile pilots.
Regulatory	Compliance burden	Documentation, conformity assessments under AI Act [16].
	Liability	Unclear responsibility for AI-induced incidents.
	Privacy	GDPR restricts data sharing needed to train models.

2.1.5. Strategies for Adoption and Transition

Common strategies to accelerate the adoption of AI-based or other innovative control systems in critical infrastructure include: (i) Phased deployment with human-in-the-loop supervision, where new controllers are first validated in simulations and rolled out in limited pilots before full automation, ensuring operators retain override capacity during early phases [93]; (ii) Explainable AI (XAI) modules, which translate model outputs into interpretable recommendations, supporting trust, operator training, and regulatory acceptance [94, 18]; (iii) Regulatory sandboxes, i.e., controlled test environments under regulator oversight that allow safe experimentation with temporarily relaxed compliance obligations [93, 95, 96]; (iv) Outcome-based procurement, in which utilities or public agencies pay for achieving performance targets (e.g., energy savings, effluent compliance) rather than buying prescriptive solutions, thus incentivising innovation [97]; and (v) Shared data platforms, to enable secure data exchange among utilities, researchers, and technology providers, improving model training and benchmarking.

Public funding, such as Horizon Europe collaborative projects, national innovation funds, and emerging EU initiatives like EIT Water, can de-risk adoption by subsidising pilots and covering part of capital costs [98, 99]. Capacity-building initiatives, including utility–research partnerships, summer schools, and living labs, help close the skill gap in data science and control engineering [100]. Together, these measures align niche innovations with regime needs and landscape pressures, in socio-technical transitions, enabling promising digital innovations to mature and integrate into the existing operational regimes under increasing sustainability and regulatory pressures [30].

2.1.6. Conclusions

To conclude, this dual-track literature review, it can be said that the state-of-the-art of technology clearly shows promising results for the application of AI and RL-based controllers in critical water infrastructure. However, there is a big simulation-to-reality gap to address, which hinders progress and trust [75].

On the other hand, Europe has built a robust but demanding regulatory framework to safeguard compliance and raise the bar for innovation. While the AI Act, GDPR, Data Act, and UWTD recast provide clarity on rights, responsibilities, and sustainability targets, they also introduce compliance costs and legal uncertainty that can slow adoption. Organisational factors - such as the ambidexterity dilemma, digital skills gaps, and procurement rigidity - further compound this challenge. However, the literature also shows possible solutions: regulatory sandboxes, outcome-based procurement, explainable AI modules, and cross-utility data platforms can align incentives, de-risk experimentation, and foster trust. In short, governance is not only a barrier but also a lever that, if properly designed, can accelerate the safe and responsible integration of AI into water infrastructure.

2.2. Identified Knowledge Gaps and Novelty of the Research

As identified by the previous literature review, nowadays, even if advancements in AI are bringing potentially large theoretical opportunities for optimising energy and process, many WWTPs continue to have basic reactive control systems (Relay or PID controllers) that fail to adapt efficiently to real-time inflow variability in their characteristics, pollution loads, and climatic conditions [101].

Therefore, this study aims to understand why this is happening.

From a *technical perspective*, RL for aeration is still novel. While recent studies show promising efficiency and compliance results, the field remains dominated by applications in simulation environments [5, 75], with limited work on transferring these results to real, full-scale plants. The application of deep RL algorithms for aeration control targeting NH_4^+ removal is still limited, and its implementation on real treatment stations is highly relevant for research. There is a need for approaches that train and validate RL agents on real historical operational data and benchmark their performance against existing control strategies under realistic disturbances, sensor noise, and actuator constraints [74, 6, 7].

From a *policy perspective*, there is little empirical evidence of how European utilities interpret and implement AI Act obligations in day-to-day operations [16, 17], how they manage liability and risk allocation, and how national strategies (e.g., France's Plan Eau, Italy's PNRR, the EU Water Resilience Initiative) influence innovation adoption [13, 14]. The system's preparedness for AI control in critical infrastructure has not been clearly explored. More research is required to safely implement these new control strategies, and a structured assessment with socio-technical transition frameworks has not been performed before. Applications of socio-technical transition frameworks such as the Transition Model Canvas (TMC) to the digitalisation of water infrastructure remain rare and underdeveloped [30]. Additionally, the literature lacks empirical evidence on how *organisational capabilities* (digital maturity, skills, procurement, investment capacity) mediate the adoption of high-risk AI controllers.

Together, these gaps strengthen the need for an integrated socio-technical analysis: one that connects technical feasibility, organisational readiness, and regulatory implementation. This research responds to these gaps by combining real data-driven RL development with interviews, multi-actor governance analysis, and the Transition Model Canvas to understand how AI control can realistically scale within European water infrastructure.

This thesis positions itself precisely in this frontier: it develops and benchmarks a deep RL controller for aeration trained on historical plant data, contributing to closing the simulation-to-reality gap, and complements it with a structured system assessment using interviews and TMC analysis [30]. By combining these two tracks, this research provides a crucial contribution to both worlds, going beyond the ”simple” proof of technical feasibility, aiming also to identify adoption pathways and policy recommendations to safely and effectively scale AI-based control across European water infrastructure, moving from research to practice.

3

Problem Statement, Research Questions, and General Approach

3.1. Problem Statement & Research Objectives

As seen in Chapter 2, even if, in theory, advancements in AI technologies are very promising, in practice, many WWTPs still have basic reactive control systems that fail to adapt efficiently to real-time inflow variability [101]. The literature review of the previous chapter highlights a clear mismatch between what is shown in academic papers and what is actually done in practice.

Therefore, the goal of this research is to understand from a holistic perspective why AI is not widely adopted to control processes in water infrastructure. Is the technology not good enough? Is AI unable to safely control water treatment processes due to its probabilistic nature? Or is the problem related to the sociotechnical system in which this innovation should be implemented? Is there a lack of trust in AI-based control systems among stakeholders? If AI-based control is so beneficial, how is it possible to push its adoption? Is this an engineering problem or a strategic and governance challenge?

All these questions have been condensed, refined, and summarised in the section below, which outlines a logical path to conduct the research.

3.2. Research Questions

Given the research objective, this study aims to answer the following overarching research question:

How can AI-based solutions be technically effective, economically viable, safe, and acceptable for sustainable water management in Europe?

To better address the primary question, the research examines the following five sub-questions:

1. *Given the state of the art of AI, is it technically feasible to control wastewater treatment processes with AI solutions in real facilities?*

   The goal of this first sub-research question is to apply state-of-the-art AI-based control to real facilities. This aims to deeply explore the technical side of this challenge, moving from applications on simulators towards training with real operational data. Additionally, this part allows one to gain knowledge on the key challenges and methods in AI-based control, resulting in a better and conscious understanding of the policy suggestions proposed in the conclusions.

2. *How effective are the proposed AI solutions in improving energy efficiency and pollutants control in WWTPs?*

   After evaluating the feasibility, it is necessary to assess whether these control systems can produce improvements in both energy efficiency and environmental sustainability. Therefore, a benchmarking study has to be conducted to compare the proposed solutions among each other and against traditional control methods.

3. *What organisational, economic, and infrastructural conditions are needed to implement AI control systems at a national scale in European countries?*

   With this research sub-question, the focus shifts towards organisational, economic, and infrastructural factors which might hinder the adoption of AI-based control systems. Investigating these is key to conducting a rigorous system assessment and producing a path to push innovation safely.

4. *How do regulators, plant operators, and engineers perceive the risks, benefits, and acceptability of AI in wastewater operations?*

   Regarding safety, since the application of AI-based control is evaluated on critical infrastructures, a special focus is mandatory. Understanding real and perceived risks can shed light on the problem and help understand resistance to adoption.

5. *How is AI-based control positioned within the wastewater sector, and which strategies and policy instruments would most effectively shift the system toward large-scale adoption?*

   Finally, the ultimate goal would be to assess the system status toward this innovation with an existing sociotechnical framework to propose effective policies to push large-scale adoption.

3.3. General Approach

It should now be clear from how the questions are structured that these reflect the dual-track of the research.

The first two sub-questions are more technical, reflecting the core engineering problem, and so they will involve modelling, developing, and benchmarking new AI-based solutions for processes in WWTPs.

On the other hand, the second two sub-questions aim to evaluate the system's preparedness to embrace and implement AI solutions in wastewater and understand how this innovation is perceived. Therefore, these will implement qualitative methods such as interviews.

In conclusion, the final sub-question integrates key insights from the preceding sections, framing the innovation within the broader context of environmental and sustainability objectives and analysing the sociotechnical transition.

Table 3.1 shows a summary of the research sub-questions together with the method used to address them. This is, however, just a high-level overview of the general methodology and should not be intended as a substitute for the two following chapters (Chapter 4 and Chapter 6), which outline in great detail how research was carried out specifically for each part.

No.	Subquestion	Research Method
1	Given the state of the art of AI, is it technically feasible to control wastewater treatment processes with AI solutions in real facilities?	Modelling and Software development; Semi-structured interviews
2	How effective are the proposed AI solutions in improving energy efficiency and pollutants control in WWTPs?	Benchmarking against traditional methods and Performance analysis
3	What organisational, economic, and infrastructural conditions are needed to implement AI control systems at a national scale in European countries?	Literature review, Semi-structured interviews with technical experts, politicians, regulators, and plant operators.
4	How do regulators, plant operators, and engineers perceive the risks, benefits, and acceptability of AI in wastewater operations?	Semi-structured interviews with technical experts, politicians, regulators, and plant operators.
5	How is AI-based control positioned within the wastewater sector, and which strategies and policy instruments would most effectively shift the system toward large-scale adoption?	Sociotechnical transition analysis (TMC), Multi-Level Perspective (MLP), and Synthesis findings.

Table 3.1: Research Subquestions and Corresponding Methods

4

Research Methodology: Engineering Approaches

4.1. Motivation of Methodological Choices

The methodological decisions underlying the engineering part of this thesis are directly derived from the first two research sub-questions formulated in Chapter 3. To answer whether AI-based control can be technically feasible for wastewater treatment processes and whether it can improve energy efficiency and pollutant removal compared to traditional control approaches, it was necessary to design a workflow capable of reproducing the operational complexity of real wastewater treatment plants while allowing systematic testing of advanced control strategies.

Model-based and data-driven approaches, as described in the following sections of this Chapter, were therefore selected. From a computational standpoint, the adoption of Python (open-source) and MATLAB is motivated by the fact that these are widely used scientific toolchains, which can ensure easy reproducibility of the research. The specific libraries adopted (such as `stable_baselines3`, `gymnasium`, and `torch`) and their purpose are summarised in Table B.1, in Appendix B. Finally, the reinforcement-learning setup was chosen to mirror the real-world trade-offs faced by plant operators.

Such an approach is aligned with the engineering goals of improving process performance while also providing essential knowledge for the following Chapters (Chapter 7 and Chapter 8).

4.2. Data Processing

As for every data science project, the first part of the methodology is related to data processing. The subsections below describe how the raw operational data of a real full-scale WWTP have been processed. After this first step, the cleaned data have been used to create training environments for the controller and to generate a buffer of historical actions and rewards to allow for warm-start training.

4.2.1. Dataset Overview

Due to non-disclosure agreements signed before the beginning of this project, the actual data cannot be shared. However, a qualitative description is now provided.

The dataset consists of high-frequency time-series data collected from the aeration tank's sensors that capture the dynamics of the biological nitrification process, including key variables such as ammonium concentration, dissolved oxygen, redox potential, and nitrate, together with operational and environmental variables such as aeration airflow, inflow rate, temperature, pH, and rainfall. Laboratory ammonium measurements, available at a lower frequency, complement the online sensors and are used for validation purposes.

Each observation corresponds to a $5\,\mathrm{min}$ sampling interval, even if various features were actually sampled each $15\,\mathrm{min}$, resulting in several `NaN`s and motivating the resampling explained after. Additional contextual variables were engineered to represent external influences, such as electricity tariff categories, allowing integrated analysis of process efficiency, energy usage, and environmental disturbances.

4.2.2. Data Cleaning and Feature Engineering

The raw dataset underwent a structured multi-stage cleaning and feature-engineering workflow implemented in Python (`pandas`, `numpy`, and `plotly`). The goal was to obtain a temporally aligned, physically consistent, and gap-free dataset suitable for process modelling and control-oriented simulation.

Step 1: Temporal alignment and resampling. All timestamps were converted to a unified datetime format and resampled to a regular $15\,\mathrm{min}$ interval, because, even if the full dataset has a $5\,\mathrm{min}$ time steps, sensors data were collected at a lower frequency (every $15\,\mathrm{min}$).

Step 2: Data aggregation and forward-filling. Rainfall data, originally recorded every $5\,\mathrm{min}$, were aggregated to $15\,\mathrm{min}$ totals. Sparse laboratory ammonium measurements were forward-filled without interpolation to maintain consistency.

Step 3: Quality filtering and range validation. Columns were pruned to retain only those relevant for process modelling. Physically implausible or incomplete records were removed.

Step 4: Data trimming. Data preceding a defined cutoff date were excluded to remove a known period affected by sensor recalibration and maintenance (which was clearly evident even by looking at the time series plots). This ensured that only stable and reliable measurements were retained.

Step 5: Feature engineering. Several derived and contextual features were generated:

- Cyclical encodings (sine/cosine) of hour, weekday, and day-of-year to capture periodic behaviour.
- A dry-weather duration counter, expressing elapsed hours since the last rainfall, used to capture ammonium spikes caused by sudden rains after a long dry period.
- A dynamic electricity price indicator, derived from time-dependent tariff rules distinguishing between peak and off-peak hours.
- Lag features for key process variables (inflow, airflow, ammonium, oxygen, and redox) at 15, 30, 60, and 90-minute delays to provide historical context.
- A binary ON/OFF indicator for aeration activity, computed by comparing airflow to the minimum operational threshold defined in the plant's operational constraints.

Step 6: Final consistency checks. All numeric columns were forward-filled to remove residual gaps, and the initial rows affected by lag creation were trimmed. The cleaned dataset was visually inspected through time-series plots and histograms to confirm range consistency and absence of artefacts.

The resulting dataset (`cleaned_full.csv`) constitutes a harmonised, feature-enriched representation of the WWTP operation.

Training and Testing Split

To evaluate data-driven models' generalisation, the cleaned dataset was chronologically divided into two non-overlapping subsets:

- Training set (`train_data.csv`): the initial portion of the time series, used for parameter fitting, feature selection, and model calibration.
- Test set (`test_data.csv`): a later, unseen period reserved for independent validation.

This temporal split preserves the natural ordering of observations and prevents information leakage between training and testing, reproducing realistic forecasting and control conditions.

4.3. Controller Objectives

As mentioned in Chapter 2, recent advances in multi-objective reinforcement learning (MORL) and Deep RL have demonstrated the potential of learning-based approaches to manage conflicting objectives in the water sector. In this work, the focus shifts from strategic basin-scale management to real-time operational control. Here, a Deep RL controller is trained to optimise multiple conflicting goals simultaneously: maintaining effluent quality, ensuring biological stability, and minimising energy consumption. Even if the temporal and physical scales differ substantially from those in [65], both works share the same principle of using reinforcement learning to optimise complex and critical water systems.

In the specific case study considered at SUEZ, the controller is designed to keep key process variables such as dissolved oxygen (DO), oxidation-reduction potential (ORP) and, optionally, ammonium (NH_4^+, not always taken into account since most stations do not measure this variable given the high costs of sensors), within optimal operational ranges to ensure stable nitrification and denitrification dynamics and compliant effluent quality. At the same time, the controller is aware of energy prices (in the sense that current and future energy prices are included within the observation space) and aims to not over-oxygenate the tank, minimising wastes while respecting the pump operational constraints (e.g. avoid turning on and off continuously, do not leave the pump inactive or active for too long/short). By observing current and lagged process states plus short-term future energy price information (observation), the controller continuously adjusts the aeration airflow (action). These high-level objectives are better explained in the two reward functions in Section 4.6.4, because their formal definition depends on the station's constraints (e.g. if the goal is to prioritise energy savings over water quality or vice versa).

4.4. Modelling the Dynamics of the Aeration Process

Before diving deep into the architecture of the deep reinforcement learning agent and how the control goals mentioned above will be used to define a reward function, it is necessary to think about the training process. Training a deep reinforcement learning model often requires an interactive model in which the agent can explore the decision space to understand rewards and outcomes, without compromising the quality of the effluents on real stations.

It is important to underline that, given the goal of bridging the gap between simulation and reality, the proposed models must be validated on real historical data and fit to station-specific parameters. Therefore, various approaches have been tried for this scope: getting a model able to calculate new ”future states” in a physically reliable manner so that the agent can understand the process and how it evolves based on its decision.

Apart from the model based on ordinary differential equations (ODEs), explained in Chapter 5 (Section 5.1.2), which was initially used also for training but resulted in extremely slow runs, probably due to the chosen integration method to solve the nonlinear system (`RungeKutta4`), a few data-driven approaches have been tried. Among these, the most relevant are:

- the Linear Regression Model;
- the ExtraTree Regression Model;
- and the LSTM-ATT Model.

All models were trained and validated on the same dataset. As expected, the ExtraTree and LSTM-ATT models achieved higher predictive accuracy given the complexity of the process.

However, it is important to align model selection with the training objective. For an initial high-level pretraining phase, a simple and computationally efficient model is preferred to estimate the next system state based on the agent's action. In this context, the Linear Regression Model provides the best trade-off, allowing the agent to learn the core dynamics of the process and its operational constraints.

During early training episodes, the controller frequently attempts extreme actions, such as continuously pumping at maximum power or keeping the pump completely off. These are far from realistic operating conditions observed in the plant's historical data. Because these situations never occur in the dataset, models like ExtraTree or LSTM-ATT cannot accurately represent such out-of-distribution behaviour. Relying on them during early training could therefore mislead the agent, causing it to infer, for instance, that constant pumping still leads to stable tank conditions, which is physically incorrect.

On the other hand, even though linear models are simple, they do extrapolate linearly beyond the data, which is much safer for the agent's exploration, and they also run orders of magnitude faster than LSTM or tree ensembles, which is crucial for deep RL training loops.

Therefore, the Linear Model (Section 4.4.1) has been selected to calculate the next state of the system. ExtraTrees and LSTM-ATT models, and their validation tests, are still available in Appendix B, and will be considered for future developments, such as further and more advanced training phases.

4.4.1. Linear Regression Model for Aeration

Three linear regression models (Equations 4.1–4.3) were fitted for dissolved oxygen (DO), oxidation–reduction potential (ORP), and ammonium concentration (NH_4).

$$\mathrm{DO}_{t+1} = \mathrm{DO}_t + \beta_0 + \beta_1 Q_t + \beta_2 F_t \tag{4.1}$$

$$\mathrm{ORP}_{t+1} = \mathrm{ORP}_t + \beta_3 + \beta_4 Q_t + \beta_5 F_t \tag{4.2}$$

$$\mathrm{NH}_{4,t+1} = \mathrm{NH}_{4,t} + \beta_6 + \beta_7 Q_t + \beta_8 F_t \tag{4.3}$$

Each variable at time $t+1$ is expressed as a linear function of its current value, the normalised blower airflow $Q_t \in [0,1]$, and the influent flow rate F_t. In all three models, the β_i coefficients are obtained by fitting the models to historical plant data.

4.5. Deep Reinforcement Learning Architectures

Deep RL agents can be distinguished by their *architecture*. For continuous control problems, such as regulating aeration flow rates in wastewater treatment plants (WWTPs), *actor–critic* methods have proven particularly effective. They combine the stability of value-based approaches with the expressiveness of policy-based learning [102, 103]. Within this family, two algorithms are particularly relevant: the *Soft Actor-Critic* (SAC) and the *Twin-Delayed Deep Deterministic Policy Gradient* (TD3). Both have been applied to WWTP optimisation [74, 5, 6] and form the technical basis of the present research.

4.5.1. Soft Actor-Critic (SAC)

The Soft Actor–Critic (SAC) algorithm [104] is an off-policy, model-free actor–critic method based on the *maximum entropy* framework. Unlike traditional RL algorithms that only maximise expected cumulative reward, SAC also encourages exploration by including a policy entropy term in its objective [105]:

$$J(\pi) = \mathbb{E}_{\tau\sim\pi}\Big[\sum_{t=0}^{T}\big(r(s_t,a_t)+\alpha\mathcal{H}(\pi(\cdot|s_t))\big)\Big], \tag{4.4}$$

where α controls the exploration-exploitation trade-off, and $\mathcal{H}(\pi(\cdot|s_t)) = -\mathbb{E}_{a_t\sim\pi}[\log\pi(a_t|s_t)]$. All terms appearing in Eq. 4.4 are defined and intuitively described in Table 4.1.

Table 4.1: Definitions and intuitive explanation of the terms in the SAC objective (Eq. 4.4).

Term	Definition	Explanation
$\pi(\cdot\|s_t)$	Stochastic policy	Outputs a Gaussian distribution over actions given the state s_t, enabling exploration.
$r(s_t, a_t)$	Reward at time t	Encodes task objectives (e.g., ORP regulation, energy minimisation). Guides learning.
$\mathcal{H}(\pi(\cdot\|s_t))$	Policy entropy	Encourages exploratory behaviour by rewarding high-entropy action distributions.
α	Temperature parameter	Balances exploitation and exploration; higher values promote exploratory behaviour.
τ	Trajectory sampled from π	Sequence of states and actions collected during an episode.
T	Episode horizon	Number of time steps in each rollout.

SAC uses two critic networks Q_{θ_1} and Q_{θ_2} to mitigate overestimation bias through the clipped double-Q target:

$$y_t = r_t + \gamma\,\mathbb{E}_{a_{t+1}\sim\pi}\Big[\min_{i=1,2} Q_{\bar{\theta}_i}(s_{t+1}, a_{t+1}) - \alpha\log\pi(a_{t+1}\mid s_{t+1})\Big], \tag{4.5}$$

where the target-network parameters $\bar{\theta}_i$ provide a slowly updated and stabilised estimate of the next-state value. The actor parameters ϕ are then updated to maximise the entropy-regularised objective:

$$L_\pi(\phi) = \mathbb{E}_{s_t\sim D,\,a_t\sim\pi_\phi}\Big[\alpha\log\pi_\phi(a_t\mid s_t) - \min_{i=1,2} Q_{\theta_i}(s_t,a_t)\Big], \tag{4.6}$$

which encourages exploratory policies while steering actions toward regions of high estimated value.

All symbols appearing in Eqs. 4.5 and 4.6 are defined and intuitively explained in Table 4.2.

This entropy-regularised approach allows SAC to perform robustly in the presence of stochastic disturbances or unmodelled dynamics, typical of wastewater processes [5, 6]. Its architecture, composed of a stochastic Gaussian actor, twin critics, target networks, and an adaptive temperature parameter, has been implemented and validated in open-source frameworks such as `Stable Baselines3` [106].

Table 4.2: Definitions and intuition behind the terms in the SAC critic-target update (Eq. 4.5) and actor update (Eq. 4.6).

Term	Definition	Explanation
$Q_{\theta_1}, Q_{\theta_2}$	Critic value networks	Provide two independent value estimates to mitigate overestimation bias through the "min" operator.
$\bar{\theta}_i$	Target-network parameters	Slowly updated copies of the critics; stabilise training by providing consistent bootstrap targets.
γ	Discount factor	Determines the relative weight of future rewards; promotes long-term control performance.
$a_{t+1} \sim \pi(\cdot \mid s_{t+1})$	Action sampled from the stochastic actor	Ensures the target incorporates the entropy-regularised next action according to the learned policy.
$\log \pi(a \mid s)$	Log-probability of selected action	Penalises low-entropy policies; maintains exploration.
y_t	Bootstrapped one-step target	The value that the critic learns to approximate; combines reward, bootstrapped value and entropy term.
D	Replay buffer	Stores transitions for off-policy learning; breaks correlation in data and improves sample efficiency.

4.5.2. Twin-Delayed DDPG (TD3)

The Twin-Delayed Deep Deterministic Policy Gradient (TD3) algorithm [107] improves upon the DDPG framework [108] through three mechanisms:

1. *Clipped Double-*Q *Learning*: two critic networks are trained, and the target uses $\min(Q_1, Q_2)$ to reduce overestimation bias;
2. *Target Policy Smoothing*: Gaussian noise is added to the target policy action to improve robustness;
3. *Delayed Policy Updates*: the actor is updated less frequently than the critics to stabilise learning.

Formally, the critic loss is:

$$L_Q(\theta_i) = \mathbb{E}_{(s_t,a_t,r_t,s_{t+1})\sim D}\Big[Q_{\theta_i}(s_t, a_t) - y_t\Big]^2, \quad y_t = r_t + \gamma\, Q_{\bar{\theta}}\big(s_{t+1},\, \mu_{\bar{\phi}}(s_{t+1}) + \epsilon\big), \tag{4.7}$$

where the actor μ_ϕ is deterministic, and target smoothing is implemented through the noise term ϵ. The critic loss $L_Q(\theta_i)$ represents the mean-squared error between the current critic prediction $Q_{\theta_i}(s_t, a_t)$ and the *bootstrapped* target y_t. The target is termed bootstrapped because it combines the observed immediate reward r_t with an estimate of future value obtained from the target critic network $Q_{\bar{\theta}}$, rather than waiting for the full return to be observed. This allows the algorithm to update the value function using its own predictions of future returns, improving learning efficiency and reducing variance.

The notation $(s_t, a_t, r_t, s_{t+1}) \sim D$ indicates that transitions are sampled from the replay buffer D, which stores experience collected by the agent (or is filled with past transitions as in the case of the "warm-start"). In practice, the expectation in Eq. 4.7 is approximated by drawing a minibatch $\tilde{D} \subset D$, enabling decorrelated updates and efficient stochastic gradient descent.

All terms appearing in Eq. 4.7 are summarised in Table 4.3.

TD3 has shown excellent performance in continuous process control, including recent applications to WWTP aeration optimisation [74, 5]. Its deterministic policy provides stable and energy-efficient operation when the process dynamics are well modelled, making it a valuable comparison baseline to SAC. As in the case of SAC, TD3 has also been implemented in `Stable Baselines3`.

Table 4.3: Definitions and intuitive explanation of the terms in the TD3 critic update (Eq. 4.7).

Term	Definition	Explanation
Q_{θ_1}, Q_{θ_2}	Critic value networks	Two independent Q-functions used to compute a conservative target via $\min(Q_1, Q_2)$, mitigating overestimation bias.
$\mu_\phi(s)$	Deterministic policy (actor)	Produces a single continuous action for each state; avoids stochastic sampling, improving control stability.
$\bar{\theta}$, $\bar{\phi}$	Target-network parameters	Slowly updated copies of the actor and critics; provide stable targets for bootstrap updates.
ϵ	Target policy smoothing noise	Small clipped Gaussian noise added to the target action to prevent exploiting narrow value-function peaks and to improve robustness.
y_t	Bootstrapped one-step target	Combines immediate reward and discounted next-state value using the smoothed target action.
γ	Discount factor	Controls the weight of future rewards in the critic update.
D	Replay buffer	Stores transitions for off-policy learning and reduces temporal correlation between samples.

4.5.3. Architectural Comparison

Table 4.4: Comparison between SAC and TD3 architectures.

Feature	SAC	TD3
Policy type	Stochastic (Gaussian)	Deterministic
Entropy regularisation	Yes (α temperature)	No
Critics	Two (clipped double-Q)	Two (clipped double-Q)
Target policy smoothing	No	Yes (Gaussian noise)
Policy update frequency	Every iteration	Delayed
Exploration mechanism	Entropy maximisation	Action noise
Robustness	High under stochasticity	High under deterministic dynamics

Both SAC and TD3 share an off-policy actor–critic foundation with experience replay and target networks. Their differences, summarised in Table 4.4 reflect complementary strengths. In WWTP applications, SAC enables adaptive control under fluctuating inflows and noise, while TD3 offers precise control under steady-state conditions.

4.6. Deep Reinforcement Learning for Aeration Optimisation

4.6.1. Custom Environment for Aeration Control

A custom Python environment (`AerationEnvLinear2`), compatible with the `Gymnasium` interface was developed to simulate the aeration basin dynamics under different control actions. This environment encapsulates a data-driven model of the basin and enforces realistic operational constraints. Key features of the environment include realistic blower (pump) behaviour with hard constraints and inertia, enforcement of minimum ON/OFF durations via a pump state tracker, centralised input/output scaling, linear state-transition models derived from regression, and exogenous profiles for inflow and energy price. The environment loads historical plant data from a CSV file (e.g. time-series of dissolved oxygen, ORP, inflow, etc.) and regression coefficients from Excel files for the process models. Each episode corresponds to a 24-hour period (96 steps at 15 min per step), initialised from a random start point in the dataset. The environment can either use the next timestep's inflow and electricity price from the internal dataset or accept user-provided profiles for these exogenous variables, enabling evaluation under specific scenarios (e.g. a given inflow pattern or time-of-day electricity tariff). All observations and actions are normalised to $[0, 1]$ using predefined scaling ranges, ensuring the learning agent operates on nondimensional inputs. Internally, physical quantities (dissolved oxygen in mg/L, oxidation-reduction potential (ORP) in mV, flow rates in Nm^3/h, etc.) are scaled based on historical min-max or statistical ranges so that the observation vector and actions reside in a unit hypercube. The environment's state-transition model advances the water quality states (DO and ORP, with ammonia NH4 optionally included) using linear regression models fitted to historical data as explained previously. At each time step t, the next state is computed as a linear function of the current state and the control action. In particular, the change in dissolved oxygen $\Delta\mathrm{DO}$ is given by a linear combination of a constant term, the current aeration rate, and the influent flow rate (among other possible features), using coefficients obtained from the regression file. Similar linear update equations are applied for ORP (redox potential) and for ammonia when enabled. This essentially forms an approximate linear state-space model: for example, $\mathrm{DO}t + 1 = \mathrm{DO}_t + \beta_0 + \beta_1\tilde{Q}_t + \beta_2\tilde{F}_t$, where $\tilde{Q}$ is the normalized blower airflow and $\tilde{F}$ is the influent flow, and β_i are regression coefficients. By using a data-driven linear model, the environment can emulate the qualitative dynamics of the aeration tank (e.g. aeration increases DO, high inflow can reduce DO, etc.) while remaining computationally light for RL training.

4.6.2. Observation Space

The observation at each time step is a vector of 10 features (in the configuration used, with ammonia disabled). This observation state is defined as follows (all variables are normalised to [0,1] ranges):

- Dissolved oxygen (DO) at current step: the current DO concentration in the basin (mg/L), scaled between a minimum and maximum typical range.
- Oxidation-reduction potential (ORP) at current step: the current ORP (mV) in the basin, scaled to [0,1].
- Influent flow rate at current step: the current influent flow (e.g. m^3/h), providing the load to the basin, normalised by the expected flow range.
- Current electricity price: the unit energy cost at the current time step (in €/kWh), normalised. This represents time-varying energy tariffs used to incentivise shifting aeration to cheaper periods.
- DO at previous step: the lagged DO measurement from the end of the last time step (normalised in the same way as current DO). This provides the agent with velocity/trend information.
- ORP at previous step: the ORP value from the previous time step (normalised), included for the same reason as above.
- Influent flow at previous step: the last-step inflow rate, capturing recent changes in load.
- Consecutive ON-duration counter: the fraction of time steps the blower has been continuously ON, scaled between 0 and 1. This is computed as $\min(n_{on}, MAX_{ON})$, where n_{on} is the count of successive steps with the blower on, and MAX_ON is a specified upper bound (e.g. 20 steps). This gives the agent a sense of how long the aerator has been running.
- Consecutive OFF-duration counter: similarly, the fraction $\min(n_{off}, MAX_{OFF})$ representing how long the blower has been continuously OFF.
- Electricity price at next step: the known price for the upcoming time step (normalised). In many practical scenarios, the energy price schedule is known ahead (e.g. day-ahead market prices), so including the next period's price informs the agent of future cost changes.

This observation space design provides the agent with both the current process state and recent historical context (previous-step values), as well as information about how long the actuator has been in its current mode and the immediate future cost signal. By including lagged state variables, the agent can infer the time derivatives or trends (e.g. whether DO is rising or falling), which is important for controlling a process with time lags. The pump ON/OFF counters embed the operational context (e.g. if the blower has been off for many steps, there may be an oxygen deficit or an impending need to aerate, whereas a long ON streak might allow a forthcoming rest). All these features are normalised, and the observation vector is a continuous 10-dimensional vector $o_t \in [0,1]^{10}$. (In a variant where ammonia is included as well, the observation space would expand to 12 dimensions accordingly, because the ammonia lag would also be included.)

4.6.3. Action Space and Actuator Constraints

The action is a single continuous command $a_t \in [0, 1]$ representing the blower air flow rate as a fraction of its operating range. This scalar action is linearly rescaled into a physical airflow Q in Nm^3/h between a minimum $Q_{\min}$ and maximum $Q_{\max}$ capacity of the blower. For example, $a_t = 0$ corresponds to $Q = 0\ Nm^3/h$ (blower off), and $a_t = 1$ corresponds to $Q = 2000\ Nm^3/h$ (maximum blower airflow). However, the real pump, even if it emits air with a gradual ramp that starts from 0, is not free to operate below a minimum airflow. Specifically, the environment defines a minimum sustainable airflow $Q_{min,run}$ (e.g. about $1500\ Nm^3/h$) such that any commanded flow below this threshold will cause the blower to shut off. In implementation, after converting a_t to a raw commanded flow, a hard constraint is applied: if $Q_{commanded} < Q_{min,run}$, then the target airflow is set to zero (blower OFF); otherwise, the target airflow is at least $Q_{min,run}$. This creates a hybrid action behaviour: the agent's continuous action effectively decides between "off" (if it outputs a very low number) or a range of "on" levels between the minimum and maximum airflow. The linear mapping and this strategy ensure that the agent can explore intermediate aeration rates but will not attempt to operate the blower in an unstable low-flow regime. In addition to the minimum flow threshold, the environment enforces rules about how frequently the blower can switch on or off, implementing a form of actuator inertia and operational policy. A `PumpTrackerV2` utility keeps track of the consecutive ON steps and OFF steps of the blower. The environment uses this to impose minimum ON/OFF durations: if the blower was just turned off, it must remain off for at least a certain number of steps before restarting; similarly, once turned on, it must stay on for a minimum duration. For example, if the minimum off duration is 2 steps, and the agent attempts to start the blower earlier, the environment will override the action to keep the blower off until the constraint is satisfied. Likewise, if the blower is on and has not yet met the minimum on-duration, any action requesting an off (or below threshold) will be disregarded and the blower kept running at least at $Q_{min,run}$. These rules prevent rapid flipping of the actuator (which could damage equipment and cause inefficient operation) and are consistent with real aeration control strategies that avoid short-cycling the blowers. They are implemented by adjusting the target airflow based on the pump state: e.g. "if last state was OFF and off-streak $< MIN_{OFF}$ steps, enforce target $Q = 0$; if last state was ON and on-streak $< MIN_{ON}$, enforce $Q \geq Q_{min,run}$." Finally, the actuator dynamics include a first-order lag to simulate the blower's inertia. The actual delivered airflow Q_{phys} cannot change instantaneously; instead, it moves towards the target command with a certain time constant τ. This is modelled as $Q_{phys,t+1} = Q_{phys,t} + \frac{\Delta t}{\tau}(Q_{target,t} - Q_{phys,t})$, which in discrete form is equivalent to an exponential smoothing. In the implementation, τ is set to a small value (on the order of minutes) based on the known characteristics of the blower. At each step, the environment updates the physical airflow by $Q_{phys} \leftarrow \alpha, Q_{phys} + (1 - \alpha), Q_{target}$, where $\alpha = \exp(-\Delta t/\tau)$ for $\Delta t = 15$ min. This simulates what happens in reality: when a large airflow is commanded, the blower takes a few steps to ramp up, and similarly ramps down when turned off. Importantly, if the target is zero (blower off) and the filtered result falls below $Q_{min,run}$, the flow is snapped to zero immediately (meaning once the blower effectively spins down near the cutoff, it goes fully off). Conversely, when ramping up, the flow is not allowed to dip below the minimum running flow if the blower is meant to be on. These measures produce a realistic, smooth actuator behaviour in the simulation.

4.6.4. Reward Function Design

The reinforcement learning agent's reward r_t at each time step is formulated to capture multiple (sometimes competing) objectives as introduced in Section 4.3. These goals are translated into custom reward functions that aggregate several penalty and bonus terms. In both cases, the reward can be viewed as a weighted sum of objectives, where the weights can be tuned to reflect the relative importance of each aspect. In this work, two related reward formulations are used: the first one prioritises water quality, while the second one places stronger emphasis on energy saving. The reward weights and the type of penalty/bonus have been defined empirically and iteratively, based on the desired goal (better water quality or higher cost efficiency), and the control behaviour showed after training.

Reward Function 1: Water-Quality-Oriented

The first reward function (`reward_function_1`) has the primary purpose of maintaining high water quality while still discouraging unnecessary energy use and pathological operating patterns. The main components of the reward are summarised below:

- DO overshoot penalty: a penalty for dissolved oxygen exceeding a desired target level. If the DO rises above this setpoint (indicating over-aeration), a quadratic penalty term kicks in, discouraging the agent from pushing DO unnecessarily high.

- ORP deviation penalty: a penalty for ORP straying from an optimal range. ORP (redox potential) is used as an indicator of the chemical environment; extremely low ORP can signal excessive anoxic conditions, while extremely high ORP may indicate over-aeration. The penalty is formulated in a Gaussian manner: it is small when ORP is near the desired level and increases as ORP deviates from this optimum. This provides a smooth negative reward that gently guides ORP toward the desired range without a hard constraint.
- Energy cost penalty: a penalty proportional to the energy consumed by aeration, accounting for real-time electricity price. The product of the blower airflow and the current price is incorporated into the reward. To keep this term bounded, the product is normalised by a reference (the maximum expected cost) so that the energy penalty is in the $[0, 1]$ range. This term pushes the agent to reduce airflow, especially when electricity is expensive, encouraging load shifting to cheaper periods if possible.
- Sub-minimum airflow penalty: a penalty for operating the blower at very low flow rates is added to deter the agent from keeping the blower barely on.
- Switching penalty: a fixed penalty incurred whenever the blower switches state (an ON-to-OFF or OFF-to-ON transition). The moment of switching is detected by the pump tracker (n_{on} or n_{off} resetting to 1). This term penalises frequent switching, addressing the wear-and-tear and inefficiency associated with toggling the blower too often.
- Duration penalties: in addition to the immediate switching penalty, the reward includes penalties for overly long continuous ON or OFF durations. If the blower remains ON beyond a certain soft limit (e.g. approaching the upper limit *MAX_ON_STEPS*), a small penalty grows with each additional step. Similarly, an excessively long OFF period accrues a penalty once it exceeds a threshold. These "smooth" duration penalties are quadratic in the normalised on/off time (which means that they increase gradually as on/off streaks become very long). The intention is to prevent pathologically long periods of no aeration (which could risk process performance) or continuous aeration (which could waste energy), beyond what is reasonably required.
- Off-state reward: a modest bonus is given for keeping the blower OFF during stable conditions. Specifically, if the aerator is off ($Q \approx 0$) *and* the process is in a safe state (e.g. ORP is above a minimum, indicating no severe reducing conditions), the agent accumulates a small positive reward that increases with the duration of the off-period (using a $\log(1 + n_{\mathrm{off}})$ form). This term explicitly encourages the agent to find opportunities to shut off aeration for energy savings when water quality permits, and to keep it off for a meaningful length of time (to avoid trivially turning it off for one step and back on).
- Operational safety bonus: a small bonus is provided for operating in a "balanced" region where both DO and ORP are within acceptable limits and the blower is either fully off or running above the minimum flow. In other words, if the system is in a desirable state (adequate treatment indicated by ORP, no DO overshoot, and no improper low airflow), a constant reward is added. This can be seen as a "living reward" for being in a good operating regime, which helps to slightly favour any action that keeps the process in a safe, efficient state.

The weighted sum of these components results in the instantaneous reward below (Equation 4.8), while Table B.6 in Appendix B contains the list of variables and parameters present in the function, together with a brief explanation for each, which would be redundant in this chapter given the extensive paragraph above, but can still be useful in case a quick identification is needed.

$$
\begin{aligned}
r_t = &- w_{\mathrm{DO}}\ \max(0, DO_t - DO^{\mathrm{set}})^2 - w_{\mathrm{Redox}}\Big[1 - e^{-\frac{1}{2}\left(\frac{ORP_t - ORP^*}{50}\right)^2}\Big] \\
&- w_{\mathrm{Energy}} \frac{p_t/p_{\mathrm{ref}} \cdot Q_{\mathrm{air},t}/Q_{\max}}{1} - w_{\mathrm{LowFlow}}\ \mathbb{1}_{Q_{\text{off}}<Q_{\mathrm{air},t}<Q_{\min}} \left(1 - \frac{Q_{\mathrm{air},t} - Q_{\text{off}}}{Q_{\min} - Q_{\text{off}}}\right)^2 \\
&- w_{\mathrm{Switch}}\ \mathbb{1}_{\text{pump just switched}} - w_{\mathrm{OverOn}}\ \max(0, n_{\mathrm{on}} - N^{\mathrm{on}}_{\max}) - w_{\mathrm{OverOff}}\ \max(0, n_{\mathrm{off}} - N^{\mathrm{off}}_{\max}) \\
&- w_{\mathrm{OnDur}} \left(\frac{n_{\mathrm{on}}}{N^{\mathrm{on}}_{\max}}\right)^2 - w_{\mathrm{OffDur}} \left(\frac{n_{\mathrm{off}}}{N^{\mathrm{off}}_{\max}}\right)^2 \\
&+ w_{\mathrm{Bonus}}\ \mathbb{1}_{\text{safe region}} + w_{\mathrm{Living}}\ \mathbb{1}_{\text{safe region}} + w_{\mathrm{OffReward}}\ \log(1 + n_{\mathrm{off}}).
\end{aligned}
\tag{4.8}
$$

Weights are summarised in Table 4.5, reflecting the relative importance of each goal.

Term	Symbol	Weight	Description
DO penalty	w_{DO}	0.8	Penalizes over-aeration above 1.9 mg/L
Redox penalty	w_{ORP}	0.6	Penalizes deviation from 200 mV
Energy cost	w_{E}	0.003	Penalizes energy use
Low-flow penalty	w_{L}	2.0	Penalizes half-on blower states
Switching penalty	w_{S}	0.03	Penalizes ON/OFF toggling
Over-ON duration	w_{ON}	1.2	Penalizes excessively long ON periods
Over-OFF duration	w_{OFF}	0.6	Penalizes excessively long OFF periods
OFF reward	w_{R}	0.5	Rewards stable safe OFF operation
Balanced-operation bonus	w_{B}	1.2	Rewards efficient stable operation
Stability reward	w_{St}	0.05	Encourages persistence in stable regime

Table 4.5: Reward weights defined in `reward_weights_5.py`.

The reward is negative for undesirable behaviour (hence called penalties) and positive for the few encouraging terms. Finally, the aggregate reward is clipped to a reasonable range (e.g. [-10, 10]) to prevent excessively large gradients from any single timestep.

Reward Function 2: Energy-Focused Scaled Variant

A second variant of the reward (`reward_function_2`) builds on the same structure but places a stronger emphasis on energy saving.

The main differences in the scaled reward are:

- Hybrid DO penalty: DO overshoot above the scaled target is initially penalised linearly and only becomes quadratic beyond a higher soft threshold. This reduces the aggressiveness of the DO term compared to the purely quadratic penalty in the physical-space reward, while still discouraging severe overshoot. This seems counterintuitive, given the fact that in this case, the goal is to further penalise overaeration. However, this will be clearer by looking at the rest of the function and the weights: the use of energy is so penalised that it is not needed to be excessively strict on DO constraints.
- Asymmetric ORP penalty: deviation of ORP from a desired band is penalised asymmetrically. Too-low ORP (overly reducing conditions) incurs a much stronger quadratic penalty than too-high ORP (overly oxidising conditions). This encodes the operational preference that temporarily high ORP is less problematic than low ORP, which can compromise biological performance.
- Stronger energy penalty: the energy term now directly penalises the (clipped) product of scaled airflow and price, with a much higher weight than in the first reward. This shifts the optimisation focus toward reducing aeration, especially during expensive price peaks, and thereby favours energy-efficient behaviour.

- Duration and switching terms: switching penalties and ON/OFF duration penalties are retained, but their relative weights are tuned so that long OFF periods are more acceptable as long as DO and ORP remain within safe bounds. Smooth quadratic penalties on normalised ON/OFF durations are still present to avoid extreme streaks.
- Off-state reward: a logarithmic OFF reward is again applied when the blower is truly OFF and the state is safe. Its weight is slightly reduced with respect to the first version.
- Balanced-operation bonuses: small bonuses for operating in a "safe region" (acceptable DO and ORP, and either OFF airflow) are preserved but with lower magnitude than in the first reward, reflecting the shifted priority toward energy savings.

The weighted sum of these terms results in the instantaneous reward below:

$$
\begin{aligned}
rV2_t = \; & - w_{\mathrm{DO}}\Big[\mathbb{1}_{\{DO_t > DO^{\mathrm{set}}\}}\big(DO_t - DO^{\mathrm{set}}\big) + \mathbb{1}_{\{DO_t > DO^{\mathrm{high}}\}}\big(DO_t - DO^{\mathrm{high}}\big)^2\Big] \\
& - \mathbb{1}_{\{ORP_t < ORP^{\mathrm{lo}}\}}\, w_{\mathrm{Redox,lo}} \left(\frac{ORP^{\mathrm{lo}} - ORP_t}{0.05}\right)^2 \\
& - \mathbb{1}_{\{ORP_t > ORP^{\mathrm{hi}}\}}\, w_{\mathrm{Redox,hi}} \left(\frac{ORP_t - ORP^{\mathrm{hi}}}{0.05}\right)^2 - w_{\mathrm{Energy}}\; \mathrm{clip}\big(p_t\, Q_{\mathrm{air},t},\, 0,\, 1\big) \\
& - w_{\mathrm{Switch}}\, \mathbb{1}_{\{\text{pump just switched}\}} - w_{\mathrm{OverOn}}\; \max\big(0,\; n_{\mathrm{on}} - N^{\mathrm{on}}_{\max}\big) \\
& - w_{\mathrm{OverOff}}\; \max\big(0,\; n_{\mathrm{off}} - N^{\mathrm{off}}_{\max}\big) \\
& - w_{\mathrm{OnDur}} \left(\min\!\left(\frac{n_{\mathrm{on}}}{N^{\mathrm{on}}_{\max}},\, 1\right)\right)^2 - w_{\mathrm{OffDur}} \left(\min\!\left(\frac{n_{\mathrm{off}}}{N^{\mathrm{off}}_{\max}},\, 1\right)\right)^2 \\
& + w_{\mathrm{Bonus}}\, \mathbb{1}_{\{\text{safe region}\}} + w_{\mathrm{Living}}\, \mathbb{1}_{\{\text{safe region}\}} + w_{\mathrm{OffReward}}\, \mathbb{1}_{\{\text{safe off}\}}\; \log\big(1 + n_{\mathrm{off}}\big).
\end{aligned}
\tag{4.9}
$$

As for `reward_function_1`, Table B.6 in Appendix B reports a description of the variables in the function.

The corresponding weights are defined in a separate configuration file (e.g. `reward_weights_5V2.py`) and are reported in Table 4.6. They reflect the design choice of making energy use comparatively more "expensive" while still protecting against biologically unsafe conditions:

Term	**Symbol**	**Weight**	**Description**
DO penalty	w_{DO}	0.02	Softer DO overshoot penalty (scaled domain)
Redox penalty (high ORP)	$w_{\mathrm{Redox,hi}}$	0.05	Mild penalty for too-high ORP
Redox penalty (low ORP)	$w_{\mathrm{Redox,lo}}$	1.0	Strong penalty for too-low ORP
Energy cost	w_{Energy}	2.0	Dominant term promoting energy savings
Low-flow penalty	w_{LowFlow}	1.5	Penalizes inefficient half-on blower states
Switching penalty	w_{Switch}	0.03	Penalizes ON/OFF toggling
Over-ON duration	w_{OverOn}	1.2	Penalizes excessively long ON periods
Over-OFF duration	w_{OverOff}	0.6	Penalizes excessively long OFF periods
OFF reward	$w_{\mathrm{OffReward}}$	0.2	Rewards stable safe OFF operation (reduced)
Balanced-operation bonus	w_{Bonus}	0.6	Rewards efficient stable operation (reduced)
Stability reward	w_{Living}	0.02	Encourages persistence in stable regime
ON-duration soft penalty	w_{OnDur}	0.2	Smooth penalty on long ON streaks
OFF-duration soft penalty	w_{OffDur}	0.2	Smooth penalty on long OFF streaks

Table 4.6: Reward weights for `reward_function_2`.

As in the first case, the aggregate reward is finally clipped (e.g. to $[-10, 10]$) to avoid excessively large gradients originating from single timesteps.

4.6.5. Agent-Environment Interaction Loop

The controller (agent) interacts with the environment in a closed loop following the standard reinforcement learning cycle. At the beginning of each episode, the environment is reset to an initial state drawn from the historical dataset (random start day/time) with the blower off. The episode then proceeds through a sequence of discrete time steps $t = 0, 1, 2, \ldots, T$ (with $T = 95$ for a 24-hour episode of 96 steps). At each time step, the following occurs:

1. Observe State: the agent receives the current state $\mathrm{s}t$ of the environment, provided as the normalized observation vector o_t described in Section 4.6.2. This includes current sensor readings (DO, ORP, etc.), the previous state, and contextual info (on/off counters, upcoming price).
2. Agent Decision (Policy): based on this observation, the agent's policy π selects an action $a_t \in [0, 1]$. In the case of a trained neural network policy, this involves a forward pass to output the action, possibly with added exploration noise if in training.
3. Apply Action with Constraints: the environment receives a_t and maps it to a physical blower airflow command Qcmd. It then enforces the pump's operational constraints: if Q_{cmd} is below the minimum running flow, it is treated as an OFF command ($Q_{\text{target}} = 0$); otherwise $Q_{\text{target}} = \max(Q_{\text{cmd}},, Q_{\text{min,run}})$. Next, the PumpTracker is consulted to enforce minimum on/off durations: if the blower has not yet satisfied a minimum OFF period, the command is overridden to OFF despite any request to start it, and conversely for minimum ON duration. This yields a possibly adjusted Q_{target} for this step.
4. Update Physical State: the blower's actual airflow Q_{phys} is updated from its previous value towards Q_{target} according to the first-order lag model. The new Q_{phys} is then used in the process model to compute the next state of the water quality. The environment's linear regression equations take the previous state (DO, ORP, etc.), the current inflow, and the achieved aeration Q_{phys} (after inertia) to calculate the new state $\mathrm{s}t + 1$. The historical dataset or provided profile is queried to update exogenous inputs for the next step (inflow rate and electricity price for $t + 1$). All resulting state variables are clipped to realistic bounds if needed (e.g. non-negativity, max sensor range).
5. Compute Reward: with the new state $\mathrm{s}t + 1$ and the action a_t, the environment evaluates the reward function described above. Internally, the raw physical values of DO, ORP, price, and Q_{phys} are used to calculate each component of the reward (DO penalty, energy cost, etc.), and these are combined into the scalar reward r_t. The PumpTracker is updated with the new pump state (on/off counters incremented or reset) before computing certain terms like switching penalties or off-duration bonuses. The final reward is returned along with the breakdown for logging.
6. Advance to Next State: the environment increments the time step. It packages the next observation o_{t+1} from $\mathrm{s}t + 1$ (including new lagged values, updated counters, etc.), and sends this $ot + 1$ to the agent, along with the reward r_t, and a flag indicating whether the episode has ended. In this episodic task, *done* will be true once $t = T$ (end of the 24h period), upon which the environment can be reset for a new episode.

This loop repeats until the episode is complete. During training, the agent uses the sequence of observations, actions, and rewards to adjust its policy (for example, by optimising a loss via backpropagation in an off-policy RL algorithm). During evaluation, the agent's actions can be recorded and the resulting performance (rewards and constraint violations) analysed.

4.7. Deep Reinforcement Learning Training

The training of the RL agents was structured in four experimental configurations, combining the two selected algorithms (SAC and TD3) with two different initialisation modes (from scratch and from a warm-start buffer). Each experiment was executed for 30.000 training steps, corresponding to approximately 300 virtual operational days, using the same environment, reward function, and process configuration described earlier in this chapter. All experiments were performed on a dedicated computational workstation equipped with an NVIDIA GPU, using the `Stable Baselines3` library within a controlled `conda` environment to ensure reproducibility.

4.7.1. Replay Buffer of Historical Actions

Before interactive training, a replay buffer was generated from the historical data (`cleaned_full.csv`). This buffer, saved as `sac_buffer_warmstart_v5.npz`, contains tuples of observations, actions, rewards, next observations, and episode termination flags collected from historical operational data, which correspond to the old station controller (in this case, a relay) and allows the RL algorithm to begin training with realistic transitions, effectively "warm-starting" the learning process. The use of a replay buffer aligns with the latest practices in off-policy learning, as both SAC and TD3 algorithms sample mini-batches of past experience uniformly from this buffer to update the policy and value networks.

4.7.2. Training Configurations

Four distinct training runs were conducted, combining algorithm type and initialisation strategy:

1. SAC (Fresh Start). In this configuration, the SAC agent was trained entirely from scratch, without any prior experience. The replay buffer was initialised empty, and the agent populated it during exploration. The environment used was `AerationEnvLinear2`, with the same linear regression model coefficients and physical constraints defined in Section 4.6. Training proceeded for 30.000 timesteps, corresponding to roughly 300 simulated 24-hour episodes. Live visualisation dashboards (implemented via `plotly.graph_objects`) monitored the evolution of dissolved oxygen, oxidation–reduction potential, airflow command, energy price, inflow, and instantaneous reward.

2. SAC (Warm Start). The second experiment reused the same model structure but initialised the replay buffer with the pre-computed transitions stored in `sac_buffer_warmstart_v5.npz`. This "warm-start" configuration allowed the SAC agent to begin learning from meaningful historical trajectories rather than from random initialisation. The architecture, hyperparameters, and episode configuration were identical to the fresh-start case. During training, the same visual dashboards were employed to track the evolution of dissolved oxygen and redox potential in response to the learned aeration policy.

3. TD3 (Fresh Start). The third experiment trained a Twin-Delayed Deep Deterministic Policy Gradient (TD3) agent from scratch under identical environmental and reward conditions. TD3, as introduced in Section 4.5.2, is a deterministic off-policy algorithm that relies on twin critic networks and delayed actor updates for stability. The same buffer size ($200\,000$), discount factor ($\gamma = 0.99$), and soft-update coefficient ($\tau = 0.005$) were used to ensure direct comparability with the SAC experiments. The total number of timesteps and evaluation frequency were also aligned with the SAC configuration. This experiment focuses on whether a noise-free actor can achieve the same balance between energy efficiency and water-quality maintenance as the stochastic SAC policy.

4. TD3 (Warm Start). Finally, the TD3 agent was trained with the same warm-start buffer used in the second experiment. This experiment combined the stabilising effect of pre-filled experience replay with the deterministic action selection of TD3. The warm-started TD3 configuration aimed to reduce the early training instability typically observed in value-based deterministic methods when starting from random exploration.

4.7.3. Hyperparameter Settings

Table 4.7 summarises the main hyperparameters adopted for all experiments. They were selected based on prior benchmarks in process control literature [5, 74] and on early pilot tests ensuring numerical stability of learning.

Parameter	Symbol / Setting	Value
Discount factor	γ	0.99
Target smoothing coefficient	τ	0.005
Learning rate	α_{LR}	3×10^{-4}
Batch size	B	256
Replay buffer size	N_{buffer}	200,000
Training timesteps	T_{train}	30,000
Policy update delay (TD3)	d_π	2
Target policy noise (TD3)	σ_{target}	0.2 (clipped to 0.5)
Entropy temperature (SAC)	α	(initial 0.2)
Evaluation frequency	f_{eval}	every 1,000 steps
Episode length	T_{episode}	96 (24 hours)
Hidden layers (actor/critic)	–	2 layers $\times$ 256 neurons (ReLU)

Table 4.7: Reinforcement learning hyperparameters used for SAC and TD3 experiments.

4.7.4. Training Procedure

At the start of each experiment, the environment was wrapped in a `DummyVecEnv` for parallel compatibility with the `Stable Baselines3` framework. Each training run proceeded in steps of $15\,\mathrm{min}$, with an episode corresponding to a 24-hour operational cycle of 96 steps. For every episode, the agent received initial conditions drawn randomly from the historical dataset and simulated the aeration control loop as described in Section 4.6. During training, the model was regularly evaluated on a separate environment (`eval_env`) to monitor policy improvement, and the best-performing models were automatically saved under `/best_model/` checkpoints.

The live dashboards implemented during training provided immediate feedback on the learning behaviour of the agent. In the initial stages, both SAC and TD3 tended to over-aerate due to the high reward sensitivity to dissolved oxygen penalties. As training progressed, the agents learned to oscillate airflow strategically, maintaining redox levels near the optimal range and reducing unnecessary blower operation during off-peak hours.

4.8. Benchmarking and Evaluation Setup

After completing the training phase, the performance of the four trained DRL controllers was systematically benchmarked against a conventional relay controller, representing the current industrial practice. The relay controller implements a simple rule-based on/off logic based on the oxidation–reduction potential (ORP) value. Aeration is switched on when ORP drops below a lower threshold, and off when it exceeds an upper threshold. The chosen bound values ensure that both controllers aim to keep ORP within the same hysteresis band. This type of control, also known as two-point relay control, is widely adopted in wastewater treatment plants due to its simplicity, robustness, and ease of implementation.

Additionally, the KPIs explained in the following subsection also allowed for comparing the four controllers with each other, as shown in Chapter 5.

4.8.1. Evaluation Methods and Key Performance Indicators

The evaluation process aimed to assess each controller under identical operating conditions to ensure a fair and reproducible comparison. For this evaluation, a common and consistent set of performance metrics was computed, following the definitions introduced in Section 4.3. These metrics quantify the ability of each controller to balance treatment efficiency, energy use, and operational safety, and include:

- Cumulative Reward per Episode (R_{tot}). Total reward accumulated across the entire episode, computed using the same multi-objective function defined during training:

  $$R_{\mathrm{tot}} = \sum_{t=1}^{T} r_t$$

  where r_t is the instantaneous reward at time step t, and T is the total number of steps per episode. The reward function jointly penalises excessive energy consumption and constraint violations, while rewarding stable and compliant operation. Therefore, R_{cum} acts as a synthetic index summarising the controller's overall performance across the other KPIs.

- Total Aerated Air Volume (Nm^3/day). Total amount of air injected into the basin over the episode, computed as the discrete-time integral of the physical airflow rate:

  $$V_{\mathrm{air}} = \sum_{t=1}^{T} Q_{\mathrm{air},t}\,\Delta t$$

  where $Q_{\mathrm{air},t}$ is the blower airflow (Nm^3/h) and Δt the sampling interval (0.25 h). This provides a direct indicator of total aeration effort and energy intensity.

- Operating Cost (€/day). Dynamic electricity cost of aeration, reflecting tariff variations throughout the episode:

  $$C_{\mathrm{op}} = \sum_{t=1}^{T} p_t\,Q_{\mathrm{air},t}\,\Delta t$$

  where p_t is the time-varying unit price of electricity (€/Nm^3). This KPI evaluates the controller's capacity to schedule aeration more intensively during cheaper tariff periods and reduce operation during cost peaks.

- Pump Violation Percentage (%). Proportion of time steps during which the blower operated below the minimum stable flow rate $Q_{\min}$:

  $$\mathrm{PV} = \frac{1}{T}\sum_{t=1}^{T} \mathbb{I}\big(0 < Q_{\mathrm{air},t} < Q_{\min}\big) \times 100$$

  where $\mathbb{I}(\cdot)$ is the indicator function. This metric quantifies mechanical or operational inefficiencies due to underflow conditions that may compromise diffuser performance.

- ON/OFF Duration Violations (count). Number of blower activation or deactivation periods that violated minimum or maximum duration constraints defined in the environment configuration:

$$N_{\text{short}}^{\text{ON}} = \sum_i \mathbb{I}(d_i^{\text{ON}} < d_{\text{min}}^{\text{ON}}), \qquad N_{\text{long}}^{\text{ON}} = \sum_i \mathbb{I}(d_i^{\text{ON}} > d_{\text{max}}^{\text{ON}}),$$
$$N_{\text{short}}^{\text{OFF}} = \sum_i \mathbb{I}(d_i^{\text{OFF}} < d_{\text{min}}^{\text{OFF}}), \qquad N_{\text{long}}^{\text{OFF}} = \sum_i \mathbb{I}(d_i^{\text{OFF}} > d_{\text{max}}^{\text{OFF}})$$

  where $d_i^{\text{ON/OFF}}$ denotes the duration (in steps) of each ON or OFF streak. To ensure reliable interpretation, the first and last streaks of each episode are excluded from this evaluation, as they may be truncated by episode boundaries.

- Mean Redox Potential (mV). Time-averaged oxidation–reduction potential across the episode:

$$\overline{ORP} = \frac{1}{T} \sum_{t=1}^{T} ORP_t$$

  Higher values within the optimal range indicate balanced nitrification–denitrification conditions, while persistently lower or higher levels suggest excessive anoxia or hyper-aeration, even if, overall, a higher value is preferable to a too low value for this KPI, since it is a comprehensive indicator of the effluent quality.

- Mean Dissolved Oxygen (mg/L). Time-averaged dissolved oxygen concentration:

$$\overline{DO} = \frac{1}{T} \sum_{t=1}^{T} DO_t$$

  This parameter reflects the controller's ability to maintain sufficient oxygen for nitrification while avoiding wasteful overaeration.

4.8.2. Design of Experiments and Scenarios

After completing the training phase, it is needed to evaluate the controllers' performance. For this purpose, the four different trained DRL agents built with the two architectures introduced in Section 4.5, were evaluated and benchmarked against two different rule-based relay controllers, which serve as the industrial baseline for aeration control. The relay controller implements a simple rule-based on/off logic based on the ORP. Aeration is switched ON when ORP drops below a lower threshold, OFF when it exceeds an upper threshold. This type of control, also known as two-point relay control, is widely adopted in wastewater treatment plants due to its simplicity, robustness, and ease of implementation.

The relay controllers used are different to maintain comparable goals with the reward function used (see Section 4.6.4).

Against `reward_function_1`, which prioritises water quality, the relay used is `relay_controller_1` with `high_ORP1` and `low_ORP1` as high and low redox thresholds, respectively.

On the other hand, against `reward_function_2`, which is more focused on energy savings, the relay used is `relay_controller_2` with `high_ORP2` and `low_ORP2`.

Obviously, `high_ORP1` is greater than `high_ORP2` and `low_ORP1` is greater than `low_ORP2`, as higher redox values are associated with better oxygenation and therefore better water quality.

The specific ORP values used cannot be shared due to the signature of NDAs before the beginning of this project. However, these values change substantially from one station to another; therefore, any numerical indication is not as useful.

All controllers were tested across five 24-hour (96 time steps at 15-minute intervals) scenarios based on real operational data that reproduced the dynamic load and tariff variations typical of real wastewater treatment plants. These scenarios were used both during evaluation and to visualise the controllers' behaviour, as illustrated later in Chapter 5.

5

Engineering Case Study and Results

5.1. Engineering Case Study

As explained previously (see Section 1.1), it has been decided to focus on the aeration process for experimenting with AI-based control strategies. Therefore, in this section, the biological processes happening in an aeration tank are explained mathematically. The proposed model is now simply used for understanding the process, which is a crucial step for designing an ad hoc control solution for it.

Figure 5.1: Water-treatment plant overview. Source: Rovatti [109].

5.1.1. Description and Biochemical Explanation of the Process

The case study focuses on a single aeration basin within a biological wastewater treatment line operating under the conventional activated sludge process (ASP). The basin receives variable influent flow and ammonium load, and its main function is to oxidize ammonium (S_{NH}) into nitrate (S_{NO_3}) through nitrification, while maintaining a suitable concentration of dissolved oxygen (C_{O_2}) to support the oxidation and the microbial metabolism of autotrophic bacteria.

The process data available from plant sensors include inflow rate, recirculation flow rates, pH, temperature, redox potential, dissolved oxygen, and concentrations of ammonium and nitrate. The airflow rate is the main manipulated variable, controlling oxygen transfer through bubble diffusion.

Only part of the relevant process states are directly measurable. While S_{NH}, S_{NO_3}, and C_{O_2} are monitored, the intermediate nitrite concentration (S_{NO_2}) and microbial activity levels are unobserved. Consequently, the output vector is:

$$y(t) = [S_{\mathrm{NH}}(t)\ S_{\mathrm{NO}_3}(t)\ C_{O_2}(t)\ E(t)]'$$

where $E(t)$ denotes the measured oxidation–reduction potential (ORP, often written as Redox potential), correlated to the ratio between oxidised and reduced nitrogen species according to a Nernst-type relationship [110, 111].

The present control strategy is rule-based: the aeration valve operates as a relay, switching the air inflow on or off according to the measured redox potential thresholds. When $E(t)$ drops below a set point, the blower activates until the potential rises above a predefined upper threshold. This on–off logic provides a basic but robust control ensuring oxidation, yet it often results in excessive energy use and limited adaptability to influent fluctuations.

Potential for intelligent control

More advanced control architectures could exploit the measured process variables within a predictive or learning-based framework. For example, Model Predictive Control (MPC) could minimise the aeration energy cost while constraining C_{O_2} and S_{NH} within acceptable limits. Similarly, Deep Reinforcement Learning (DRL) controllers could learn near-optimal dynamic policies balancing oxygen supply and pollutant removal under uncertainty. These strategies would replace the binary relay logic with smooth, adaptive actuation, improving both process stability and energy efficiency, a direction pursued in later stages of this research.

5.1.2. Dynamic Model of the System

The dynamic model adopted for the aeration basin follows the conventional structure of nitrification models [112], restricted to a single aerobic tank. The state vector is defined as:

$$x(t) = [S_{\mathrm{NH}}(t)\ S_{\mathrm{NO_2}}(t)\ S_{\mathrm{NO_3}}(t)\ C_{O_2}(t)]'$$

The system inputs are the inflow rate Q_t, influent ammonium concentration $S_{\mathrm{NH,in}}$, and air flow rate Q_{air}, while the measured outputs are $S_{\mathrm{NH}}, S_{\mathrm{NO_3}}, C_{O_2}$, and E.

The model dynamics are governed by a set of coupled non-linear ordinary differential equations:

$$\frac{dS_{\mathrm{NH}}}{dt} = \frac{Q_t}{V}\big(S_{\mathrm{NH,in}} - S_{\mathrm{NH}}\big) - r_{\mathrm{AOB}}, \tag{5.1a}$$

$$\frac{dS_{\mathrm{NO_2}}}{dt} = \frac{Q_t}{V}\big(S_{\mathrm{NO_2,in}} - S_{\mathrm{NO_2}}\big) + r_{\mathrm{AOB}} - r_{\mathrm{NOB}}, \tag{5.1b}$$

$$\frac{dS_{\mathrm{NO_3}}}{dt} = \frac{Q_t}{V}\big(S_{\mathrm{NO_3,in}} - S_{\mathrm{NO_3}}\big) + r_{\mathrm{NOB}}, \tag{5.1c}$$

$$\frac{dC_{O_2}}{dt} = \frac{Q_t}{V}\big(S_{O_2,\mathrm{in}} - C_{O_2}\big) + k_{La}(Q_{\mathrm{air}})\big(C^* - C_{O_2}\big) - 3.43\, r_{\mathrm{AOB}} - 1.14\, r_{\mathrm{NOB}} - OUR_{\mathrm{end}} \tag{5.1d}$$

The reaction rates are described by double Monod kinetics:

$$r_{\mathrm{AOB}} = AOR_{\max} \frac{S_{\mathrm{NH}}}{K_{\mathrm{NH}} + S_{\mathrm{NH}}} \frac{C_{O_2}}{K_{O_2,\mathrm{AOB}} + C_{O_2}}, \tag{5.2a}$$

$$r_{\mathrm{NOB}} = NOR_{\max} \frac{S_{\mathrm{NO_2}}}{K_{\mathrm{NO_2}} + S_{\mathrm{NO_2}}} \frac{C_{O_2}}{K_{O_2,\mathrm{NOB}} + C_{O_2}}. \tag{5.2b}$$

The oxygen transfer coefficient depends on the air inflow rate according to a semi-empirical correlation [113]:

$$k_{La} = k_{La,0} + \alpha\, Q_{\mathrm{air}}, \qquad C^* = C^*(T). \tag{5.3}$$

The ORP may be approximated using a logarithmic Nernst-type relationship between reduced and oxidized nitrogen species [110, 111]:

$$E(t) = a + b\,\mathrm{pH}(t) + c\, \log_{10}\left(\frac{S_{\mathrm{NH}}(t)}{S_{\mathrm{NO_2}}(t)}\right). \tag{5.4}$$

Parameter identification. The kinetic and transfer parameters were estimated using a Least Squares Method (LSM) applied to measured time series of S_{NH} and C_{O_2}. The identification minimizes the squared prediction error between simulated and measured outputs over short data windows:

$$\min_{\mathbf{p}} \sum_{t\in\mathcal{T}} \big\| y_{\mathrm{meas}}(t) - y_{\mathrm{sim}}(t,\mathbf{p}) \big\|_2^2, \quad \mathbf{p} = [AOR_{\max}\ K_{\mathrm{NH}}\ K_{O_2,\mathrm{AOB}}\ k_{La,0}\ \alpha\ \ldots]'. \tag{5.5}$$

The identified model reproduced the main short-term trends of dissolved oxygen and ammonium concentrations, but the resulting parameters exhibited time variability when the estimation window was shifted in time. This behaviour suggests slow nonstationary dynamics (e.g., biomass adaptation and temperature-dependent kinetics), implying that constant-parameter LSM identification provides only a local approximation of the process.

Despite these limitations, the LSM results confirm the physical plausibility of the model and its potential integration in a model-based control framework. In future developments, hybrid identification combining recursive LSM or adaptive estimation with data-driven corrections could improve tracking performance under varying process conditions.

Identification Results. The LSM was applied to experimental data, using measured time series of ammonium, dissolved oxygen, and redox potential. The optimization window covered several consecutive 15 min intervals during quasi-steady aeration periods. The resulting identified parameters are reported in Table 5.1.

Table 5.1: Parameter estimates obtained by LSM identification.

Parameter	Description	Estimated Value	Unit
$AOR_{\max}$	Max. ammonium oxidation rate (AOB)	5.3	mg N L^{-1} h^{-1}
$NOR_{\max}$	Max. nitrite oxidation rate (NOB)	0.74	mg N L^{-1} h^{-1}
K_{NH}	NH_4^+ half-saturation constant	0.62	mg N L^{-1}
$K_{\mathrm{NO_2}}$	NO_2^- half-saturation constant	1.15	mg N L^{-1}
$K_{O_2,\mathrm{AOB}}$	O_2 half-saturation constant (AOB)	0.12	mg O_2 L^{-1}
$K_{O_2,\mathrm{NOB}}$	O_2 half-saturation constant (NOB)	0.45	mg O_2 L^{-1}
$k_{La,0}$	Base oxygen transfer coefficient	0.010	min^{-1}
α	Aeration scaling factor	5.0×10^{-4}	m^{-1}
C^*	DO saturation concentration	8.5	mg O_2 L^{-1}
OUR_{end}	Endogenous oxygen uptake rate	0.10	mg O_2 L^{-1} h^{-1}

The identified parameters are consistent with those reported in the literature for aeration tanks operating at low dissolved oxygen concentrations [113], confirming the validity of the simplified dynamic model. Among all parameters, $k_{La,0}$ and α exhibited the highest temporal sensitivity, suggesting that the oxygen transfer efficiency changes dynamically with operating conditions such as air flow, sludge rheology, and diffuser fouling. Similarly, $AOR_{\max}$ and $NOR_{\max}$ showed slow variations over time windows of several hours, reflecting biological adaptation of the nitrifying biomass to fluctuating load and temperature. This behaviour confirms that the system is partially time-varying and that a static parameter set provides only a local approximation of the underlying process dynamics.

The simulated trajectories of dissolved oxygen and ammonium reproduced the measured signals with a mean absolute error (MAE) of about 0.3 mg L^{-1} for DO and 0.25 mg L^{-1} for NH_4^+. The model successfully captured the typical oscillatory dynamics caused by the relay aeration control, with a characteristic rise of C_{O_2} during blower activation and a progressive decline when aeration was off. The predicted redox potential $E(t)$ obtained from (5.4) qualitatively matched the measured signal, showing the expected cyclic variations synchronised with aeration periods and the shift of redox balance between nitrifying and denitrifying phases.

The identification confirms that the proposed dynamic model, with parameters estimated via LSM, captures the essential behaviour of the aeration tank while revealing a significant time-varying component in both oxygen transfer and biological kinetics. This intrinsic variability suggests that fixed-parameter or model-based controllers may not always adapt effectively to changing operating conditions. Consequently, the observed time dependence provides a strong motivation for introducing adaptive control schemes such as reinforcement learning (RL), which can learn and update optimal aeration policies directly from process data in real time. These results therefore establish a quantitative baseline for the development of hybrid model-based predictive and data-driven control strategies aimed at optimising aeration efficiency under non-stationary operating conditions.

5.2. Engineering Results: Models and Controlling Performance

This section presents and discusses the results obtained in the engineering part of this research from the experimental evaluation of the Deep DRL controllers as described in Section 4.8.

The discussion follows the same set of Key Performance Indicators (KPIs) defined in Section 4.8.1, ensuring consistency in evaluation. These results aim to answer the first two research sub-questions introduced in Chapter 3, which are:

- *Given the state of the art of AI, is it technically feasible to control wastewater treatment processes with AI solutions in real facilities?*
- *How effective are the proposed AI solutions in improving energy efficiency and pollutants control in WWTPs?*

5.2.1. Quantitative Outcomes Across Five Scenarios

As explained in Section 4.8.2, the Deep RL agents were evaluated against two different relay controllers on five 24-hour evaluation scenarios, each representing a distinct daily inflow and electricity price profile generated from the reference plant data, ensuring exposure to realistic operational variability.

A visualisation of the five test scenarios is shown in Figure 5.2 and Figure 5.3.

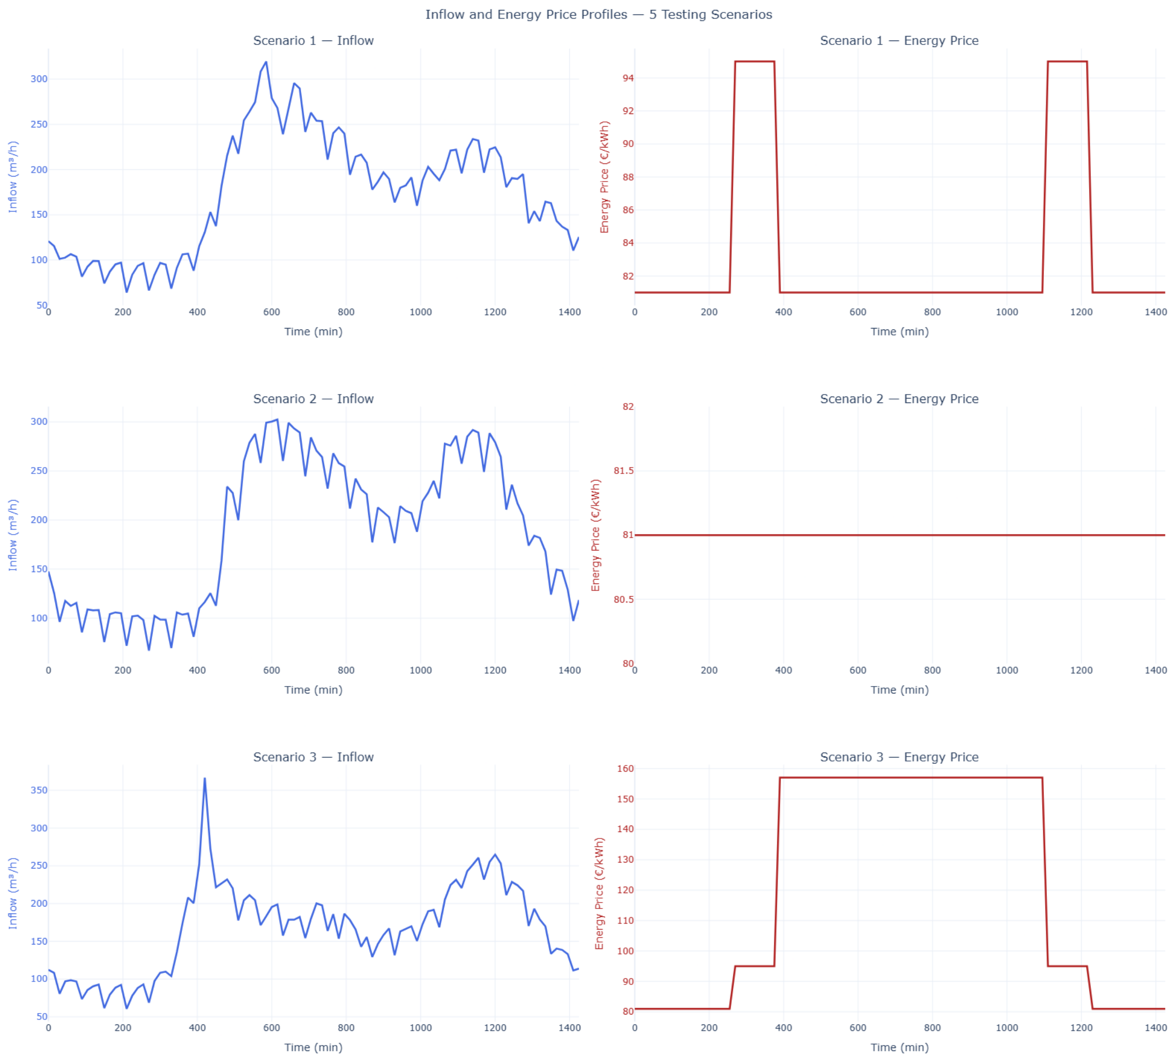

Figure 5.2: Inflow and Energy Price Profiles — Testing Scenarios (Part 1)

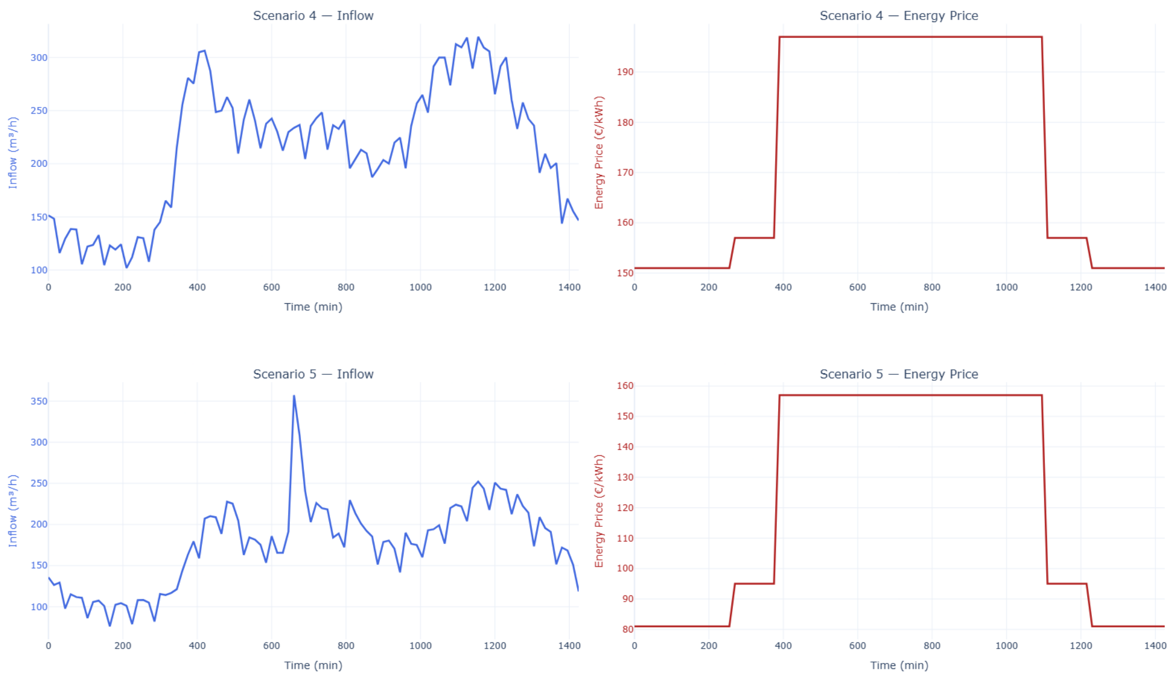

Figure 5.3: Inflow and Energy Price Profiles — Testing Scenarios (Part 2)

Results with Reward Function 1 for the Deep RL Agents

Table 5.2 shows the outcomes of the tests conducted on the five scenarios using `reward_function_1` and `relay_controller_1`.

Table 5.2: Results of the five evaluation scenarios for all controllers using `reward_function_1`.

Agent	Sc.	$R1_{\mathrm{tot}}$	V_{a} [Nm^3/day]	C [€/day]	$\overline{ORP}$ [mV]	$\overline{DO}$ [mg/L]	PV [%]	DV[-]
SAC-fresh	1	70.99	16308	2.44e6	225.7	0.93	0.0	0
SAC-warm	1	65.84	16294	2.43e6	225.0	0.92	0.0	0
TD3-fresh	1	73.88	16116	2.41e6	222.8	0.91	0.0	0
TD3-warm	1	70.69	16347	2.44e6	227.4	0.92	0.0	0
Relay1	1	48.84	15789	2.38e6	192.8	0.86	0.0	0
SAC-fresh	2	70.59	15752	2.37e6	225.2	0.93	0.0	0
SAC-warm	2	65.07	15743	2.39e6	228.7	0.94	0.0	0
TD3-fresh	2	73.94	15722	2.39e6	225.8	0.93	0.0	0
TD3-warm	2	70.33	15594	2.37e6	226.4	0.93	0.0	0
Relay1	2	49.92	15789	2.40e6	194.2	0.86	0.0	0
SAC-fresh	3	67.22	15678	2.71e6	225.9	0.94	0.0	0
SAC-warm	3	62.87	15899	2.75e6	223.7	0.93	0.0	0
TD3-fresh	3	68.69	15794	2.75e6	233.1	0.96	0.0	0
TD3-warm	3	68.73	15639	2.73e6	219.1	0.91	0.0	0
Relay1	3	44.77	15789	2.75e6	193.2	0.87	0.0	0
SAC-fresh	4	65.04	16493	1.37e6	207.2	0.79	0.0	0
SAC-warm	4	66.10	16277	1.35e6	238.0	0.85	0.0	0
TD3-fresh	4	67.40	16429	1.37e6	198.0	0.79	0.0	0
TD3-warm	4	66.10	16589	1.38e6	228.8	0.82	0.0	0
Relay1	4	44.06	16745	1.38e6	181.6	0.75	0.0	0

Continues on next page

Agent	Sc.	$R1_{\text{tot}}$	V_{a} [Nm^3/day]	C [€/day]	$\overline{ORP}$ [mV]	$\overline{DO}$ [mg/L]	PV [%]	DV[-]
SAC-fresh	5	71.77	15908	2.78e6	225.6	0.91	0.0	0
SAC-warm	5	65.87	15897	2.78e6	224.4	0.90	0.0	0
TD3-fresh	5	69.30	16152	2.80e6	229.6	0.91	0.0	0
TD3-warm	5	67.83	16483	2.85e6	213.6	0.85	0.0	0
Relay1	5	49.85	15789	2.75e6	191.6	0.83	0.0	0

Note: $R1_{\text{tot}}$ = cumulative reward; V_{a} = total aerated air volume (Nm^3/day); C = operating cost (€/day), given by $V_{\text{air}} \times EnergyPrice$; $\overline{ORP}$ = mean redox potential (mV); $\overline{DO}$ = dissolved oxygen (mg/L); PV = percentage of sub-minimum airflow steps; DV = number of steps that violated minimum or maximum pump duration constraints.

Results with Reward Function 2 for the Deep RL Agents

Table 5.3 shows the outcomes of the tests conducted on the five scenarios using `reward_function_2` and `relay_controller_2`.

Table 5.3: Results of the five evaluation scenarios for all controllers using `reward_function_2`.

Agent	Sc.	$R1_{\text{tot}}$	V_{a} [Nm^3/day]	C [€/day]	$\overline{ORP}$ [mV]	$\overline{DO}$ [mg/L]	PV [%]	DV[-]
SAC-fresh	1	-21.23	14593	1.21e6	117.9	2.33	0.0	0
SAC-warm	1	-5.31	15152	1.30e6	112.8	0.53	0.0	0
TD3-fresh	1	-4.24	14952	1.27e6	105.4	0.76	0.0	0
TD3-warm	1	-5.47	14902	1.27e6	107.8	0.80	0.0	0
Relay	1	-34.58	14976	1.28e6	96.9	0.60	0.0	0
SAC-fresh	2	-10.24	15293	1.24e6	120.1	0.52	0.0	0
SAC-warm	2	-7.85	15459	1.30e6	116.3	0.50	0.0	0
TD3-fresh	2	-12.78	15323	1.28e6	112.1	0.54	0.0	0
TD3-warm	2	-7.69	15294	1.24e6	116.3	0.90	0.0	0
Relay	2	-47.44	15671	1.27e6	90.1	0.60	0.0	0
SAC-fresh	3	-18.80	14566	1.77e6	103.8	0.51	0.0	0
SAC-warm	3	-21.60	14432	1.84e6	114.3	1.84	0.0	0
TD3-fresh	3	-19.58	14530	1.80e6	115.5	0.58	0.0	0
TD3-warm	3	-18.71	14645	1.78e6	100.9	0.55	0.0	0
Relay	3	-43.30	14716	1.82e6	97.4	0.60	0.0	0
SAC-fresh	4	-57.79	15798	2.76e6	97.4	0.84	0.0	0
SAC-warm	4	-70.29	15587	2.74e6	138.1	0.54	0.0	0
TD3-fresh	4	-59.43	15780	2.77e6	121.0	0.52	0.0	0
TD3-warm	4	-58.52	15802	2.78e6	110.3	0.53	0.0	0
Relay	4	-77.45	16053	2.83e6	92.1	0.63	0.0	0
SAC-fresh	5	-18.68	14857	1.82e6	105.1	1.03	0.0	0
SAC-warm	5	-26.61	14942	1.87e6	124.7	0.52	0.0	0
TD3-fresh	5	-17.76	14581	1.80e6	109.7	0.51	0.0	0
TD3-warm	5	-16.97	14712	1.80e6	106.3	0.51	0.0	0
Relay	5	-46.79	15098	1.86e6	96.0	0.58	0.0	0

Note: $R2_{\text{tot}}$ = cumulative reward; V_{a} = total aerated air volume (Nm^3/day); C = operating cost (€/day), given by $V_{\text{air}} \times EnergyPrice$; $\overline{ORP}$ = mean redox potential (mV); $\overline{DO}$ = dissolved oxygen (mg/L); PV = percentage of sub-minimum airflow steps; DV = number of steps that violated minimum or maximum pump duration constraints.

Consideration on Operational Safety

As can be seen, in both cases, all DRL controllers maintained safe operation throughout the experiments, with zero pump or duration violations (PV = 0%, $N_{\text{viol}} = 0$).

5.2.2. Average Performance Across Scenarios

Table 5.4 and Table 5.5 report the average KPI values computed across the five scenarios using both `reward_function_1` and `reward_function_2`.

Table 5.4: Average KPI values across all five scenarios (reward_function_1).

Agent	$\overline{R1_{\mathrm{tot}}}$	$\overline{V_{\mathrm{air}}}$ [Nm^3/day]	$\overline{C}$ [€/day]	$\overline{ORP}$ [mV]	$\overline{DO}$ [mg/L]	$\Delta R1_{\mathrm{tot}}$ vs Relay [%]
SAC-fresh	69.52	16028	2.33e6	222.9	0.90	+42.5
SAC-warm	65.15	16062	2.34e6	227.5	0.91	+35.4
TD3-fresh	70.65	16043	2.34e6	221.7	0.90	+45.7
TD3-warm	68.35	16330	2.36e6	223.9	0.88	+40.3
Relay1	47.89	15980	2.33e6	190.7	0.84	—

Table 5.5: Average KPI values across all five scenarios (reward_function_2).

Agent	$\overline{R2_{\mathrm{tot}}}$	$\overline{V_{\mathrm{a}}}$ [Nm^3/day]	$\overline{C}$ [€/day]	$\overline{ORP}$ [mV]	$\overline{DO}$ [mg/L]	$\Delta R2_{\mathrm{tot}}$ vs Relay [%]
SAC-fresh	-25.75	15021	1.76e6	108.9	1.05	+47.6
SAC-warm	-26.73	15194	1.81e6	121.2	0.79	+45.6
TD3-fresh	-22.76	15033	1.78e6	112.7	0.58	+53.6
TD3-warm	-21.07	15071	1.77e6	110.3	0.66	+57.1
Relay	-49.11	15283	1.81e6	94.5	0.60	—

5.2.3. Control Behaviour

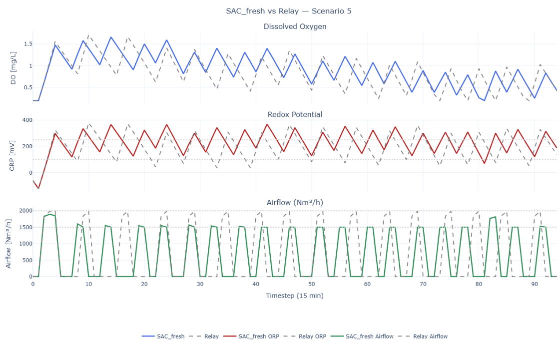

(a) SAC_fresh vs Relay (RF1)

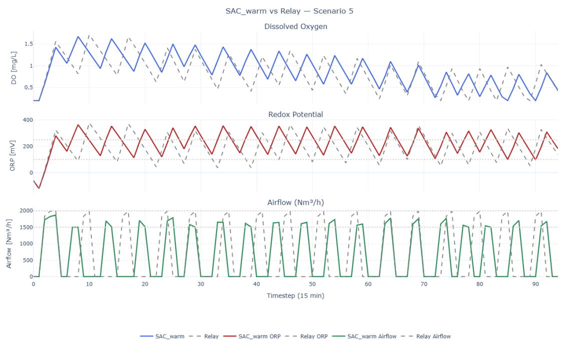

(b) SAC_warm vs Relay (RF1)

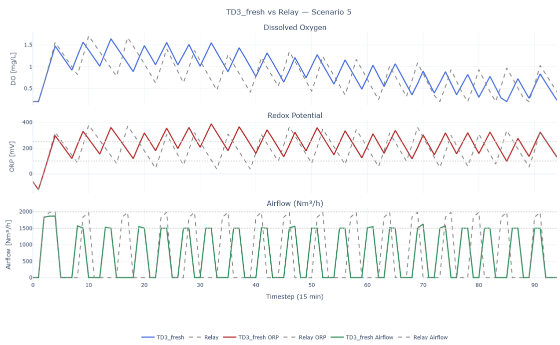

(c) TD3_fresh vs Relay (RF1)

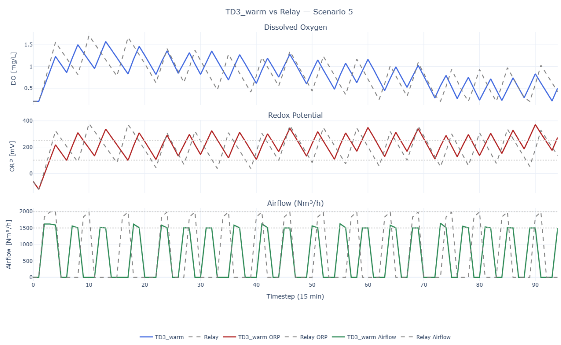

(d) TD3_warm vs Relay (RF1)

Figure 5.4: Comparison of SAC and TD3 controllers trained with `reward_function_1`, both freshly trained and warm-started, against the baseline relay control under Scenario 5. Each subplot shows the temporal evolution of dissolved oxygen (DO), oxidation–reduction potential (ORP), and airflow rate across 24 hours. Solid lines represent DRL controllers, while dashed lines indicate the relay baseline.

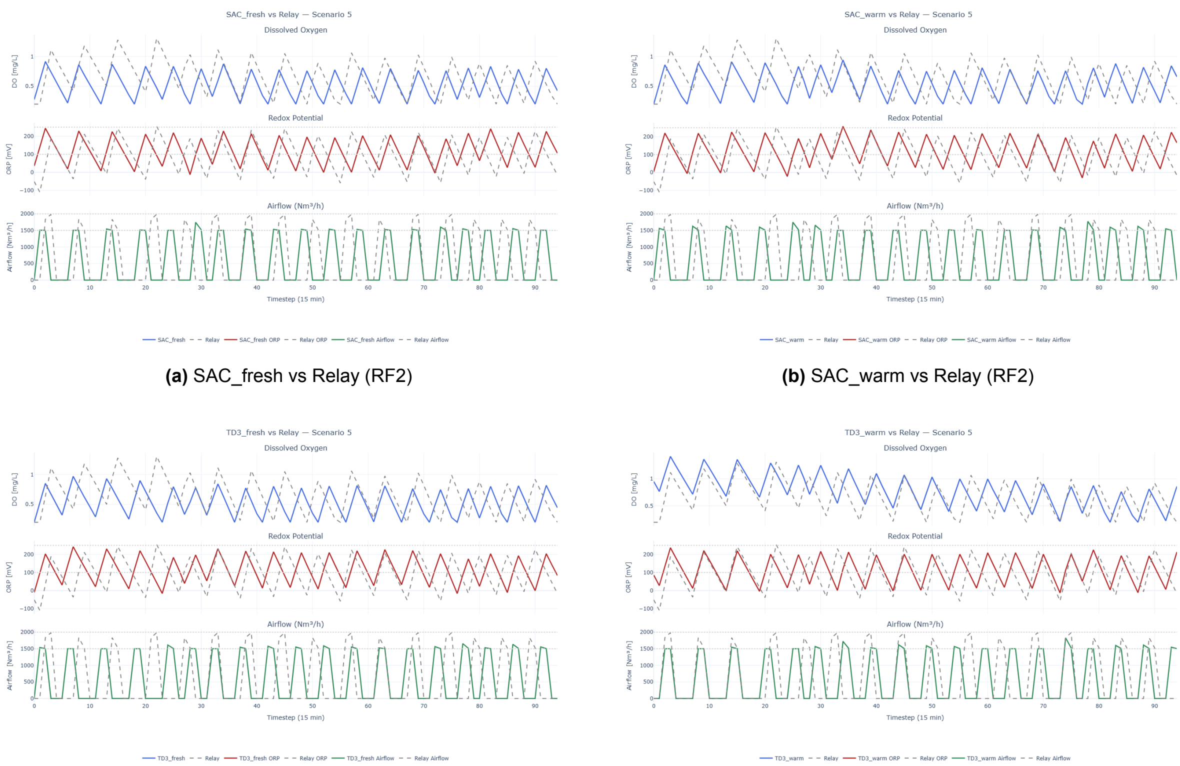

(a) SAC_fresh vs Relay (RF2)

(b) SAC_warm vs Relay (RF2)

(c) TD3_fresh vs Relay (RF2)

(d) TD3_warm vs Relay (RF2)

Figure 5.5: Comparison of SAC and TD3 controllers trained with `reward_function_2`, both freshly trained and warm-started, against the baseline relay control under Scenario 5. Each subplot shows the temporal evolution of dissolved oxygen (DO), oxidation–reduction potential (ORP), and airflow rate across 24 hours. Solid lines represent DRL controllers, while dashed lines indicate the relay baseline.

The behaviour of the DRL controllers under Scenario 5 is illustrated in Figure 5.4 and Figure 5.5, as an example. The first group of plots (Figure 5.4) shows the SAC and TD3 agents trained with `reward_function_1`, both in their fresh and warm-start versions, compared against the relay baseline. The second group (Figure 5.5) presents the corresponding results when the agents are trained with `reward_function_2`. In both sets of figures, each subplot displays the 24-hour evolution of dissolved oxygen (DO), oxidation–reduction potential (ORP), and airflow rate. Solid lines represent the DRL controllers, while dashed lines show the baseline relay control.

5.2.4. Interpretation of Results

Clear trends emerge regarding how the DRL agents exploit the different optimisation objectives encoded in the two reward functions.

Reward Function 1 (Water-Quality-Oriented). All DRL agents achieve substantially higher cumulative rewards than the relay controller, with improvements ranging from +35% to +46%. This reflects a more effective regulation of ORP and DO levels, as the agents learn to time aeration phases more strategically while respecting pump constraints. Average aeration volume and operating cost remain comparable to the relay, indicating that the DRL controllers do not reduce the total amount of aeration but redistribute it more efficiently. The mean ORP increases by approximately 30 mV relative to the baseline, suggesting more stable nitrification dynamics, while mean DO remains slightly higher (0.88–0.91 mg/L versus 0.84 mg/L). Differences between fresh-start and warm-start variants remain small, confirming that the reward structure itself is sufficient for effective learning.

Reward Function 2 (Energy-Focused). When the optimisation objective prioritises energy savings, the DRL agents markedly reduce the total aeration effort. Across scenarios, airflow decreases by approximately 4–8% compared to the relay, producing a 2–3% reduction in operating cost. Despite the reduction in aeration, the agents maintain biologically acceptable ORP and DO values, thanks to the strong penalisation of overly reducing conditions in the reward. The mean ORP remains above the relay benchmark, and the DO stays within a safe operating range. The largest reward gains (+45% to +57%) are obtained under this formulation, reflecting the dominance of the energy-penalty term. Among all agents, TD3-warm achieves the best aggregate trade-off between low aeration, low cost, and biological stability.

Cross-Reward Comparison. The two reward functions lead to distinct operating regimes. Under `reward_function_1`, DRL improves water-quality indicators without substantially altering total aeration demand. Under `reward_function_2`, DRL explicitly reduces aeration and cost while preserving acceptable effluent-quality proxies. In all cases, constraint violations remain zero, demonstrating that the learned policies consistently comply with operational safety requirements.Overall, the averaged KPIs indicate that DRL controllers are capable of learning different aeration strategies depending on the encoded objectives: either stabilising water-quality dynamics at equal energy use, or reducing energy consumption while preserving treatment performance.

Visualisation of Result. Figure 5.6 and Figure 5.7 visualise the normalised KPIs of each DRL controller against the relay baseline.

For metrics where higher raw values indicate *better* performance (e.g., Reward, DO, ORP), the normalised value is used directly. For metrics where higher raw values indicate *worse* performance (e.g., Cost and Airflow V_{air}), is plotted $KPI^* = 1 - KPI$.

This ensures a consistent visual interpretation across the radar charts: a point further from the centre always represents a more desirable outcome.

Figure 5.6: Radar comparison between each DRL controller and the relay baseline under `reward_function_1`. Each subplot shows the normalised KPIs for one DRL agent versus the relay controller (black). Note that $Cost^*$ and V^*_{air} are shown as $Cost^* = 1 - Cost$ and $V^*_{\mathrm{air}} = 1 - V_{\mathrm{air}}$ so that a wider area always correspond to a better performance.

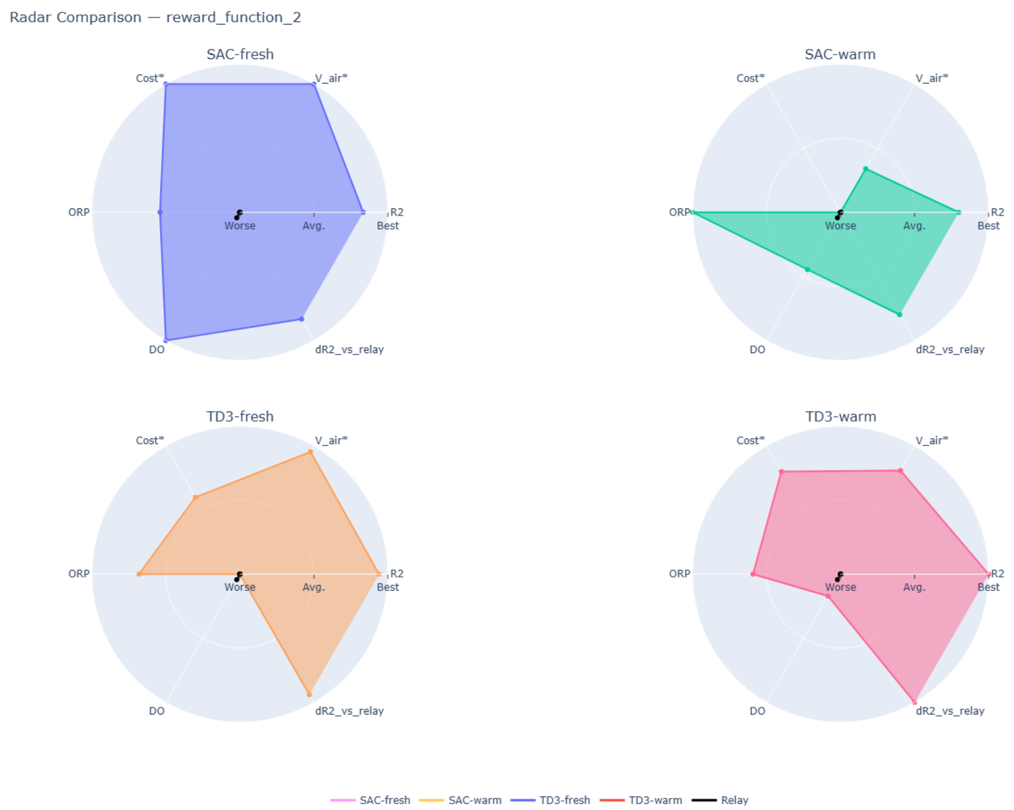

Figure 5.7: Radar comparison between each DRL controller and the relay baseline under `reward_function_2`. Each subplot shows the normalised KPIs for one DRL agent versus the relay controller (black). Note that $Cost^*$ and V^*_{air} are shown as $Cost^* = 1 - Cost$ and $V^*_{\text{air}} = 1 - V_{\text{air}}$ so that a wider area always correspond to a better performance.

By mapping all ”better” performance toward the outer region of the chart, the area enclosed by each agent’s polygon becomes an intuitive indicator of overall operational quality.

5.2.5. Feasibility and Benefits Compared to the Baseline Controller

The results obtained across all evaluation scenarios demonstrate that the application of Deep RL to aeration control is technically feasible and operationally robust. All trained agents, under both reward functions and for both SAC and TD3, did not violate the operational constraints, confirming that the learned policies can operate safely within the strict mechanical and operational requirements of real wastewater treatment systems.

When compared to the relay baseline, the DRL controllers consistently achieve superior performance, but in a way that depends on the optimisation objective encoded in the reward. With the water-quality-oriented reward, DRL policies increased DO and ORP with similar costs and energy consumption, improving nitrification behaviour. Under the energy-focused reward, the same agents learn to reduce airflow and operational costs while preserving acceptable surrogate indicators of treatment performance. These behaviours confirm that DRL can adapt its strategy according to different plant priorities.

Overall, the experiments show that DRL-based aeration control is feasible, safe, and capable of delivering measurable benefits over conventional rule-based strategies. These findings provide a solid engineering justification for exploring how such AI systems could be integrated into the governance and operational decision-making of critical water infrastructures. The next chapters, therefore, expand the analysis beyond technical performance, examining the institutional, regulatory, and organisational conditions that influence the adoption of AI-based control solutions in real wastewater utilities.

6

Research Methodology: Policy Analysis Approaches

6.1. Motivation of Methodological Choices

The methodological choices adopted in this part of the research stem directly from the third, fourth, and fifth sub-questions defined in Chapter 3. These questions address the organisational, regulatory, and societal conditions required for the safe integration of AI into wastewater infrastructures, the perceptions of stakeholders regarding risks and benefits, and the systemic factors that may enable large-scale adoption. Answering them requires a framework capable of linking technological change to institutional dynamics and stakeholder behaviour. For this reason, the analysis combines semi-structured interviews with decision-making and transition-theory frameworks: the *Decision-Making Levels*, the *Multi-Level Perspective* (MLP), and the *Transition Model Canvas* (TMC) [114, 30].

Semi-structured interviews were chosen to balance comparability across respondents with flexibility to explore emerging insights. In such settings, the ability to adapt follow-up questions is crucial to uncover context-specific insights without losing analytical consistency [115]. This makes semi-structured interviews especially valuable for governance and innovation studies involving complex socio-technical systems where stakeholder expertise and perspectives vary widely.

On the other hand, the analytical frameworks provide a structured way to connect technological innovation with institutional, organisational, and societal dynamics within these systems [114, 30]. In practice, in this research, they served multiple purposes: (i) they guided the selection of interviewees by ensuring representation across different socio-technical layers, (ii) informed the formulation of interview questions by linking them to relevant transition dimensions, and (iii) later provided the conceptual basis for integrating qualitative insights with technical findings in Chapter 8. This integration allowed the qualitative phase to remain theoretically grounded while directly contributing to the broader understanding of how AI-based control systems can be embedded in sustainable water governance.

A final motivation for this methodology arises directly from the knowledge gaps identified in Chapter 2. As the literature synthesis shows, research on AI-based control in the water sector shows highly promising technical results, still largely confined to simulations [5, 75], and a rapidly evolving policy and regulatory landscape, where utilities and regulators have limited experience with high-risk AI systems [17, 16]. This highlights the need for an empirical, actor-centred exploration of how governance, regulation, organisational capabilities, and technical constraints interact in practice. The methodological approach adopted was therefore explicitly designed to bridge these gaps by connecting the technological opportunities identified in Chapter 2 with the institutional and organisational realities that shape AI adoption. Overall, this design also reflects the interdisciplinary nature of the thesis: while the engineering approach in Chapter 4 and Chapter 5 tests what AI can achieve, the policy analysis investigates how such innovation can be responsibly adopted within real socio-technical systems.

6.2. Analytical Frameworks

To explore how AI can be integrated into critical water infrastructures, this study adopts the *Transition Model Canvas* (TMC) from [30] as its overarching analytical lens, which builds upon the *Multi-Level Perspective* (MLP) from [114], translating its concepts into an operational framework applicable to policy analysis. This research also uses the *Decision-Making Levels* framework, which is relevant to understanding organisations' structures and the interviewees' sampling. These concepts are explained in the following subsections, but their final implementation is shown in the integrative analysis in Chapter 8.

6.2.1. Decision-Making Levels

A first concept is that of *Decision-Making Levels*, often distinguished as *strategic, tactical, and operational*. This three-level framework was originally developed in management control theory by Anthony [116] and later applied to complex socio-technical systems and risk management by Rasmussen [117]. At the *strategic level*, long-term policy goals and regulatory frameworks are defined, typically by national governments and supranational bodies such as the European Union. The *tactical level* translates these frameworks into concrete regulations, plans, and standards through ministries, regulatory agencies, and authorities. Finally, the *operational level* involves utilities, plant operators, and municipalities, who implement regulations on the ground by managing WWTPs and ensuring compliance. Innovation, such as AI-based control, interacts differently with each level: strategic actors shape long-term legitimacy, tactical actors define compliance conditions, and operational actors handle daily feasibility.

6.2.2. The Multi-Level Perspective (MLP)

The *Multi-Level Perspective* (MLP) developed by [114] provides a conceptual framework to analyse how large socio-technical systems evolve and transform over time. It integrates insights from evolutionary economics and science–technology studies to explain how technological, institutional, and behavioural changes co-evolve. The MLP conceptualises transitions as multi-layered processes involving interactions between three analytical levels that differ in scope, stability, and rate of change but are continuously interlinked:

- Landscape: the broad, exogenous environment encompassing slow-changing trends and external pressures such as societal values, policy agendas, demographic shifts, or global crises;
- Regime: the dominant socio-technical configuration of infrastructures, technologies, institutions, and user practices that stabilises current systems and guides incremental innovation;
- Niche: the protected spaces where radical innovations, new technologies, or organisational forms can emerge and mature before challenging the existing regime.

These three levels and their interactions are modelled in the diagram in Figure 6.1 as an adaptation from [114, 30].

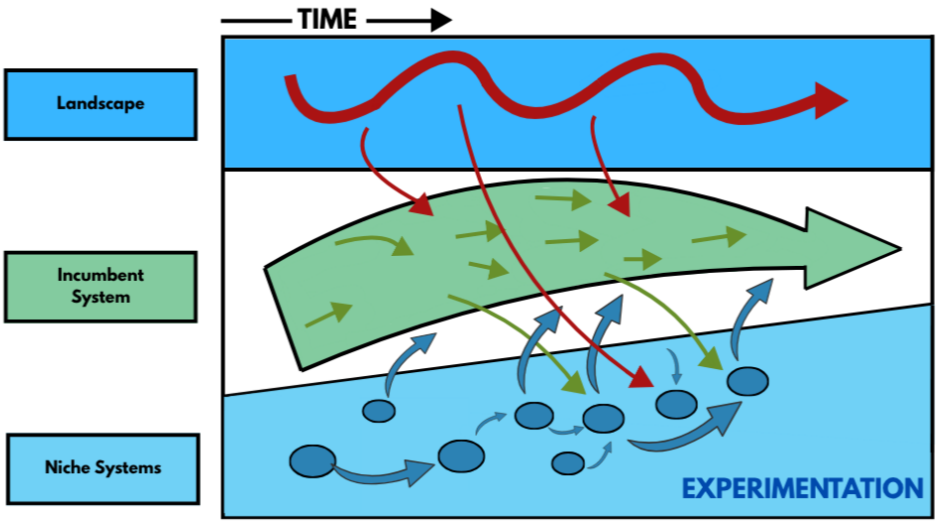

Figure 6.1: Simplified representation of the Multi-Level Perspective, adapted from [114, 30]. Transitions emerge through the interaction of landscape pressures, regime stability, and niche experimentation.

Transitions occur when developments at these three levels align. Pressures or opportunities at the landscape level destabilise the existing regime, while learning and network-building in niches generate alternatives that can exploit these openings. As [114] describes, transitions represent not merely technological substitutions but broad processes of *reconfiguration*, where technical, institutional, and cultural elements realign around new socio-technical configurations.

In this thesis, the MLP provides the foundational logic for analysing the governance of digital transitions in the water sector. It frames AI-based optimisation systems as niche innovations within a relatively stable wastewater treatment regime characterised by standardised operations, regulatory compliance, and risk aversion. Meanwhile, evolving landscape drivers (e.g. the EU AI Act, climate-neutrality targets, and the digital transition agenda) create external pressures that may open “windows of opportunity” for these innovations to scale and reshape incumbent structures. This layered view allows the analysis to trace how technological advances interact with institutional change across levels of governance.

6.2.3. The Transition Model Canvas (TMC)

Building on the MLP, the *Transition Model Canvas* (TMC) proposed by [30] operationalises transition theory into a practical analytical tool. According to [30], its value lies in its accessibility and adaptability, allowing transition analysis to move from theoretical reflection to practical application, and that is why it has been selected for this research.

As can be noticed in Figure 6.2, the TMC preserves the multi-level logic of niches, regimes, and landscapes but presents it through canvas blocks (see Figure 6.2) designed to support systematic analysis and stakeholder dialogue. These enable researchers and practitioners to map transition dynamics by distinguishing between four interrelated components:

1. Transition goal: the overarching transformation objective that defines the desired end state of the system;
2. Incumbent system: the dominant configuration of actors, infrastructures, and institutional rules stabilising the current regime, including its strengths, vulnerabilities, and defensive strategies;
3. Niche system: emerging innovations, networks, and coalitions that challenge or complement the incumbent regime;
4. Landscape: the wider political, economic, and societal context that shapes both the incumbent and niche systems through gradual trends or sudden shocks.

It is important to underline that the TMC goes beyond the simple identification of structural elements by adding also the *strategies* actors use to influence transitions (e.g. defending or destabilising the incumbent system, or strengthening the niches), therefore bridging analytical and practical perspectives.

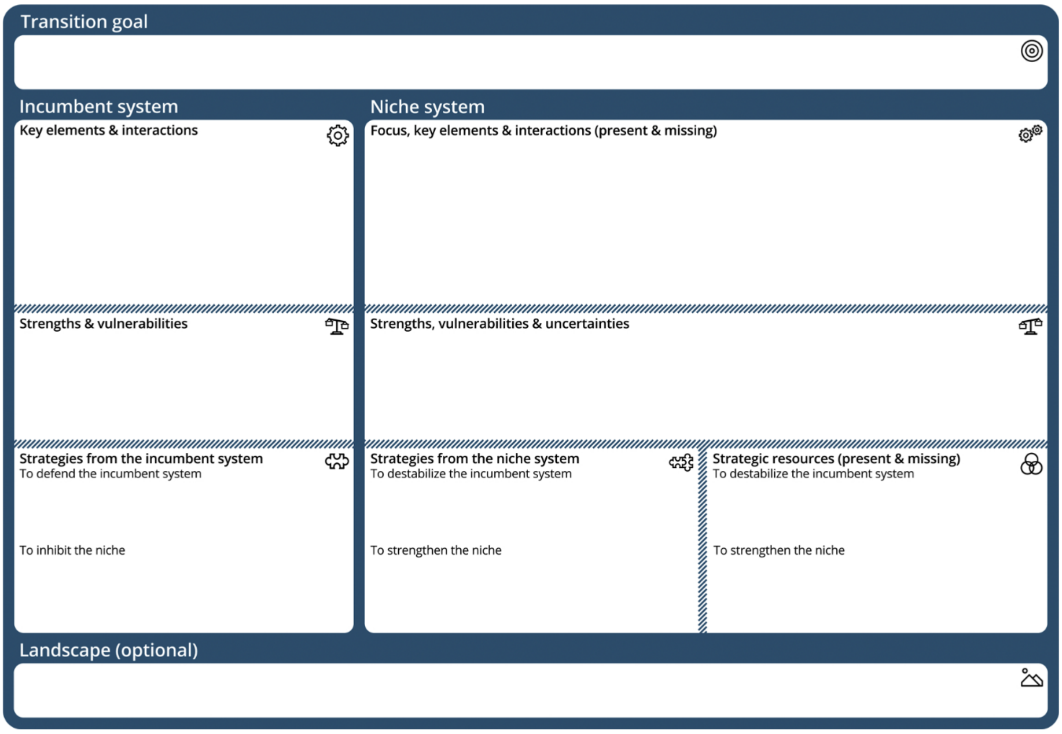

Figure 6.2: Structure of the Transition Model Canvas, from [30].

In this research, the TMC is used to design and code the semi-structured interviews, linking stakeholders' perceptions to the systemic conditions identified through the MLP. The complete and populated canvas will later serve as the central instrument for integrating policy and engineering insights in Chapter 8.

6.3. Transition Context

After having introduced the TMC, it is useful to position the emerging use of artificial intelligence in critical water infrastructures within the framework, to clarify and motivate how its components have been identified.

It is important to understand that these statements must be introduced at this point in the discussion, even if the final implementation will be part of Chapter 8. This is because filling the TMC has been an iterative process. It started with an initial idea, grounded in literature, of what the key TMC components could be, and this will be confirmed (or not) with the interviews and the technical results, providing also suggested strategies for implementation, corrections, and improvements.

Therefore, Table 6.1, below, summarises the preliminary configuration of the sociotechnical system through the TMC.

Table 6.1: Preliminary Transition Model Canvas framing for AI integration in wastewater treatment.

Transition goal	Moving from outdated reactive control systems towards a safe and effective integration of AI-based optimisation and control systems in wastewater infrastructures to improve energy efficiency, resilience, and environmental performance while maintaining regulatory compliance.
Incumbent system	SCADA/PLC-based automation embedded in a compliance-oriented regime of utilities, regulators, and technology providers. Guided by prescriptive standards, risk-averse organisational cultures, and mature but inflexible infrastructures [3, 13].
Niche system	Emerging AI-driven control and monitoring technologies (soft sensors, deep-RL controllers) developed in protected environments [5, 72]. Supported by evolving norms on transparency and safety (EU AI Act [16]).
Landscape	External pressures from the European Green Deal and AI regulation, coupled with climate-neutrality and digital-transition agendas [14, 16].

This configuration indicates that the European water sector is currently in an *emergent niche accumulation* phase [114], in which AI-based control solutions demonstrate technical feasibility but remain peripheral. The alignment between niche developments and landscape pressures could accelerate transition dynamics if regulation evolves to recognise AI as an enabler of sustainability rather than a source of operational risk. In TMC terms, this transition is marked by niche actors consolidating coalitions and advocating for regulatory adaptation, while incumbent actors cautiously explore integration pathways through controlled pilots and sandboxes [93, 96].

6.4. Interviewee Distribution and Sampling Strategy

Interviewees were selected strategically according to their position in the water and digital innovation ecosystem from industries and institutions classified by power and interest in Chapter 7 (Section 7.1). The goal was to capture perspectives across all decision-making levels and institutional sectors, ensuring representation from both the technological and policy domains. Table 6.2, summarises the interviewees' profiles, conveying the intersections between the analysed categories.

The interviewees were first approached via email, LinkedIn, or mutual contacts (such as suggestions from the graduation committee members). Forward snowballing and questions like "who should I interview next in your opinion?" during the interviews were also used to identify other participants.

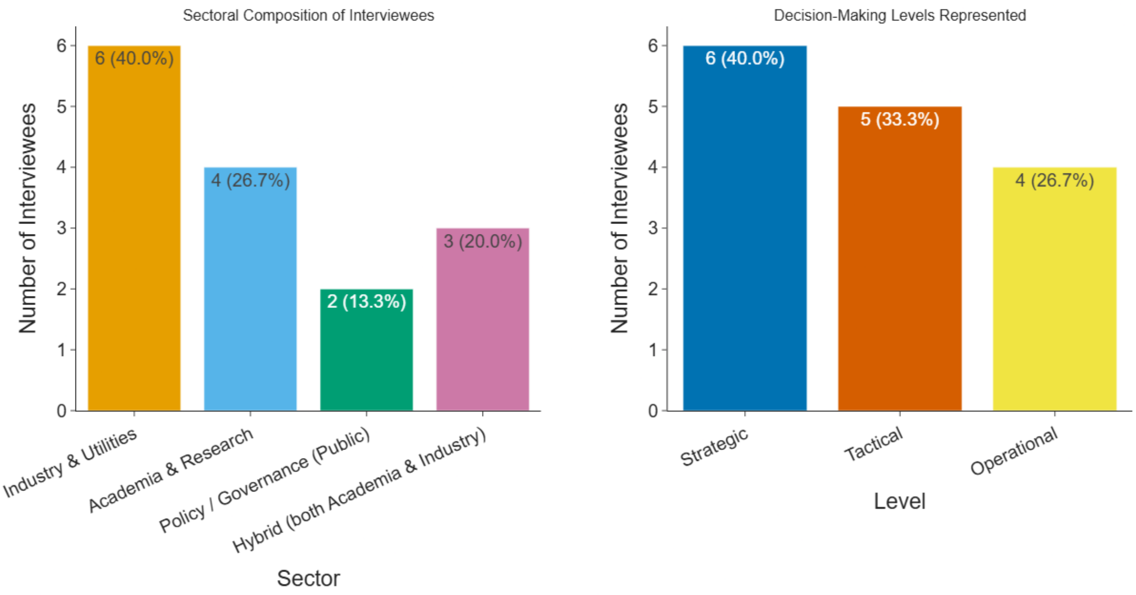

Figure 6.3: Sectoral composition of interviewees sample (left) and decision-making levels represented

Figure 6.3 shows the institutional and hierarchical diversity of the sample. To ensure consistency with the *Decision-Making Levels* framework introduced in Section 6.2, each participant was positioned according to the type of decisions they influence in their own organisational setting, and not according to seniority alone. While this distinction is intuitive for utilities, ministries, or companies, applying it to research institutions requires an explicit clarification. In universities, for instance, high-level academics typically define research directions and allocate resources, whereas early-career researchers carry out the hands-on work within those strategic boundaries.

Most interviewees operate in *industry and utilities* (40%), followed by *academia and research* (27%), with smaller but significant representation from *policy and governance* (13%) and *hybrid profiles* working across sectors (20%). Decision-making levels were also well distributed:

- 40% of participants act at the *strategic* level. These are, for example, university professors, members of companies' Boards of Directors, and national-level policymakers shaping long-term research or policy strategies.
- 33% of the sample is involved in decisions at the *tactical* level. These people are mid-level professionals, engineers, project managers, and policy advisors who translate strategic goals into actionable frameworks and projects.
- 27% of the sample acts on the *operational* level. This is the case of early-career engineers, operators, PhD and EngD students, who are directly executing or researching technical work and performing on-site control tasks.

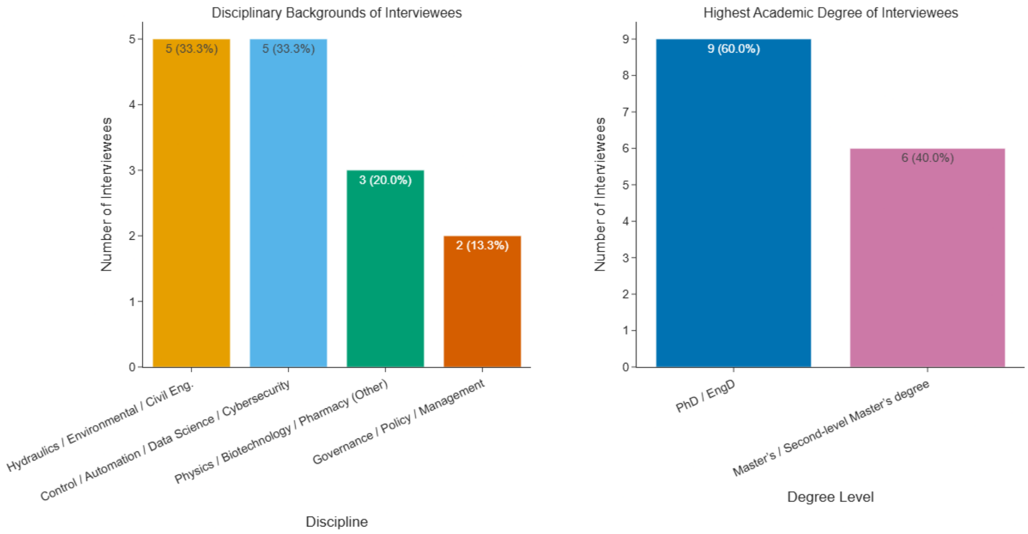

Figure 6.4: Academic and disciplinary profiles of the interviewed stakeholders.

Figure 6.4 summarises the academic and disciplinary composition. Two-thirds of participants have a background in engineering disciplines, either *hydraulics, environmental, or control and automation engineering*, while others contribute expertise in *physics, biotechnology, pharmacy*, and *governance or management*. The high academic qualification of the sample is evident, with 60% holding a PhD or EngD and the remaining 40% holding Master's or post-laurea degrees (e.g. MBAs).

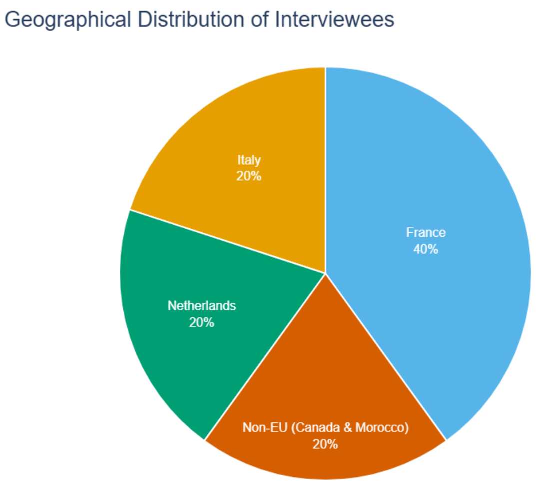

Figure 6.5: Geographical distribution of interviewees by country of work.

Finally, Figure 6.5 illustrates the geographical distribution of the interviewees by country of work. The majority are based in *France* (40%), with additional representation from *Italy* (20%) and the *Netherlands* (20%), complemented by two extra-European perspectives from *Canada* and *Morocco* (20%), used to expand the idea horizon and compare Europe to different realities. This distribution ensures coverage of different regulatory, infrastructural, and organisational contexts within and beyond the EU.

Overall, the stakeholder sample reflects a balanced sample of actors across institutional, disciplinary, and geographical dimensions, allowing the analysis to integrate technical feasibility considerations with policy and governance perspectives in the transition towards AI-enabled water management systems.

Table 6.2: Overview of the interview sample by sector, decision-making level, disciplinary background, and country of work.

Interview	Sector	Decision-Making Level	Disciplinary Background	Highest Degree	Country
Interview 01	Industry & Utilities	Operational	Hydraulics/ Environmental/ Civil Eng.	Master's/ Second-level Master's	France
Interview 02	Industry & Utilities	Operational	Hydraulics/ Environmental/ Civil Eng.	Master's/ Second-level Master's	France
Interview 03	Academia & Research	Strategic	Physics/ Biotechnology/ Pharmacy	EngD/ PhD	Morocco
Interview 04	Industry & Utilities	Tactical	Control/ Automation/ Data Science/ Cybersecurity	Master's/ Second-level Master's	France
Interview 05	Industry & Utilities	Strategic	Hydraulics/ Environmental/ Civil Eng.	Master's/ Second-level Master's	Italy
Interview 06	Hybrid	Strategic	Hydraulics/ Environmental/ Civil Eng.	EngD/ PhD	Italy
Interview 07	Hybrid	Strategic	Physics/ Biotechnology/ Pharmacy	EngD/ PhD	Italy
Interview 08	Hybrid	Operational	Hydraulics/ Environmental/ Civil Eng.	EngD/ PhD	Netherlands
Interview 09	Academia & Research	Strategic	Control/ Automation/ Data Science/ Cybersecurity	EngD/ PhD	Canada
Interview 10	Industry & Utilities	Tactical	Physics/ Biotechnology/ Pharmacy	EngD/ PhD	France
Interview 11	Academia & Research	Tactical	Control/ Automation/ Data Science/ Cybersecurity	EngD/ PhD	Canada
Interview 12	Academia & Research	Operational	Control/ Automation/ Data Science/ Cybersecurity	EngD/ PhD	Netherlands
Interview 13	Industry & Utilities	Tactical	Control/ Automation/ Data Science/ Cybersecurity	EngD/ PhD	France
Interview 14	Policy & Public Governance	Strategic	Governance/ Policy/ Management	Master's/ Second-level Master's	France
Interview 15	Policy & Public Governance	Tactical	Governance/ Policy/ Management	Master's/ Second-level Master's	Netherlands

6.5. Interview Design and Protocol

The selected people were interviewed according to a rigorous design of the interview guide, which ensures, as proposed by [118], methodological transparency and enhances the credibility of qualitative findings by systematically linking interview themes to research objectives. The semi-structured format also allows the interviewer to probe unanticipated topics while maintaining thematic coherence across interviews, an approach particularly suited to multi-actor systems where perspectives vary in expertise, institutional position, and strategic interest.

The interviews aimed to explore four interrelated themes: *technical feasibility*, *trust and explainability*, *incentives and barriers to adoption*, and *institutional and stakeholder alignment*. These themes were derived from the Transition Model Canvas (TMC) dimensions (technological, organisational, institutional, and societal), aiming to assess the system's status and preparedness while also finding strategies that the *niches* can implement to strengthen themselves and destabilise the *incumbent regime*. In line with the goals of this research, the guide sought to elicit both factual assessments (e.g., technical feasibility, regulatory conditions) and value-oriented reflections (e.g., perceived risks, trust, and societal desirability). The full interview question list is provided in Appendix C.

All the interviews were conducted online, one-on-one, and conversations were carried out on Microsoft Teams in English, Italian, or French according to the interviewee's preference, to avoid language-related misunderstandings and ensure conceptual clarity. Each interview lasted between 35 and 70 minutes and was recorded (with consent) using Microsoft Teams.

Ethical procedures were strictly followed according to the TU Delft Human Research Ethics Committee (HREC) guidelines. The research obtained HREC approval on 22-Sep-2025 (application ID: 6015). All participants received an informed-consent form explaining the project's purpose, data-handling procedures, anonymity, and voluntary participation. Consent to record and process data for research purposes was explicitly obtained before each interview, and the full informed consent form is available upon request.

6.6. Interpretation of Qualitative Data: Coding & Analysis

Microsoft Teams transcripts of the interviews were anonymised and manually adjusted to correct minor mistakes, accordingly to what was said in the original interview. Interviews conducted in Italian and French were translated into English using DeepL and subsequently reviewed by a linguistically proficient evaluator (native Italian speaker with working proficiency in French). This verification step ensured that conceptual nuances and domain-specific terminology were accurately preserved during translation.

After this preprocessing, the qualitative data were analysed by using ATLAS.ti, using Braun and Clarke's six-phase thematic analysis framework, which includes (1) familiarisation with data, (2) generating initial codes, (3) searching for themes, (4) reviewing themes, (5) defining and naming themes, and (6) writing the report [119].

The coding was done on the processed data, but with constant referencing to the original transcript to double-check translation correctness.

A total of 15 semi-structured interviews were coded using a hybrid deductive-inductive approach. Deductive codes were derived from the four primary themes, as described in Section 6.5.

In parallel, inductive codes were allowed to emerge iteratively through close reading of the transcripts, leading to new subthemes such as "future perspectives" or "AI misconception".

Each meaningful sentence or concept was assigned one or more thematic codes. The codes were then grouped into main categories that were then used as section titles in the result Chapter, specifically in Section 7.3. Coding was conducted directly on the English translations of each interview, with all quotes timestamped and anonymised to preserve confidentiality. This structured coding framework enabled the extraction of stakeholder concerns and expectations, which were then synthesised and later used to inform the policy analysis in Chapter 8.

7

Policy Analysis Case Study and Results

7.1. Policy Analysis Case Study

Similar to what was said for the engineering case study description, the next subsections explain the case study considered for the policy dimension of this research. As introduced in Section 1.1, the focus of this research is on European countries, and particularly on France, Italy, and the Netherlands. Each country's water governance involves a complex stakeholders' system and some peculiarities in its organisation, allowing for a European focus and a comparison between different realities.

A more detailed explanation of how the water governance works in each of these countries is provided using power/interest matrices to frame the actor system. In stakeholder-management frameworks, "power" is defined as the stakeholder's ability to influence decision-making processes (e.g., policy, investment, regulation), while "interest" refers to the stakeholder's level of concern or involvement with the outcome under consideration (e.g., organisational success, service delivery). The conceptual origins of this two-dimensional grid, which is now widely adopted for mapping stakeholder engagement strategies, can be traced back to the seminal work by [120], later reviewed in depth by [121].

In the case considered, power and interest must be considered with respect to the sociotechnical transition, which is, in this case, a renewal of the water critical infrastructure to integrate modern AI-based controllers. "Power" refers to the ability to influence or control water policy and infrastructure decisions, while "Interest" denotes the degree of stake or focus on water services and regulation. So the specific questions that one needs to ask oneself to place an actor on the grid are: "How interested is such an actor in this transition? What benefits/concerns does it have? And how much power does it have to influence the sociotechnical system and therefore the transition?".

In the following subsections, power/interest matrices are provided to map the water governance in each of the three countries object of study, together with a ranking of each actor's influence, and a discussion of their roles. These maps include key stakeholders in broad water governance (e.g. drinking water, wastewater, flood management, etc.) and emerging actors (innovators, pilot projects, new institutions), which are classified based on the matrix quadrant they fall into.

7.1.1. Mapping the Italian Governance System

The power interest matrix in Figure 7.1 represents the stakeholders for the Italian water governance. On the next page, it follows an overview of the system and an explanation of the key stakeholders represented.

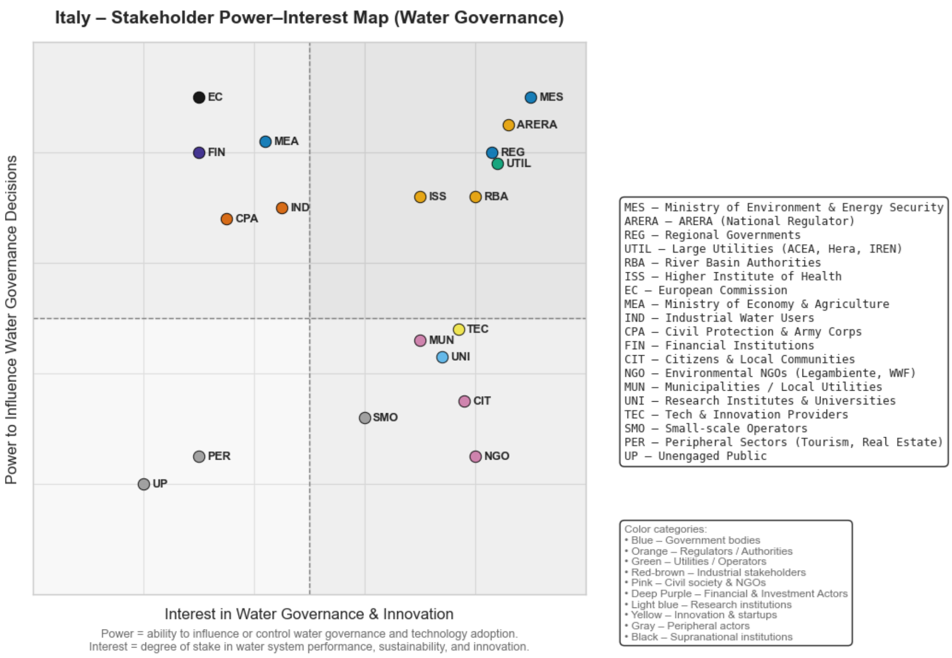

Figure 7.1: Stakeholder power–interest grid for Italy

In Italy, water governance is decentralised and fragmented across national, regional, and local levels. Reforms in the 1990s introduced the Integrated Water Service (*Servizio Idrico Integrato*) to consolidate water supply and sanitation within Optimal Territorial Areas (*Ambiti Territoriali Ottimali*, ATO) grouping many municipalities [122]. A national independent regulator (since 2011 the Regulatory Authority *ARERA*) oversees service standards and tariffs. Despite these reforms, implementation is incomplete since the number of water service operators has decreased from about 7,800 in 1999 to about 2,110 in 2022, still indicating significant fragmentation in several regions [123]. This fragmentation and multi-level structure mean that influence is spread among various actors.

High Power / High Interest Stakeholders in Italy

In this quadrant of the matrix, there are the key players, those who can really influence the system and have a great interest in its performance. An explanation of their role is reported below in Table 7.1.

Table 7.1: Stakeholders in the High Power / High Interest Quadrant (Italy).

High Power / High Interest	
Ministry of the Environment and Energy Security (MES)	Formerly *Ministry of Ecological Transition*, this actor sets national water and environmental policies and transposes EU directives (WFD, UWWTD). It wields legislative power and oversight interest in water services and resources, historically shaping water quality standards and tariff guidelines. Today, it coordinates national strategy (e.g. drought plans, water infrastructure funding) and leads Italy's input to EU water policies [124, 125].
National Regulatory Authority (ARERA)	Regulates integrated water services since 2011. It has high power to set tariffs, technical standards, and monitor utilities, as well as a high interest in efficient and sustainable water service delivery. ARERA's role brought more uniform oversight after years of local tariff-setting; it aims to reduce the historically fragmented approach [126]. Its decisions strongly influence utilities' investment and operation decisions.
Regional Governments (REG)	Hold constitutional powers over environmental planning, water-resource management. They translate national frameworks to regional plans and often host regional environmental agencies. Regions create or supervise governing bodies (*Autorità d'Ambito*) responsible for planning and controlling local water services. Thus, regions have high power and strong interest due to environmental and territorial development outcomes [125].
Large Integrated Water Utilities (UTIL)	Major companies operate water and wastewater services for major cities and areas, giving them high stakes. Acea S.p.A. is a prime example, managing integrated water supply and sewage for 11 million people in Italy. Other examples include HERA, IREN, A2A, and SMAT. These utilities have high interest (water service is a core business) and considerable power both operationally and economically. Their scale gives them weight in national discussions on water infrastructure [127].
River Basin Authorities (RBA)	River Basin Authorities (*Autorità di Bacino Distrettuale*) have a high interest in water. These actors coordinate water resource management and pollution control plans at the basin scale, having planning power and scientific expertise. They ensure that utilities and regions meet water quality and ecological flow targets, directly influencing infrastructure needs (e.g. by requiring WWTPs upgrades to meet river health objectives) [124, 125].
Higher Institute of Health (ISS)	The Higher Institute of Health (*Istituto Superiore di Sanità* ISS) oversees and monitors drinking and bathing water safety. Its suggestions and analyses influence political decisions and regulatory standards.

High Power / Low Interest Stakeholders in Italy

These are powerful influencers with a secondary focus. They are not directly interested in Italian water governance, but they have the power to influence it. Table 7.2 briefly explains their roles.

Table 7.2: Stakeholders in the High Power / Low Interest Quadrant (Italy)

High Power / Low Interest	
European Commission (EC)	The European Commission shapes Italian water governance mainly through the Water Framework Directive (2000/60/EC) and the Urban Waste Water Treatment Directive (91/271/EEC), exerting influence via legislative, financial, and enforcement mechanisms. Recent infringement actions over Italy's non-compliance with wastewater obligations highlight this leverage, including financial penalties for several agglomerations [128, 129, 130]. While the EU's power is high, through rule-making and sanctions, its interest in Italian day-to-day water policies remains moderate.
Ministry of Economy and Agriculture (MEA)	Influence water governance indirectly through budgets, incentives, agricultural rules, and infrastructure financing. High resource-mobilisation power but low sustained interest, subordinate to fiscal and agricultural priorities.
Large Industrial Water Users (IND)	Energy, chemical, and agro-food industries can lobby on abstraction licences, discharge limits and pricing frameworks. Their interest peaks in regulatory changes or when scarcity affects operations, but water policy is not their core mission.
Civil Protection Authorities (CPA)	The *Dipartimento della Protezione Civile* exercise strong emergency powers during floods, droughts and dam-safety events. Outside emergencies, their involvement in water governance is limited [131].
Financial Institutions (FI)	Major investments in Italian infrastructure often hinge on conditions set by financiers such as the European Investment Bank and, domestically, Cassa Depositi e Prestiti. Recent loans to Acea and regional utilities demonstrate how funding can steer resilience upgrades, leakage reduction and compliance with EU sustainability goals, evidencing high power via conditional finance despite a lower interest in governance processes [132, 133, 134].

Low Power / High Interest Stakeholders in Italy

This quadrant contains stakeholders who are active with but have limited influence. Among those, the following groups can be identified, as listed in Table 7.3.

Table 7.3: Stakeholders in the Low Power / High Interest Quadrant (Italy).

Low Power / High Interest	
Citizens and Local Communities (CIT)	These have a high stake in affordability, quality, and continuity of water services, as well as in broader environmental outcomes. Italy has a tradition of civic engagement in water governance, as illustrated, for example, by the 2011 national referendum, in which an overwhelming majority voted to keep public control over water services and prevent profit-driven management. Despite this, individual citizens have limited institutional power, and their influence is primarily exerted through public opinion, referenda, participation in basin planning consultations, or involvement in agreements such as *Contratti di Fiume, i.e. River Contracts* [135].
Environmental NGOs (NGO)	Organisations such as *Legambiente*, *WWF Italia*, and the *Forum Italiano dei Movimenti per l'Acqua* advocate for the protection of water quality, aquatic ecosystems, and the human right to water. They also promote integrated and participatory governance through social mobilisation and information dissemination. Although they possess limited power, these NGOs can shape the policy agenda by mobilising public support, initiating legal actions, and engaging in basin-level consultations. In recent years, their successful campaigns highlighted gaps in wastewater treatment compliance and ecological flow protection.[136].

Continued on next page

Low Power / High Interest (continued)	
Municipalities and Local Utilities	Small/medium municipalities and in-house providers often fall into this quadrant. In many regions, especially in southern Italy, the integration of the *Servizio Idrico Integrato* (Integrated Water Service) remains incomplete. Consequently, local governments retain a strong operational interest in ensuring the provision of safe drinking water and wastewater services but lack financial and regulatory leverage when acting individually. However, collective organisations such as the National Association of Italian Municipalities (ANCI) and inter-municipal consortia can amplify representation [135, 126].
Research Institutes and Universities	Italian research bodies, including *IRSA-CNR* (the Water Research Institute of the National Research Council), *Politecnico di Milano*, and several universities (e.g., Genoa, Padova, Bologna), demonstrate a high interest in water systems from both technical and governance perspectives. They contribute to innovation in water treatment, digital management, and control systems, and frequently provide policy advice through commissioned studies. Although their direct power is limited, they influence governance through knowledge production, pilot demonstrations, and EU-funded projects (e.g., Horizon Europe, LIFE, and Water4All initiatives) [136].
Tech and Innovation Providers	Firms supplying digital solutions, data platforms, and analytics tools for the water sector show high interest in the Italian market due to utilities' increasing need for digitalisation and performance optimisation. However, evidence from a recent Italian digital-transformation case study demonstrates that technology providers operate in a supporting role: successful implementation depends on close collaboration with utility staff, and the most demanding tasks require full alignment with utility-defined data structures, operational procedures, and business priorities. As external actors, they contribute technical expertise but cannot influence regulatory requirements, tariff structures, or investment decisions, which positions them as high-interest yet structurally low-power stakeholders [137].
Small-Scale Operators (SMO)	This category includes individuals or small entities (e.g., private well owners, small farmers outside consortia, small businesses) which are directly affected by regulatory and environmental changes (e.g. tariff adjustments or contamination events) without having the ability to influence them. These actors generally exhibit low interest in water governance beyond the expectation of a reliable and affordable supply, and they possess minimal institutional or lobbying power. They seldom participate in planning or decision-making processes and are usually represented indirectly through agricultural or trade associations [138, 139].

Low Power / Low Interest Stakeholders in Italy
In this quadrant, there are the peripheral actors, with low power and low interest, as shown in Table 7.4.

Table 7.4: Stakeholders in the Low Power / Low Interest Quadrant (Italy).

Low Power / Low Interest	
Peripheral Economic Sectors (PER)	Sectors like tourism, real estate, and transport typically show low interest in water governance except during events that affect their economic activity. Their power is therefore situational and limited, and their interest in governance remains typically low. However, as climate variability increases, these sectors are expected to become more exposed to water-related risks and may require stronger coordination with authorities [140, 141].

Continued on next page

Low Power / Low Interest (continued)

Unengaged Public (UP)	A large portion of the Italian population remains disengaged from water-governance issues except during crises. Urban residents, in particular, tend to take access to safe water for granted and are often unaware of the roles of utilities, regulators, or basin authorities. From a governance perspective, this is problematic, as public scrutiny and informed participation are central to improving transparency and responsiveness in the water sector. Recent surveys show that less than 30% of Italians feel well informed about local water management structures [142, 138]. Addressing this awareness gap through education, communication, and participatory initiatives remains an important challenge for achieving more inclusive governance.

Notes and Emerging Trends in Italy

Italy's water sector is still addressing historical governance gaps. One key challenge identified is improving stakeholder engagement. Efforts like the adoption of River Contracts (*Contratti di Fiume*) for participatory basin planning and the spread of water-oriented living labs are aiming to transform passive stakeholders into active co-managers of water resources. Additionally, Italy's governance is marked by north–south disparities: some southern regions still lack fully consolidated service operators, meaning fragmentation is region-specific (Calabria, Campania, Molise, and Sicily are often cited as having persistent fragmentation) [123]. Meanwhile, large utilities (mainly in central/northern Italy) push for modernisation and economies of scale. These dynamics suggest that stakeholder alignment and capacity-building, especially in weaker regions, will remain a focus in the coming years.

7.1.2. Mapping the French Governance System

The power interest matrix in Figure 7.2 shows the stakeholders for the French water governance system, together with an explanation for the actor's positioning in the map.

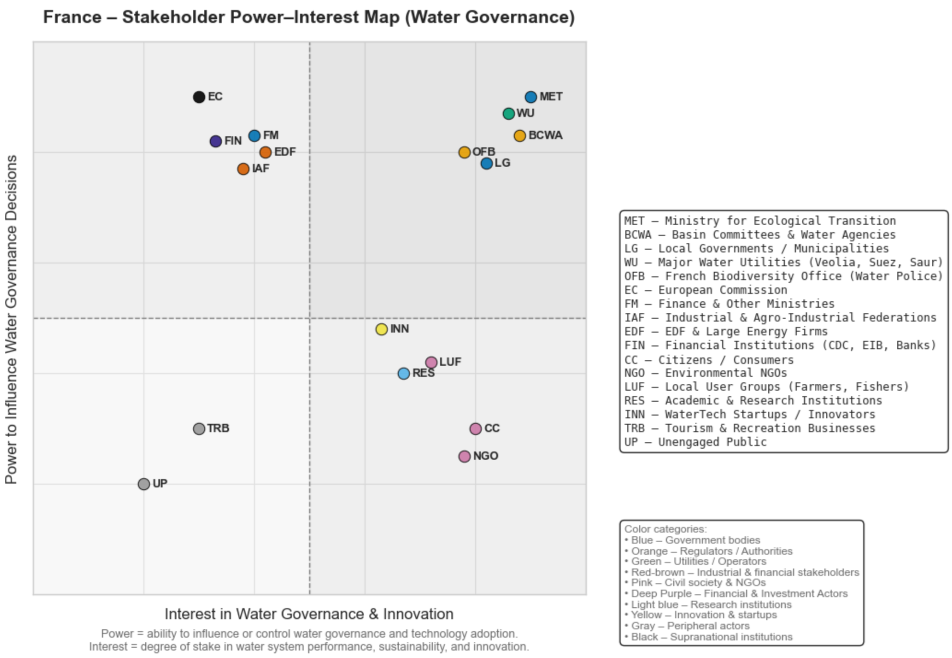

Figure 7.2: Stakeholder power–interest grid for France

France's water governance is often cited as a model of basin-oriented, multi-level governance [143]. Since the Water Law of 1964, France has divided its mainland territory into six large river basin districts, each with a Basin Committee (a "Water Parliament") and a Water Agency. These institutions practice integrated, participatory management at the river basin level, financing water projects and coordinating stakeholders' input. Meanwhile, responsibilities for drinking water supply and sanitation lie with municipalities or their groupings, which often delegate to or partner with specialised water companies. This creates a shared power structure: national authorities set overall policy and ensure EU directives are implemented; basin agencies plan and fund projects; and local authorities and companies implement services. In the following subsections, an explanation of each quadrant of the matrix can be found.

High Power / High Interest Stakeholders in France

The stakeholders in this quadrant are explained in Table 7.5, below.

Table 7.5: Stakeholders in the High Power / High Interest Quadrant (France).

High Power / High Interest	
Ministry for Ecological Transition (MET)	The *Ministère de la Transition Écologique* defines national water policy, implements environmental regulation, and transposes EU directives. Holds high legislative and regulatory power and a strategic interest in water resources, aquatic ecosystem health, and public water services. The ministry issues national water legislation (e.g. the Water Law 1992, the Water and Aquatic Environments Law 2006), oversees nationwide programs, and supervises enforcement agencies (see OFB below). In essence, the Ministry sets the overall goals that all other stakeholders work towards.
Basin Committees and Water Agencies (BCWA)	Each major river basin has a Basin Committee (100 members representing the state, local authorities, users, NGOs) and a corresponding Water Agency (a public financial institution). The Basin Committee deliberates and adopts a strategic plan: the Master Plan for Water Development and Management (*SDAGE*). The Water Agency raises funds via water abstraction and pollution fees and finances projects to meet the SDAGE goals. They have high interest and high power: while agencies cannot mandate projects, their funding power and expertise strongly influence actions.
Local Governments (e.g., Communes and Inter-Municipal Syndicates, LG)	Communes (municipalities) in France traditionally hold responsibility for drinking water supply and sanitation services, resulting in high interest. Often, small ones are aggregated into syndicates or metropolitan authorities for higher efficiency. They can operate services directly (*régie municipale*) or contract them out to water companies. Collectively, they influence a huge part of the water sector (e.g., Greater Paris's water syndicate (SEDIF)). Given their number, their power can be fragmented, even if large urban governments are highly influential in setting standards and rates.
Major Water Utility Companies (WU)	France is home to global leaders in the water industry (Veolia, Suez, and Saur). These firms have high interest and high power in practice, operating thousands of infrastructures under municipal contracts. They also influence policy through associations and dialogue with regulators (e.g. providing input for performance indicators or standards due to their know-how). While they do not make laws, their ability to implement changes, or conversely their warnings about impractical rules, carries weight. Additionally, they invest heavily in R&D, driving trends like digital water and resource recovery.
Office Français de la Biodiversité (OFB)	This national agency, formed by merging the water police (*ONEMA*) and various biodiversity offices in 2020, acts as France's enforcement and ecological expertise arm in water and environment matters. OFB has a high interest in healthy aquatic ecosystems and proper water management, and holds significant power in enforcing environmental regulations. Its agents can inspect facilities, issue fines, and ensure compliance with laws. While water is only part of its broader biodiversity mission, OFB's authority and expertise make it a key player in implementing water policy on the ground.

High Power / Low Interest Stakeholders in France

In the French system, several actors wield substantial influence over water policy and investments, but do not have water governance as their primary or continuous focus. They intervene through broader regulatory, sectoral, or financial mandates, which justifies their positioning, as explained in Table 7.6.

Table 7.6: Stakeholders in the High Power / Low Interest Quadrant (France).

High Power / Low Interest	
European Commission (EC)	Similar to Italy, the EU holds high power over France's water policy through its directives and enforcement mechanisms. However, its interest is broad and not specifically tailored to French local circumstances. The EU acts as a powerful background regulator whose intervention becomes salient mainly when compliance gaps arise.
Finance & Other National Ministries (FM)	Central government ministries beyond the environment portfolio (e.g. finance, regional development, agriculture) influence water governance indirectly through budget transfers, territorial reforms, and sectoral policies. Colon et al. show how the state has progressively repositioned itself as a regulator and coordinator, while local authorities and basin institutions implement water policies on the ground [143]. These ministries thus retain significant *power* over resources, rules and territorial organisation, but their *interest* in water is subordinate to broader fiscal, economic, or regional-development objectives.
Industrial & Agricultural Federations (IAF)	Large industrial and agricultural users (e.g. energy producers, manufacturing sectors, irrigation interests) account for a substantial share of water abstractions in France, especially for energy cooling and irrigation [143]. Their federations exert *power* via lobbying and representation in basin committees or professional bodies, particularly when regulations or allocation rules affect production costs. However, their *interest* is episodic and predominantly instrumental: water governance is one constraint among many, rather than their core mission.
EDF and Large Industrial Water Users (EDF)	Companies such as EDF and other major industrial/bottling users depend on reliable access to water for energy production, cooling, or industrial processes. Colon et al. document the importance of energy and industrial withdrawals and the way sectoral disputes can trigger regulatory adjustment [143]. These actors can influence regulatory debates and investment priorities when new norms or scarcity conditions threaten operations, indicating high *power*. Yet their *interest* in water governance is focused on securing permits and predictable rules, not on continuous engagement with broader water-policy objectives.
Financial Institutions (FIN)	Within the French model, local authorities responsible for water and sanitation services finance operations and investments largely from user charges, but they are also "entitled to resort to bank loans and public funds to finance investments" [143]. Commercial banks and public financial institutions thus wield *high power* at the project level, as their lending conditions shape the feasibility and timing of infrastructure upgrades. Nevertheless, they do not participate in day-to-day water-policy deliberation or basin planning; their *interest* is confined to the bankability and risk profile of individual projects rather than to the substance of water governance.

Low Power / High Interest Stakeholders in France

These stakeholders, as explained in Table 7.7, are the ones that do not have high power in water governance, but are heavily affected by the decisions in the field.

Table 7.7: Stakeholders in the Low Power / High Interest Quadrant (France).

Low Power / High Interest	
Local Citizen Groups and Users Associations (CC)	Citizens and local user groups display strong interest in water issues—service continuity, pricing, pollution incidents, and drought responses. In cases of perceived environmental or service failures, local groups often form to demand greater transparency or improved management, as illustrated by recurrent disputes around river pollution or water shortages. Colon et al. highlight that such disputes have historically acted as catalysts for institutional adjustments in France, demonstrating the high engagement capacity of local communities [143]. However, their formal decision-making power remains limited, relying primarily on public mobilisation, opinion pressure, and participation in local consultation processes.
Environmental NGOs (NGO)	Environmental organisations such as France Nature Environnement, WWF France, and river-protection associations hold a high interest in water management, river restoration, and pollution control. They participate in basin committees and local water commissions, contributing technical knowledge and environmental advocacy. While lacking statutory authority, NGOs can influence governance by mobilising the public, engaging in litigation, and contesting projects that threaten aquatic ecosystems. [143] show that conflicts and NGO mobilisations have been key drivers of regulatory evolution in France, reinforcing their influence despite limited formal power.
Local User Groups: Farmers and Fishers (LUF)	Farmers and small-scale irrigators have a high interest in water allocation rules, restrictions, and pollution regulations, as agricultural productivity strongly depends on water availability. Individually they possess limited power, but collective representation through bodies such as FNSEA, Chambers of Agriculture, and local river committees grants them structured participation channels. Colon et al. emphasise that conflicts between agricultural and environmental interests have historically shaped basin planning and regulatory reforms [143]. Complementarily, Amblard's research demonstrates that successful collective action in France depends on institutional support and the alignment between rules, local contexts, and perceived fairness, influencing farmers' compliance with water-quality measures [144].
Academic and Research Institutions (RES)	Universities, engineering schools, and research bodies contribute high-level expertise to water-management innovation, monitoring, and ecological assessment. Their interest is high due to longstanding involvement in modelling, treatment technologies, and evaluation of governance reforms. Colon et al. describe how scientific input has supported the shift toward integrated basin planning and the modernisation of water governance instruments [143]. Despite this, research institutions have limited formal power, influencing governance indirectly through knowledge production, studies commissioned by agencies, and participation in consultative processes.
WaterTech Startups and Innovators (INN)	Digital water and WaterTech firms (e.g. remote-sensing, leak detection, analytics providers) have a high interest in the sector due to growing demand for monitoring, optimisation, and compliance tools. However, Colon et al. show that France's water governance system is dominated by the "triangle" of State-Agencies-Local Authorities [143], leaving limited space for new entrants. Startups, therefore, depend on adoption by utilities and municipalities and have low direct influence over governance or regulatory design. Their interest is strong, but their structural power remains low.

Low Power/Low Interest Stakeholders in France

These are the peripheral actors, explained in Table 7.8, in the following page.

Table 7.8: Stakeholders in the Low Power / Low Interest Quadrant (France).

	Low Power / Low Interest
Small Private Water Operators (SPO)	A small number of private operators continue to manage drinking water or sanitation services in specific niche contexts (e.g., small resorts, rural communes, or tourism facilities), typically through short *affermage* or service contracts, where municipalities delegate daily operation but retain responsibility for investment and strategic decisions. As a result, these actors have *low power* in governance terms and limited strategic relevance. Evidence from the remunicipalisation trend shows that their market share has declined as many communes have reverted to public or inter-municipal *régies*, reinforcing their residual and unstable position in France's water landscape [145]. Their *interest* in broader water governance remains low, focused primarily on contractual service obligations.
General Public (Unengaged) (UP)	While French citizens express strong preferences regarding the quality and affordability of water services, a large share of the population remains weakly engaged in water governance. [143] highlight that despite France's long-standing consultative structures, everyday users rarely participate in water-policy processes unless mobilised by disputes or crises. Outside these episodes, most residents do not follow governance debates or interact with basin institutions. Consequently, they hold *low power* and *low sustained interest*, becoming visible mainly during high-profile events.

Notes and Recent Developments in France

Over the decades, France's basin management model has generally been successful in improving water outcomes while balancing stakeholder voices. However, challenges remain. Climate change is causing more frequent droughts and floods, testing the limits of existing governance arrangements. There are ongoing discussions about creating more storage for water supply versus preserving rivers, which pit agricultural and environmental interests against each other in new ways. Additionally, the 2019 merger of Suez's water activities into Veolia (completed in 2022) has raised questions about maintaining competition and innovation in the water services sector. France is also working to enhance transparency and citizen engagement; for example, recent "Assises de l'eau" (Water Conferences) have been held to gather input from a wide array of stakeholders on future water policy. Overall, the French system continues to evolve, but its core principle of basin-level co-management remains a reference point in global water governance discussions.

7.1.3. Mapping the Dutch Water Governance System

The Netherlands is internationally recognised for its advanced and highly institutionalised water governance system, particularly regarding flood protection, delta management, and integrated water resources planning. A distinctive feature of the Dutch model is the presence of the *Regional Water Authorities* (*Waterschappen*), which are functional local governments responsible exclusively for water management, including flood defense, regional water-level regulation, and wastewater treatment. These water boards are among the oldest democratic institutions in Europe, with origins tracing back to the Middle Ages, and remain directly accountable through dedicated elections. They possess independent fiscal authority, levying water management taxes on residents and businesses, and operate with a high degree of autonomy within the national legal framework [146].

At the national level, the Ministry of Infrastructure and Water Management (*Ministerie van Infrastructuur en Waterstaat*) sets overarching water policy and supervises large-scale projects through its executive agency, *Rijkswaterstaat*. The ministry ensures coordination with European directives and manages interregional planning instruments like the National Water Program (*Nationaal Water Programma*) 2022–2027 [147]. Rijkswaterstaat itself is responsible for the main rivers, major canals, coastal zones, and the operation of storm surge barriers, forming the backbone of national flood defense and hydraulic infrastructure.

Dutch water governance is characterized by institutionalized collaboration and consensus building—the so-called *polder model*. Coordination occurs through formal administrative agreements between government levels and water authorities, such as the *Bestuursakkoord Water* (Administrative Agreement on Water). This cooperative framework aims to achieve efficiency, resilience, and public participation while balancing competing interests of safety, ecology, and spatial development. Figure 7.3 presents the stakeholder power–interest grid for the Dutch system, followed by a detailed explanation of the main actors involved.

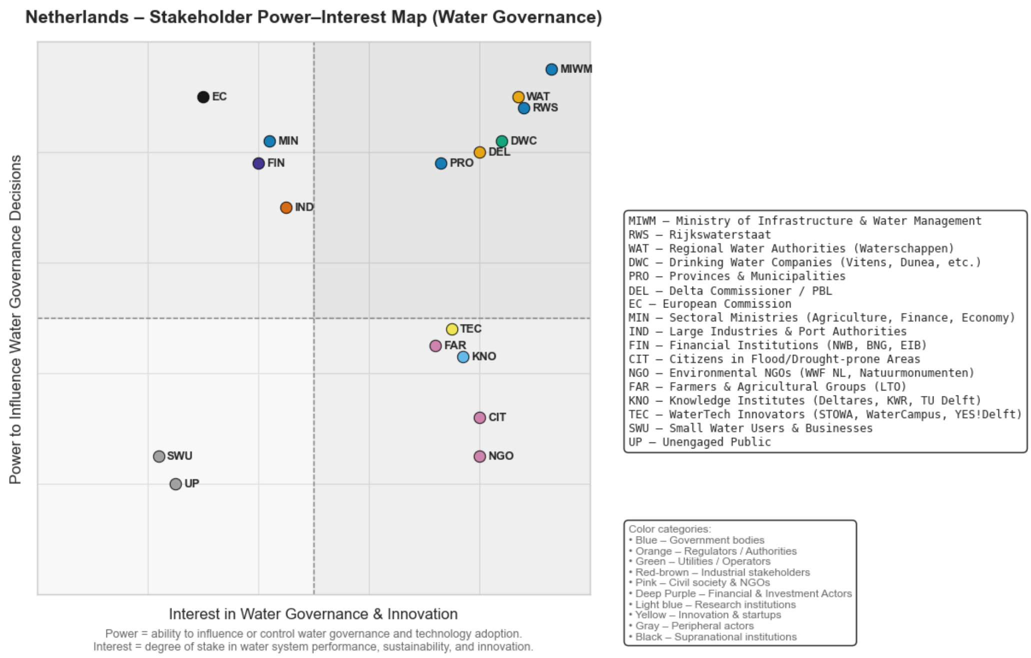

Figure 7.3: Stakeholder power–interest grid for the Netherlands

High Power / High Interest Stakeholders in the Netherlands

Table 7.9: Stakeholders in the High Power / High Interest Quadrant (Netherlands).

High Power / High Interest	
Ministry of Infrastructure and Water Management (MIWM) & Rijkswaterstaat (RWS)	The national government, through MIWM and its executive agency RWS, sets national water policy (e.g. the National Water Plan, the Delta Programme) and implements EU directives. RWS manages the main rivers, canals, dikes, and flood-defence infrastructure, reflecting the existential importance of flood safety for the Netherlands. These actors coordinate drought responses, oversee long-term climate adaptation, and ensure alignment across administrative layers. Their interest is high given national safety, navigation, and freshwater supply priorities, while their power derives from regulatory authority and control over strategic assets.
Regional Water Authorities (Waterschappen, WAT)	The 21 water authorities are independent public bodies with taxation powers and legally mandated responsibilities for flood protection, water-level management, and urban wastewater treatment. The OECD describes them as the “backbone” of Dutch water governance [148]. Their fiscal autonomy, elected boards, and technical competence give them major influence at the regional level. With annual expenditures of roughly €2.7 billion, they are central investors in water infrastructure and innovation. Their “functional democracy” ensures stakeholder participation (farmers, industries, NGOs) reinforcing both their legitimacy and their strategic role in water governance.

Continued on next page

High Power / High Interest (continued)

Drinking Water Companies (Vitens, Dunea, Waternet, etc., DWC)	Ten publicly owned regional drinking water companies operate under strict licensing and oversight. They manage critical infrastructure for water extraction, treatment, and distribution, and hold strong influence due to their monopoly positions and essential public-service function. Through VEWIN (their national association), they shape policy debates, participate in administrative agreements (e.g. the 2011 *Bestuursakkoord Water*), and collaborate with provinces to protect groundwater resources. Their interest is high (service reliability and quality), and their practical power is substantial due to infrastructure ownership, technical expertise, and long-term planning roles.
Provinces and Municipalities (PRO)	Provinces regulate spatial planning, groundwater licensing, and exercise oversight over water authorities (e.g. approving their by-laws). They integrate water considerations into regional development, making them medium-to-high power actors. Municipalities manage sewerage and urban drainage, cooperate with water boards for wastewater treatment and flood prevention, and shape land-use decisions that determine local water retention capacity. While water is one of many responsibilities, their interest is heightened in areas facing drought, subsidence, or urban flooding. Municipal representation on water authority boards further embeds them in strategic water governance.
National Water Advisory Bodies (Delta Commissioner DEL/PBL)	The Delta Commissioner provides independent oversight for the multi-billion-euro Delta Programme, ensuring long-term continuity in flood safety and freshwater planning beyond political cycles. PBL (Netherlands Environmental Assessment Agency) delivers scientific policy analysis supporting strategic decisions. These expert bodies possess a high interest in water management and hold strong soft power: their reports and recommendations are highly influential and broadly followed by Parliament and ministries. Through evidence-based analysis, they help steer national priorities in flood-risk reduction, drought adaptation, and climate resilience.

High Power / Low Interest Stakeholders in the Netherlands

In Table 7.10, the influential but secondary-focus stakeholders of the Dutch water governance system are summarised.

Table 7.10: Stakeholders in the High Power / Low Interest Quadrant (Netherlands).

High Power / Low Interest

European Commission (EC)	As elsewhere, the EU sets key directives that the Netherlands must implement. The Netherlands often meets or exceeds EU requirements (it has been a frontrunner in many areas of water management), but EU law still provides the baseline and some pressure. It provides funding for cross-border projects (e.g., Meuse–Scheldt cooperation) but is not a daily participant in national water governance forums.

Continued on next page

High Power / Low Interest (continued)	
Other Ministries (MIN)	Ministries such as Finance, Economic Affairs and Climate (EZK), and Agriculture, Nature and Food Quality (LNV) influence water governance indirectly through sectoral policies. LNV shapes manure, fertiliser, irrigation, and land-use rules that directly affect water quality and availability. Finance determines the fiscal space in which water authorities operate. EZK steers innovation policy (e.g. Topsector Water & Maritime), positioning water technology within the broader competitiveness agenda. Their interest in water is secondary, but their power is high when setting frameworks that condition local governance. As observed by van Dijk, national ministries often "set overarching ambitions but rarely provide a clear division of roles or ownership," requiring coordination by regional actors [149].
Large Industries and Port Authorities (IND)	Industries in major zones (e.g., Rotterdam port, petrochemical hubs, large food-processing plants) depend on a reliable water supply and predictable discharge regulations. Their interest in governance is intermittent, but their relevance grants them influence on specific issues. Associations such as VEMW lobby for large industrial water users. While they do not shape long-term governance agendas, they co-finance projects and participate in consultations, especially when regulatory changes affect operations.
Financial Institutions (FIN)	Financial institutions such as the Nederlandse Waterschapsbank (NWB), BNG Bank, and the European Investment Bank (EIB) provide long-term loans that underpin investments in flood protection, wastewater treatment and climate adaptation. Their financing conditions and risk assessments decisively affect which projects water authorities, municipalities and drinking water companies can implement, giving them high structural power over investment trajectories. However, they do not participate in day-to-day water governance or basin planning, and their interest remains focused on the financial soundness of projects rather than the design of water policy.

Low Power / High Interest Stakeholders in the Netherlands

In Table 7.11, the active but low-power stakeholders in Dutch water governance are summarised.

Table 7.11: Stakeholders in the Low Power / High Interest Quadrant (Netherlands).

Low Power / High Interest	
Farmers and Agricultural Interests (FAR)	Farmers have a high interest in drainage, irrigation reliability, and water-quality regulations. While individual power is limited, collective influence through LTO Nederland and regional branches (e.g. ZLTO) remains significant. Tightening environmental obligations (WFD, nitrogen norms, drought measures) creates recurring tensions between top-down regulation and experiential knowledge. Evidence from Noord-Brabant shows farmers often perceive environmental monitoring as predetermined rather than collaborative [150]. Yet, negotiated, context-sensitive programmes—such as *Schoon Water voor Brabant*—demonstrate that co-designed measures can improve compliance and strengthen long-term cooperation.
Environmental NGOs (NGO)	Organisations such as Natuurmonumenten, WWF–Netherlands, and regional ecological federations advocate for river restoration, wetland protection, and biodiversity conservation. Although lacking statutory power, they exert influence through the Netherlands' consensus-based governance model, participating in water-authority roundtables and basin-level programmes such as *Room for the River*. They also strategically use the legal system to challenge environmental decisions and shape debates on nature-based solutions and agricultural runoff through public campaigns.

Continued on next page

Low Power / High Interest (continued)

Citizens and Residents' Associations in At-Risk Areas (CIT)	Residents living in flood-prone polders, drought-affected regions, or socio-economically vulnerable urban districts show consistently high interest in water governance because their safety, livelihoods, and property values depend directly on water-level and climate-adaptation decisions. Their formal power is limited, but influence is exercised through public consultations, objection procedures, legal appeals, and participation in water-board elections, a unique feature of Dutch institutional design. Recent research highlights frustration with voluntary drought-adaptation measures and a perceived gap between strategic plans and local lived realities. As Augustijn et al. (2025) observe, "the current policy mix raises concerns about whether voluntary instruments can steer drought adaptation if uptake remains low" [151]. Residents often self-organise into neighbourhood associations when confronted with recurrent pluvial flooding, rising groundwater, or drainage failures. Evidence from South Rotterdam shows that such groups increasingly link water insecurity to broader social inequalities, calling for more inclusive dialogue and co-production methods. As Esteban et al. (2024) emphasise, community workshops reveal the need to "start conversations within communities" to help municipal actors understand residents' diverse perspectives [152]. Although lacking statutory authority, these citizen networks act as persistent watchdogs and legitimacy-shaping actors, influencing the design or timing of dike-raising and water-storage projects.
Knowledge Institutions and Innovative Firms (KNO & TEC)	Research institutes (Deltares, KWR), universities (e.g. TU Delft), and water-tech innovators (STOWA, WaterCampus, YES!Delft) demonstrate high interest in advancing digital water management, nature-based solutions, and climate-adaptation strategies. Their influence is primarily indirect: they support governance through pilot projects, modelling expertise, advisory roles, and participation in innovation networks. While dependent on high-power actors (water boards, ministries) for adoption and scaling, these institutions act as essential drivers of experimentation, technological development, and long-term learning in the Dutch water sector.

Low Power / Low Interest Stakeholders in the Netherlands

These are the peripheral stakeholders, summarised in Table 7.12.

Table 7.12: Stakeholders in the Low Power / Low Interest Quadrant (Netherlands).

Low Power / Low Interest

Small Water Users & Businesses (SWU)	Small commercial actors (e.g. campsite operators, small manufacturing sites, or individual groundwater users) typically have limited involvement in Dutch water governance. Their interest remains low unless directly affected by disruptions (e.g., low water levels, discharge restrictions). They lack collective organisation, hold no formal representation in water-board governance, and exert little structural influence. Their interactions with authorities are mostly operational (permits, inspections) rather than strategic.
General Public (Apathetic) (UP)	A large portion of Dutch residents remain disengaged from water governance because the system works reliably and water risks appear well-managed. They pay water-board taxes and water bills but rarely participate in elections or consultations. Their interest increases only during crises and their power is low apart from aggregate voting, where low turnout periodically raises concerns about the democratic legitimacy of water boards.

Notes and Emerging Trends in the Netherlands

The Netherlands' long-established water governance framework provides an enabling environment for innovation. Water boards have clear mandates and motivation to adopt new solutions, knowledge institutions and companies are closely integrated into the sector, and multi-level governance coordination is strong (facilitated by instruments like the National Administrative Agreement on Water). One current trend is the increasing focus on climate adaptation: the traditional emphasis on flood defence is expanding to include drought management and nature-based solutions, which brings in new stakeholders (like urban planners and landscape NGOs) into water dialogues. Another development is the push for energy neutrality and circular economy in the water sector – for example, wastewater treatment plants are being transformed into "resource factories," and water boards aim to be energy self-sufficient. These goals require cooperation between water authorities, municipalities, and industry (for waste reuse and energy projects), further blurring the lines between sectors. Overall, the Dutch model shows a high capacity to incorporate new priorities, but it also faces future challenges such as maintaining its consensus approach in the face of climate pressures and ensuring continued public engagement (voter turnout and awareness) in water governance.

7.2. Final Considerations on Policy Analysis Case Study

Across Italy, France, and the Netherlands, water governance involves a diverse set of stakeholders with varying degrees of power and interest. Italy grapples with uniting a fragmented service landscape, making stakeholder engagement and regulatory cohesion critical. France leverages its basin-centric model to include many voices, balancing strong state and industry actors with active local and environmental input [143]. The Netherlands benefits from long-empowered regional water authorities and a culture of consensus, which together enable adaptive and innovative water management [148].

For a researcher or practitioner looking to deploy advanced control systems, these maps are crucial to highlight where and how such a project must interface with governance. First, align the initiative with the goals of high-power stakeholders, whether it is meeting the regulators' performance standards in Italy, contributing to basin agencies' environmental targets in France, or serving the water boards' efficiency and climate objectives in the Netherlands. Second, secure buy-in from high-interest groups. This could mean informing and educating local communities and NGOs about the benefits of the new technology, as well as training utility staff and assuring all stakeholders that the innovation is not a threat but a value-add. Third, navigate the multi-level dynamics, understanding that national and EU incentives (funding, directives) can catalyse local action or, if overlooked, become barriers. Each country demonstrates the importance of connecting a local pilot's success to broader policy frameworks (for example, feeding results into basin management plans or national strategy debates).

By mapping out who holds influence and who cares most deeply about water decisions, one can devise an implementation strategy for new technologies that is technically sound *and* governance-aware. Ultimately, successful water innovation is as much about effective stakeholder management as it is about engineering.

7.3. Policy Analysis Results: Stakeholder Perspectives

Similar to what was done for the engineering part of this research, the following sections describe what was found during the interview process, divided into key insights that match each coding group.

In this case, the main goal is to answer the research sub-questions three and four, even if Section 7.3.1 is still referring back to the first research sub-question, to strengthen the technical results. Research sub-questions three and four are reported as a reminder:

- *What organisational, economic, and infrastructural conditions are needed to implement AI control systems at a national scale in European countries?*
- *How do regulators, plant operators, and engineers perceive the risks, benefits, and acceptability of AI in wastewater operations?*

7.3.1. Technical Feasibility and Innovation

Across the fifteen interviews, most participants recognised that AI-based control for water and wastewater systems is *technically feasible*, but only under favourable conditions of data quality and digital infrastructure. Technical experts from utilities and academia (Interviews 01, 07, 08, 09, 10, 12) stressed that the underlying algorithms are mature; the true bottlenecks lie in:

- sensor reliability, quality, and price;
- data availability and integration;
- and infrastructure heterogeneity (”*each network is particular with its own particular materials, problems, malfunctions, and flaws*” from Interview 02)

For example, Interviewee 01 explained that ”*if we have good sensors and good data, it would be easy to apply AI control*”, referring to wastewater treatment control, while several researchers (Interviews 08 and 12) highlighted that modern datasets enable realistic prediction of energy consumption and treatment performance.

However, this technical optimism was tempered by strong realism among field engineers and policy-oriented interviewees. A recurrent observation was that the possibility of applying AI-based control (or other modern control strategies) depends on the physical readiness of the plant, and that most facilities, especially in Southern Europe (Interviews 07, 08), remain ”*obsolete, not designed for digital integration*”. The interviewee 07 specifically refers to the southern Italian situation as ”*dramatic,..with leaks up to 48–60%*” in the context of water networks, stressing the fact that ”*even important companies such as IREN still have 20% leakages*”, which is *unacceptable*. Two academic experts (Interviews 06 and 09) noted that, in their opinion, control algorithms cannot generalise across different stations due to heterogeneity. Smaller utilities were viewed as incapable of supporting the required computational and maintenance costs.

Perceived opportunities were consistent: ”*energy efficiency*”, ”*optimal control*”, ”*predictive maintenance*”, and ”*fault detection*” dominated. Yet even the most enthusiastic experts, such as the hydrology professor in Interview 06, warned that economic and institutional barriers outweigh technical ones: ”*Before thinking of intelligent control, we must first maintain the networks in good condition.*” Overall, approximately nine of fifteen participants viewed AI as technically sound in principle, but only three described it as currently operationally ready.

7.3.2. Trust and Explainability

All participants agreed that human-in-the-loop supervision is indispensable. Even those with high technical confidence rejected full autonomy in critical infrastructures. The prevailing expectation was that AI should "*assist, not replace*" operators (Interviews 01, 03, 05, 10, 11, 12). Concerns centred on explainability of not understanding the AI-control decisions, which could trigger unsafe actions, especially in safety-critical contexts such as chemical dosing or dam management.

Trust levels varied systematically. Engineers and data scientists (Interviews 01, 05, 09, 10, 11) displayed moderate-to-high trust but demanded traceability of decisions. Operators and junior engineers (Interviews 02, 06) expressed only conditional trust, fearing that AI could be "*not smart enough to understand certain problems*" (Interview 02). The cybersecurity specialist (Interview 04) offered the most cautious stance, warning that "*when you add AI, you just add another entry point for attackers.*" This reflects a recurrent duality: AI is trusted for performance, but not yet for critical control tasks. This suggests that further proof could be required to gain people's trust in AI-based control systems.

Explainability was widely framed as a social rather than purely technical property: interviewee 07 argued that black-box models undermine institutional confidence, while interviewee 10 warned that overtrust can produce complacency: "*people just click 'yes, yes' without reading.*" This highlights a key insight: trust erosion may occur both from lack of understanding and from excessive familiarity.

7.3.3. Organisational and Economic Factors

Every interviewee linked AI adoption to organisational readiness rather than algorithmic maturity. Skill shortages and fragmented governance were identified as the most persistent barriers. At least eight participants emphasised that existing staff lack digital competence or familiarity with data-driven tools (Interviews 01, 05, 06, 07, 08, 09, 10, 11, 15). Particularly, interviewee 15 underlined that in the Netherlands, training on AI is a policy that is currently being implemented in the 21 water boards of the country.

The interviews reveal a clear public–private divide. Large utilities and research centres possess in-house expertise, sensor networks, and financial resources, while small operators face prohibitive costs and lack specialised personnel (Interviews 02, 05, 06): "*we already have a data engineer at our company, and they can barely afford him. It's a small company*" (Interview 02). The consultant in Interview 05 summarised this asymmetry: "*The real gap is in competences, not technology.*" Public utilities often prioritise short-term maintenance and regulatory compliance over innovation, resulting in chronic underinvestment in automation infrastructure. Economic discussions also surfaced a conceptual tension between CAPEX (capital expenditures, such as investment in automation and sensors) and OPEX (operational expenditures). Practitioners argued that digital investments stabilise long-term energy costs, but political and tariff constraints discourage upfront spending. Several interviewees (05, 06, 07) insisted that without state-supported training and incentives, small operators risk permanent technological lag.

7.3.4. Regulatory and Institutional Barriers

Views on regulation diverged sharply. Roughly half of the participants (01, 05, 07, 09, 11) described the regulatory environment as restrictive, especially in the European context, where environmental rules and procurement procedures slow experimentation. On the AI Act, there are ambivalent perceptions. Technical actors viewed it as bureaucratic and risk-averse, whereas policy-oriented interviewees recognised its value for trust-building and safety assurance (Interviews 06, 08). Overall, regulation is seen as necessary for legitimacy but insufficiently informed by technical expertise. An engineer lamented that regulators "*put rocks on the road instead of helping optimisation*" (Interview 01), while interviewee 07 criticised "*laws written by people who don't understand the technology*". A similar concept also occurs in Interview 12, where it was mentioned that the EU AI Act "*lacks technical clarity - too many lawyers, not enough engineers*". Confirming these points, interviewee 15 avoided answering most of the technical-related questions, even on a high level, such as discussing main infrastructural requirements, showing uncertainty and a lack of knowledge on the topic, even if directly involved in its governance. These last three opinions (01, 07, 12) are explicitly calling for mixed expert-legislator committees to translate engineering realities into adaptive governance. These opinions contrast with the moderate optimism of non-EU respondents (03, 09, 11), who characterised the Canadian and Moroccan frameworks as supportive but still immature.

Procurement and institutional fragmentation were equally problematic. Interviewee 06 pointed to Italy's "*800 water operators versus eight in the UK*" as a structural impediment to coordinated innovation. Similarly, Interviewee 14 stressed the differences that appear between rural and urban water systems in France. Across contexts, the absence of common data standards, inter-agency coordination, and consistent funding mechanisms was described as the main non-technical bottleneck.

7.3.5. Transition and Future Outlook

Nearly all respondents anticipate gradual but inevitable AI integration over the next decade. The dominant expectation is a hybrid transition: AI will control routine operations under human oversight. Interview 01 envisioned "*fully functioning treatment sites with AI control, but still monitored by humans.*" Interviews 06 and 07 both emphasised that progress depends on infrastructure consolidation and regulatory flexibility.

Pilot projects were the most frequently cited transitional tool (Interviews 01, 05, 06, 07, 08), followed by regulatory sandboxes (Interviews 03, 11, 12). Participants generally rejected the idea of abrupt automation; instead, they favoured incremental scaling through tested demonstrators and shared learning platforms. Training and education were universally endorsed as essential enablers, both for operational staff and policymakers. Interview 05 mentioned the need for a "*Joint educational program between companies and institutions*" in the matter of AI.

Critically, the transition narratives reveal a paradox: while all actors recognise the inevitability of digital transformation, few foresee coherent governance to steer it. Without cross-sector coordination between engineers, regulators, and cybersecurity experts, AI risks amplify, rather than reduce.

7.3.6. AI Misconception: "AI = ChatGPT"

A striking insight from the interviews is the widespread conflation of AI with generative chatbots. It should be clear from reading the previous chapters that the goal of this work is not to let OpenAI control critical water stations through chatbots like ChatGPT.

Even if a detailed explanation of the research was made at the beginning of every interview, approximately one-third of interviewees (mostly non-specialists or junior engineers) initially equated "AI" with ChatGPT or text-based assistants. Several participants explicitly clarified that "*people think AI is ChatGPT, but that's not the same as control algorithms*" (Interview 02). As a matter of fact, interviewee 13 stated that "*AI is a buzzword*" and prefers to refer to algorithms such as Deep Reinforcement Learning as "*advanced statistical methods*".

The confusion has critical implications. First, it inflates expectations: some anticipate conversational or autonomous reasoning in contexts where AI control is, in reality, deterministic and data-constrained. Second, it distorts regulatory focus; legislators may overemphasise ethical or linguistic issues while neglecting control-system safety and data infrastructure. Finally, it reinforces cultural resistance among operators who associate AI with unpredictable or intrusive behaviour, given the fact that sometimes chatbots can give wrong answers.

Some technical experts (*Interviews 01, 07, 08, 09, 11, 12, 13*) consistently distinguished between deep reinforcement learning (RL) and generative AI, but not all of them, and this is concerning. This evidence underscores a profound knowledge gap between research-oriented actors and everyday practitioners, suggesting that awareness campaigns and education are as crucial as technological innovation.

7.3.7. Critical Reflection

Across the interviews, a consistent theme emerges: technological feasibility is outpacing institutional capacity. Experts largely agree that the science of AI-based optimisation, particularly in aeration, energy management, and predictive maintenance, is mature enough for deployment. Yet adoption is slowed by socio-technical factors: uneven literacy, regulatory rigidity, fragmented governance, and divergent perceptions of risk.

The findings indicate that successful transition will depend less on algorithmic performances (even if they are crucial to boost public trust) and more on the creation of shared cognitive infrastructure done through training programmes, cross-disciplinary coordination, and regulatory experimentation that connect technical feasibility with societal legitimacy.

8

From Potential to Adoption: Integrative Analysis and Policy Recommendations

8.1. Technology-Policy Gap

As was initially mentioned in the technical literature review in Section 2.1.2, and subsequently confirmed by the engineering results in Section 5.2 and the interview results in Section 7.3, in the previous Chapter, applying AI to water treatment plants (or similarly to other pieces of critical water infrastructure), is technically feasible, more optimal, and relatively safe.

However, as highlighted with the literature review conducted in Chapter 2, and specifically as was mentioned in Section 2.2, many WWTPs continue to have basic reactive control systems that fail to adapt efficiently to real-time inflow variability in their characteristics, pollution loads, and climatic conditions [101]. This gap between what modern AI-based controllers can actually do and how these are implemented in real world facilities could be given by a misalignment between technology and policies.

The table below Table 8.1 shows a comparison that links quantitative model outcomes (e.g., energy optimisation, process stability, cost reduction) with institutional and organisational constraints (e.g., trust, data governance, skills), synthesising the dual-track findings. From a technical standpoint, DRL-based control systems have demonstrated robustness, efficiency, and compliance potential under realistic disturbances. Yet, institutional inertia—manifested in fragmented governance, cautious regulation, and skill shortages, preventing the translation of such advancements into operational practice.

Table 8.1: Cross-analysis of technical and institutional readiness for AI-based wastewater optimisation.

Dimension	Technical Findings	Institutional Insights
Energy optimisation	DRL controllers showed up to 46% higher R_{tot} than the relay, maintaining effluent stability and lowering operating costs.	Regulatory frameworks promote energy neutrality (UWTD recast 2024) but risk aversion and rigidity pilot deployment.
Process stability and safety	No constraint violations observed; controllers demonstrated stable behaviour under variable inflow and tariff conditions.	Lack of certified validation protocols for AI-based control; operators remain accountable for any process deviation, discouraging automation.
Data and sensing	Controller performance and adaptability increase with sensor reliability and real-time data access.	Data fragmentation, ownership ambiguity, infrastructural limits (sensors), and privacy obligations (GDPR, Data Act) restrict model sharing across utilities.
Trust and transparency	RL behaviour interpretable through KPIs and reward functions; smooth control increases human confidence.	Operators and regulators require explainable dashboards and clear human-override mechanisms to ensure accountability.
Skills and capacity	Model implementation requires multidisciplinary expertise (process, data science, control).	Small and medium utilities lack competencies and financial capacity for AI implementation.
Innovation governance	DRL allows adaptive optimisation aligned with sustainability goals (energy, GHG reduction).	Regulatory mechanisms not yet tailored to adaptive control; no clear pathways exist.

This misalignment creates what can be defined as a *technology–policy gap*: while algorithms can already deliver measurable sustainability gains, governance frameworks are not yet adaptive enough to legitimise, certify, and scale them. To interpret this misalignment as part of a wider systemic transition, the following section applies the Transition Model Canvas (TMC) from [30] at a European level, building on the Multi-Level Perspective introduced in Chapter 6.

8.2. Adoption Readiness

As highlighted by the literature review in Chapter 2 and the technical results in Section 5.2, algorithmic maturity is not the limiting factor for the adoption of AI-based control systems. As a matter of fact, the first experiments on learning machines can be traced back to 1959 with the studies conducted by [153], even if, in recent years, the breakout of generative AI and large language models (LLMs) made machine learning a ”trend”. With this being said, the *technical readiness* (in the sense of ”algorithmic maturity”) will not be further discussed in this section, as it was widely proven in different ways and various points of this thesis.

Therefore, the causes of the missing adoption of AI-assisted control in WWTPs must be searched elsewhere. The previous section (Section 8.1) has already identified the technology-policy gap, but there is more. It is important to consider the entire sociotechnical system and its multiple dimensions. The goal of the next subsections is to assess the system's readiness for adoption by building on the results of Chapter 5 and Chapter 7.

8.2.1. Organisational and Human Readiness

Organisational readiness remains the most decisive factor for adoption. As emerged from the interviews, the diffusion of AI-based control requires a *digital competence base* that many utilities currently lack. While large operators such as SUEZ or Veolia already employ data engineers and process analysts, small and medium utilities are often struggling to keep up due to high costs required for AI R&D. Training, cross-departmental collaboration, and awareness of AI's operational logic are prerequisites for meaningful adoption. Interviewees (01, 05, 06, 07, 08, 09, 10, 11, 15) repeatedly stressed that even in technologically advanced contexts, the staff lacks competencies and understanding, revealing the need for targeted educational programmes that combine process and machine learning engineering, data literacy, and AI ethics.

Trust and explainability are also part of human readiness. The majority of experts endorsed hybrid configurations where “*AI assists but does not replace humans*”, underlying the relevance of the human-in-the-loop. Such configurations align with the EU AI Act's emphasis on human oversight for high-risk systems. Explainable dashboards, anomaly detection modules, and traceable logs are therefore not only design features but cultural enablers that foster acceptance and confidence among operators.

8.2.2. Regulatory and Institutional Readiness

Regulatory readiness is progressing but still fragmented. The new Urban Wastewater Treatment Directive (UWWTD recast 2024) establishes energy-neutrality targets and encourages digitalisation, yet it does not provide guidance on the certification or validation of adaptive control algorithms. Similarly, the EU AI Act [16] defines wastewater control systems as *high-risk*, mandating risk management, transparency, and post-deployment monitoring, but stops short of specifying sector-specific standards. Additionally, several interviewees (01, 07, 12) specifically called for mixed expert–legislator committees capable of translating operational realities into adaptive governance. Until such mechanisms are in place, utilities will continue to adopt a “wait-for-standard” strategy, delaying innovation until liability and compliance frameworks are clarified. Moreover, what was said in Section 8.1 remains valid also for assessing the regulatory and institutional readiness in this section,

8.2.3. Technical and Infrastructural Readiness

As mentioned before, the maturity of AI-based control algorithms has developed enough to consider their implementation in water treatment. However, it needs the support of digital and physical infrastructures. Modern DRL or MPC-based controllers rely on dense, reliable data streams. However, many plants, especially in Southern Europe (Interviews 07 and 08), or rural realities (Interview 02), still operate with outdated sensors, heterogeneous SCADA systems, and weak communication protocols. As emphasised by Interview 01 (“*if we have good sensors and good data, it would be easy to apply AI control*”), the bottleneck lies in instrumentation rather than computation. Adoption thus depends on upgrading the sensing infrastructure (for instance, DO, ORP, NH_4), implementing redundancy to guarantee fault tolerance, and ensuring interoperability between SCADA and external optimisation layers. Cybersecurity and privacy compliance (under the Data Act and GDPR) also shape infrastructural readiness, requiring secure data pipelines and clear data ownership models.

8.2.4. Overall Assessment

When combined, these insights reveal that European WWTPs are at an *intermediate readiness stage*:

- Technological readiness: high, proven through simulations and pilots (see Section 5.2);
- Organisational and human readiness: moderate, constrained by skill gaps and resistance to automation;
- Regulatory readiness: emerging, pending certification and liability mechanisms;
- Infrastructural readiness: uneven, depending on national and regional investment capacity.

8.3. European TMC for AI-based Wastewater Infrastructure

Following the identification of misalignments between technological feasibility and the systems' readiness highlighted in the previous sections, this section explores how innovation can be effectively and safely advanced. To do this, as introduced in Chapter 6, the analysis implements the analytical lenses of the *Multi-Level Perspective* (MLP) [114] and the *Transition Model Canvas* (TMC) [30], which are applied to the European wastewater sector. The filled-in TMC can be seen in Figure 8.1, below, while the next subsections are used to explain its main components in greater detail, with attention to the national variations, but still keeping a European point of view.

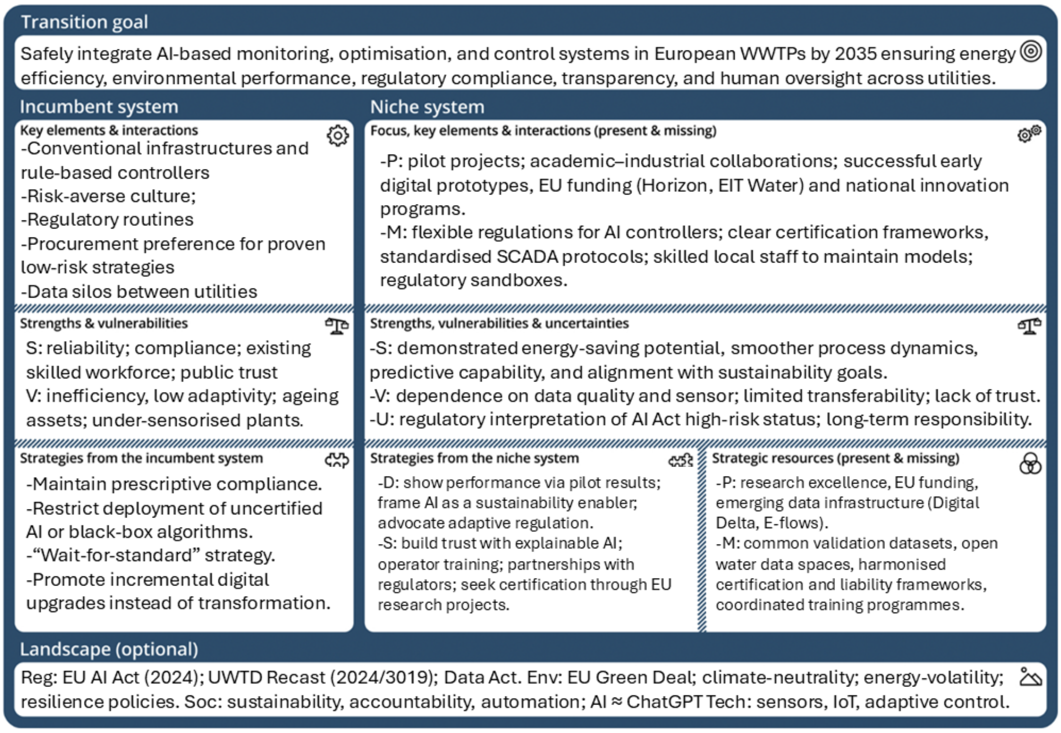

Figure 8.1: Integrated Transition Model Canvas for AI-based wastewater infrastructure in Europe.

8.3.1. Transition Goal

As confirmed by the results of Chapter 5 and Chapter 7, AI-based control systems are a promising novelty in water infrastructures when compared to traditional reactive control systems.

Therefore, an obvious *transition goal* - directly linked to the final research sub-question introduced in Chapter 3 - emerges. This is summarised at the top of Figure 8.1 with the following sentence: *safely integrate AI-based monitoring, optimisation, and control systems in European wastewater treatment plants (WWTPs) by 2035, ensuring energy efficiency, environmental performance, regulatory compliance, transparency, and human oversight*.

8.3.2. Incumbent System

In the terminology of [114], the *incumbent system* represents the stabilised socio-technical regime that currently structures, in this case, wastewater treatment across Europe. It is characterised by conventional infrastructures, rule-based control logic, and a strongly risk-averse culture. Most treatment plants rely on Supervisory Control and Data Acquisition (SCADA) and Programmable Logic Controller (PLC) architectures, reflecting decades of incremental technological evolution that has prioritised compliance and reliability over digital adaptability [101, 70]. These systems are rooted in engineering reliability rather than digital transformation [3].

Organisationally, as also shown by the interview results of Chapter 7, utilities tend to operate within risk-averse cultures, where reliability, effluent quality, and energy cost containment outweigh experimentation with algorithmic control [11].

This conservatism is reinforced institutionally: regulatory bodies consolidate regime stability through binding, prescriptive frameworks. An example is the Urban Wastewater Treatment Directive (UWWTD) [13], which defines deterministic performance thresholds, sampling procedures, and mandatory penalties for non-compliance. The recast Directive explicitly mandates that *"urban wastewater treatment plants shall meet the requirements in Table 1 of this Annex"* and that *"Member States shall lay down the rules on penalties applicable to infringements … ensuring they are effective, proportionate and dissuasive"* [13]. Additional provisions on energy audits and nutrient removal [79] exemplify how EU regulation maintains a command-and-control paradigm that privileges compliance and predictability over adaptive experimentation, thereby constraining technological flexibility [35, 100]. As a result, feedback loops between technology and regulation consolidate a "compliance-first" mindset across utilities, privileging reliability and control over innovation.

From a TMC perspective, this incumbent configuration is maintained by a self-reinforcing alignment of three pillars:

- Actors: utilities, engineering firms, regulators, and large equipment manufacturers that stabilise operational norms and procurement cycles;
- Institutions: rule-based permitting systems and public procurement frameworks that favour proven, low-risk technologies;
- Infrastructures: mature mechanical–electrical systems (aeration blowers, sensors, SCADA networks) optimised for stability rather than flexibility or data-driven adaptation.

This alignment provides efficiency and reliability but also generates *path dependency*, limiting the sector's absorptive capacity for disruptive innovation [154, 155]. The wastewater domain thus constitutes a highly regulated yet digitally immature regime, an archetypal *niche-resistant system* in transition-management terms, where innovation trajectories are bounded by the interplay of regulatory prescription, infrastructural lock-in, and institutional inertia.

8.3.3. Niche System

As shown in the literature review in Chapter 2, several studies conducted in academic, industrial, and experimental environments report promising results for AI-based control and monitoring systems in water critical infrastructures. These include machine-learning soft sensors [72, 73], digital twins, and deep RL controllers [5, 74, 6], which demonstrate the capacity to dynamically optimise aeration and nutrient removal while maintaining process stability and regulatory compliance [4, 5]. Yet, consistent with the findings in Chapter 2, most of these applications remain at the pre-operational stage, requiring further experimentation on real facilities before they can meaningfully challenge the incumbent system [89, 75].

From a transition-theory perspective, these developments constitute an emerging *niche system*. Following the Multi-Level Perspective (MLP) [114], niches are protected spaces—such as research labs, pilot sites, regulatory sandboxes, or digital-simulation environments—where radical innovations can develop without immediately conforming to incumbent constraints. In the European wastewater sector, this niche is fuelled by academic–industrial collaborations and EU-funded programmes (Horizon Europe, EIT Water, LIFE), which provide temporary protection and resources for experimentation. As shown in Section 5.2, niche demonstrators increasingly provide evidence of smoother aeration dynamics, improved effluent quality, and reduced operational variability.

Using the Transition Model Canvas (TMC), the niche configuration can be described through three interrelated components:

- Actors: research groups, technical universities, start-ups, start-up incubators (e.g. YES!DELFT), digital-solution providers, and innovation-oriented utilities forming collaborative consortia;
- Institutions: emerging governance norms on explainability, human oversight, and ethical AI, influenced by broader frameworks such as the EU AI Act [16], which, although originating at the landscape level, are progressively translated by niche actors into operational design principles;
- Infrastructures: experimental digital twins, IoT sensor networks, and data platforms that enable training, validation, and benchmarking of AI-driven control solutions.

Niche actors pursue dual strategies of *strengthening the niche*-through EU-funded R&D projects, demonstrators, and inter-utility knowledge sharing [98, 99]-and *regime disruption*, by framing AI as a pathway to energy-neutrality, sustainability, and improved compliance. At the same time, they face structural constraints: dependence on incumbent infrastructures, the absence of clear certification pathways, limited interoperability, and persistent skill gaps within utilities [87, 156].

Overall, the niche is technologically advanced but institutionally fragile: a promising foundation for transition that still awaits stronger systemic alignment, clearer standards, and broader real-world validation.

8.3.4. Landscape Drivers

The *landscape level*, as defined by [114], encompasses broad exogenous trends that exert pressure on established regimes while creating opportunities for transition. In the European wastewater sector, these pressures are particularly pronounced due to the concurrent green and digital transformations shaping infrastructure governance.

Four major categories of landscape drivers can be identified:

- Regulatory drivers: the EU AI Act [16] and the recast Urban Wastewater Treatment Directive (Directive 2024/3019) [13] jointly redefine compliance boundaries for the sector. The AI Act classifies AI control systems in critical infrastructure, such as WWTPs, as *high-risk*, enforcing documentation, transparency, traceability, and human-oversight requirements, which might weaken the niches. At the same time, the UWWTD introduces performance-based targets for energy neutrality and nutrient removal. Together, these frameworks reshape the risk calculus of utilities and create incentives for responsible digitalisation. Complementary instruments such as the European Green Deal, the Data Act, and national recovery plans further reinforce a multi-layered push toward digitalisation, resilience, and carbon neutrality.
- Environmental drivers: climate neutrality and circular economy commitments under the European Green Deal [14] amplify the demand for energy-efficient and resource-recovering operations. Aeration alone may account for a large part of a wastewater treatment plant's total electricity use [2], positioning AI-based optimisation as a key enabler of environmental and economic sustainability.
- Societal drivers: public expectations for sustainability, transparency, and accountability are increasing, reinforcing the importance of explainability in AI deployment [91, 18]. As shown from the interview results Section 7.3, opinions on trust in AI remain quite divided, posing the social landscape in a sort of ”unstable equilibrium” position.
- Technological drivers: at the same time, advances in sensor reliability, machine learning, and data interoperability are accelerating technological readiness and enabling real-time optimisation capabilities previously unattainable within conventional frameworks.

These interacting forces can influence both the incumbent and the niche system by either reinforcing one or the other. A good alignment of the landscape drivers might also create a *window of opportunity* for niche innovations to scale beyond the pilot stage and reconfigure the regime toward adaptive, AI-enabled control [114].

8.3.5. National Variations

Italy The Italian water sector remains highly fragmented, governed by regional authorities (EGATO, ARERA, ARPA). While digitalisation is uneven, EU Recovery and Resilience Facility (RRF, which is applied in Italy as PNRR, which stands for "*Piano Nazionale di Ripresa e Resilienza*") funding and ARERA's innovation incentives provide leverage for pilot adoption. Barriers include heterogeneous SCADA systems, low interoperability, and limited AI literacy among operators. The transition thus depends on targeted regional modernisation and skills development.

France France features a centralised governance model with strong incumbents (SUEZ, Veolia) and national coordination through the Ministry of Ecological Transition. Procurement rigidity and certification barriers slow innovation, but national AI strategies, innovation clusters, and pilot programmes by the Agences de l'Eau or OFB offer promising entry points. Here, integration efforts focus on explainability, liability, and standardised validation frameworks.

Netherlands The Netherlands stands out for its decentralised yet highly competent water boards, collaborative culture, and openness to digital innovation. Initiatives like Digital Delta and the Water & Energy Factory already embed AI and data analytics in operations. Remaining challenges include uneven AI expertise and liability clarity, but trust, interoperability, and environmental accountability make the Netherlands a frontrunner in the transition.

8.4. Barriers and Leverage Points

As summarised in Table 8.2, the diffusion of AI-based control in wastewater systems is constrained by intertwined technical, organisational, regulatory, and economic barriers. Yet, each dimension also reveals actionable leverage points that can inform transition pathways, used in the next section to develop an actionable policy for the safe implementation of AI-based control systems.

Table 8.2: Summary of systemic barriers and leverage points for AI-based wastewater control adoption in Europe.

Dimension	Main Barriers	Leverage Points / Enablers
Technical & Infrastructural	Outdated and limited sensors (e.g., mostly missing for NH_4); rule-based control systems; data fragmentation; cybersecurity and privacy constraints (GDPR, Data Act).	Incremental sensor and infrastructure upgrades; deployment of soft sensors and digital twins; secure data pipelines; integration with national and EU data spaces.
Organisational & Human	Skill shortages and low digital literacy; resistance to automation; misconceptions about AI (e.g., "AI = ChatGPT, as seen in Section 7.3"); lack of interdisciplinary collaboration; lack of trust in AI-based systems; cultural resistance to innovation; safety concerns.	Training and reskilling programmes; hybrid human–AI control schemes; cross-disciplinary communities of practice; educational campaigns to improve AI understanding and trust; pilot projects to prove safety and feasibility.
Regulatory & Institutional	Procurement rigidity; liability uncertainty; lack of certification standards for adaptive control; fragmented governance across EU and regional levels.	Regulatory sandboxes under AI Act supervision; harmonised and flexible certification and validation schemes; mixed expert–legislator committees.
Economic & Market	High upfront development costs (CAPEX) for digital upgrades; limited budgets in small utilities; weak financial incentives for innovation.	Public–private partnerships; innovation grants and tariff-based incentives (e.g., ARERA schemes); circular-economy and energy-efficiency funds; shared-cost demonstration pilots.

8.5. Policy and Design Recommendations for Safe Transition

Building upon the gaps, readiness assessment, and transition analysis discussed in the previous sections, this final part of the chapter translates the findings into concrete policy and design recommendations. Whereas Section 8.1 highlighted the persistent misalignment between technical feasibility and institutional adoption, and Section 8.2 assessed the system's preparedness across organisational, regulatory, and infrastructural dimensions, this section focuses on *how* to bridge that divide.

The recommendations derive from the integrated perspective offered by the TMC in Section 8.1, which positioned AI-based control as a niche able to grasp a *window of opportunity* in the European regulatory and sustainability landscape but constrained by incumbent routines. They are therefore intended to help policymakers, regulators, and utilities create an enabling environment for the responsible, explainable, and scalable deployment of AI in wastewater treatment plants (WWTPs).

It is also important to mention that this research keeps a European-level focus; therefore, the following recommendations are high-level, providing just minor specifications for each state.

Addressing the systemic lock-in. As summarised in Section 8.4, the European wastewater regime remains trapped in a self-reinforcing loop: conservative procurement and cultural resistance limit digital upgrades, reducing data quality and algorithm reliability, turning into a higher regulatory caution. Breaking this lock-in requires coordinated interventions along three complementary axes: governance, human capacities, and infrastructure.

Leverage Points. Still referring back to Section 8.4, the integrative analysis identifies several key levers that can be deployed to weaken the regime and strengthen the niches, in order to accelerate the transition towards the implementation of AI-based control systems, without compromising safety or accountability.

Similarly to what was done for Section 8.2, the proposed recommendations, which are derived from the results, are divided into the same three groups, to build on the current status identified.

8.5.1. Organisational & Human Factors Recommendations

The organisational and human dimension represents one of the most decisive factors for the safe adoption of AI-based control systems. As shown in the readiness analysis (Section 8.2), the key challenge lies in equipping people and organisations with the necessary knowledge, skills, confidence, and routines to work effectively with intelligent automation.

A first step is to introduce continuous training and reskilling programmes tailored to the wastewater sector. Several interviewees (01, 05, 06, 07, 08, 09, 10, 11, 15) pointed out that many plant operators and even some regulators lack familiarity with how AI-based control systems function, interpret data, or make optimisation decisions. Training modules should therefore combine technical content, such as AI interpretability, anomaly detection, and reinforcement learning basics, with operational case studies illustrating how digital control can enhance, rather than replace, human decision-making. Establishing permanent training tracks, possibly coordinated by national water associations or EU initiatives, would help professionalise the workforce and reduce resistance to change.

A complementary measure concerns the deployment of hybrid human–AI control protocols, in which operators remain at the centre of supervision while AI systems assist in optimisation. This "human-in-the-loop" configuration was consistently emphasised during interviews as a condition for acceptance: operators trust systems that they can override and understand. Transparent dashboards, interpretable KPIs, and clear traceability logs should therefore be embedded in the design of AI-based controllers. Beyond technical transparency, this approach also has a symbolic dimension—preserving human agency and accountability, in line with the EU AI Act's provisions for high-risk systems.

Equally important is the need to counter the widespread misconception that 'AI = ChatGPT'. As emerged repeatedly from interviews (e.g., 03, 07, 09, 10), this confusion persists even among technical personnel and contributes to unwarranted scepticism about control algorithms. Targeted education and communication campaigns can correct these misunderstandings by highlighting the differences between generative models and control-oriented machine learning. Workshops, site demonstrations, and visual communication materials can demystify AI and emphasise its concrete, safety-oriented applications in wastewater optimisation.

Finally, communities of practice across utilities, academia, and regulators would institutionalise learning from early adopters. These networks, modelled on existing European initiatives, could facilitate cross-pilot knowledge exchange, discuss shared challenges (e.g., sensor reliability, algorithm certification) and develop common guidelines. These communities would reduce the fragmentation that is currently holding technical innovation in a separate sphere from institutional reform by creating continuous interaction among engineers, data scientists, and policy makers.

In summary, strengthening organisational and human capabilities is a precondition for any technological transition in the water sector. Structured training, human-centred design, accurate communication, and collaborative learning platforms form the social infrastructure that will enable AI-based control systems to move from technical potential to trusted daily practice.

8.5.2. Policy & Governance Recommendations

To adapt the policy and the governance system to support AI innovation, the two most relevant actions to take first are to establish better mixed expert-legislator committees and provide technical education to high-level strategic decision makers, so that the proposed policies can have more practical relevance and that the system is understood by the governance on a deeper level. This is a clear suggestion that emerged from interview results (particularly Interviews 01, 07, 12, 15).

Another policy suggestion, this time coming from OECD, is the one to utilise regulatory sandboxes and pilot projects to guide AI innovation [93]. These are controlled experimentation environments that allow supervised trials of AI controllers under regulator oversight and relaxed compliance constraints for defined periods and plants. This approach can also help in the definition of future certification standards. As a matter of fact, data collected under controlled conditions can inform future certification schemes, as mentioned by [93]: "*The link between AI regulatory sandboxes and AI certification is clear: testing an AI product in a regulatory sandbox can help an agency decide to certify its quality or safety*". This pathway builds confidence, generates reference cases, and aligns with the EU AI Act's risk-management requirements.

Another critical policy lever is the development of data-sharing frameworks and interoperability standards. This would facilitate benchmarking and model validation, lower the cost of developing AI-based controllers, and contribute to the emergence of common performance metrics. Such shared repositories would also ease collaboration with academic partners, allowing universities to train and validate models on representative datasets while respecting privacy and security constraints. In this context, the European Commission has already encouraged the creation of data spaces for key sectors through the Data Act and the Digital Europe Programme [16, 157].

Finally, aggregation of small utilities represents a context-specific recommendation, especially relevant for Italy. Interviewees highlighted that small operators often lack the financial and human resources to invest in advanced control or data infrastructure individually. Policy instruments promoting joint digitalisation programmes, inter-utility consortia, or shared service platforms could allow these entities to reach critical mass, benefit from economies of scale, and access European innovation funding (e.g., Horizon Europe, LIFE, or Digital Europe calls). In this sense, aggregation is not merely an administrative reform but a digital enabler, creating a more coherent governance landscape capable of sustaining AI-based operations.

In summary, these governance-oriented recommendations converge on a common objective: to make Europe's regulatory environment both *protective and enabling*. By combining interdisciplinary governance, evidence-based experimentation, harmonised certification, data interoperability, and structural aggregation, policymakers can transform the current cautious stance toward AI into a structured pathway for safe and transparent innovation.

8.5.3. Physical & Digital Infrastructure Recommendations

The infrastructural foundation of European wastewater treatment plants remains one of the most critical enablers of AI adoption. As discussed in Section 8.2, many utilities—particularly small and medium ones—still operate with heterogeneous Supervisory Control and Data Acquisition (SCADA) systems, limited sensor coverage, and weak communication protocols. These limitations restrict both the quality and quantity of data available for model training and real-time optimisation, thereby undermining the reliability of AI-based control. Consequently, modernising the digital and physical infrastructure constitutes a strategic priority for safe transition.

The most immediate intervention is to fund incremental upgrades of sensor networks and infrastructure. This measure was repeatedly highlighted during the interviews (especially 01, 05, 08) as the most pragmatic and cost-effective entry point for innovation. Even partial improvements can unlock significant process visibility and enable hybrid AI–human optimisation. The process follows a reinforcing loop: "investments lead to more sensors which allow for more data, resulting in better AI models for monitoring and control (softsensors). This ultimately improves optimisation and lowers costs, which can result in greater organisational buy-in.

This virtuous cycle illustrates how digital infrastructure and organisational readiness co-evolve. Early investment in sensing capabilities is thus not only a technical improvement but also a confidence-building mechanism that accelerates cultural adoption.

In parallel, the sector should invest in secure data pipelines to support real-time distributed control while ensuring compliance with cybersecurity and privacy regulations. Local computing capacity allows inference and optimisation to occur directly at the plant level, reducing latency and dependency on cloud infrastructure. At the same time, secure communication protocols and encryption standards aligned with the EU Cybersecurity Act and the Data Act are necessary to protect sensitive operational data and ensure trust among regulators, utilities, and technology providers.

Finally, it is crucial to align digital investments with the broader objectives of the European Green Deal and the recast Urban Wastewater Treatment Directive (Directive 2024/3019). These policy frameworks jointly promote energy neutrality, greenhouse-gas reduction, and circular-resource management. AI-based control systems can directly contribute to these goals by optimising processes (e.g. aeration), reducing electricity demand, and improving nutrient recovery. Funding programmes under the Green Deal and Digital Europe Programme should therefore explicitly recognise wastewater digitalisation as a contributor to climate and resource-efficiency targets.

In conclusion, it can be said that it is only by reinforcing these infrastructural foundations can AI-based optimisation can move from models to a resilient and scalable component of Europe's critical water infrastructure.

8.6. Indicative Roadmap and Transition Pathways

Building on the preceding analysis, a roadmap is now proposed to guide the application of the previously identified recommendations. The transition unfolds through three phases, applicable to the short, medium, and long term.

8.6.1. Phase 1 (0-12 months)

In this first phase, five main actions are required:

1. Incentivise the creation of mixed expert-legislator committees to support the activity of regulatory bodies and knowledge diffusion, as shown in the interview results in Chapter 7, the lack of information and culture on AI is the most relevant problem identified, also in regulatory bodies.
2. Launch trainings and educational programmes on how AI actually works and how it should be used, to fill the skill gap and knowledge problems, given the results of Chapter 7, which highlighted how the skill gap is often the first barrier to adoption.
3. Constitute regulatory sandboxes (also this suggestion is grounded in the interview results of Chapter 7), especially to enable supervised experimentation of AI controllers under regulator oversight. National authorities, such as ARERA, OFB, or Rijkswaterstaat, could authorise pilot deployments in selected WWTPs, temporarily relaxing compliance requirements. This will generate operational evidence for certification and risk-management frameworks [16, 93].
4. Use the sandboxes to launch pilot projects under regulatory supervision.
5. Financially incentivise the acquisition of modern IoT sensors and overall modernisation of the infrastructure.

These are fundamental as they create a solid base for the next phases.

8.6.2. Phase 2 (12-24 months)

As lessons from pilot projects consolidate, utilities can progressively integrate AI optimisation modules alongside conventional rule-based control. Operators, that already gone through the training phase can now retain oversight on AI decisions through explainable dashboards, anomaly detection, and manual-override functions, ensuring accountability and greater user trust.

This can be viewed as a sort of ”acceleration phase”. Public-private partnerships can now provide the organisational and financial capacity to scale these initiatives across plants and regions. Additionally, fragmented realities (e.g., Italy) should be pushed towards aggregation in this phase, as mentioned in Interview 06.

It is also useful to define unified and detailed European standards and methods for algorithm evaluation and benchmarking to ensure maximum trust in the last phase, while keeping up the technical progress at the same time. As mentioned in Interview 14, ”*the role of politics (in this context) is to reassure and explain* and ”*Europe must keep its sovereignty (on AI) and should not take an imposed American or Chinese vision*”.

8.6.3. Phase 3 (24 months-onward)

At this point, the sociotechnical system should not require further intervention besides long-distance monitoring. The full-scale implementation should now happen naturally and complete the transition to AI-based control systems, which will, in the meantime, continue their improvements thanks to the increased data availability.

9

Conclusions, Reflections and Limitations

9.1. Conclusions and Answers to the Research Questions

This thesis was set out to explore how AI can enhance the technical efficiency, economic viability, and institutional acceptability of wastewater treatment processes. The overarching research question asked:

How can AI-based solutions be technically effective, economically viable, safe, and acceptable for sustainable water management in Europe?

And it was answered through the five following sub-questions answers:

1. *Given the state of the art of AI, is it technically feasible to control wastewater treatment processes with AI solutions in real facilities?*
   The research confirmed the technical feasibility of AI-based control through the development and benchmarking of DRL controllers trained on real operational data from a SUEZ wastewater treatment plant. The custom digital-twin environment reproduced realistic actuator dynamics and constraints, allowing safe pre-training before future deployment. Results showed that data-driven controllers can successfully optimise aeration dynamics and respond adaptively to inflow disturbances, proving readiness for pre-operational validation under industrial conditions. In addition, the interview results also showed agreement on technical feasibility, further proving this point.
2. *How effective are the proposed AI solutions in improving energy efficiency and pollutants control in WWTPs?*
   Quantitative results demonstrated that DRL controllers can achieve double-objective optimisation: reducing energy use while maintaining effluent quality. Compared to baseline and rule-based control strategies, the trained agents achieved significant reductions in aeration energy demand without compromising ORP and DO set-points. This confirms that AI can contribute meaningfully to both economic and environmental performance goals of WWTPs.
3. *What organisational, economic, and infrastructural conditions are needed to implement AI control systems at a national scale in European countries?*
   The analysis of France, Italy, and the Netherlands revealed that successful adoption depends on three readiness pillars: (i) data and sensor infrastructure, (ii) organisational capacity and operator training, and (iii) financial and regulatory enablers. Centralised systems exhibit stronger industrial readiness but slower procurement and certification pathways, while fragmented contexts require first regional modernisation and interoperability programmes. The Netherlands, with its integrated governance model, stands out as a pioneer of adaptive, innovation-driven water management.
4. *How do regulators, plant operators, and engineers perceive the risks, benefits, and acceptability of AI in wastewater operations?*
   Interviews across multiple stakeholder categories highlighted a cautious optimism. AI is perceived as a promising tool to improve process stability and sustainability, yet concerns persist about safety, transparency, and accountability in high-risk infrastructures. Stakeholders emphasised the importance of human oversight, operator explainability, and certification frameworks. Interestingly, many interviewees noted that the integration of engineering and policy perspectives, as developed in this thesis, is essential to building trust and ensuring responsible deployment.
5. *How is AI-based control positioned within the wastewater sector, and which strategies and policy instruments would most effectively shift the system toward large-scale adoption?*
   Using the Transition Model Canvas, the study positioned AI-based control as a niche innovation beginning to challenge the incumbent paradigm of rule-based operation. Landscape pressures, such as the EU AI Act and the recast Urban Wastewater Treatment Directive, are opening institutional windows for change. The proposed transition pathway, outlined in Chapter 8, suggests progressive alignment between technical pilots, certification sandboxes, and human skills. These measures collectively chart a path from experimentation to responsible adoption.

To conclude, this thesis demonstrated, by combining a DRL-based engineering case study with a multi-level policy analysis, that AI-based control of wastewater treatment is technically viable and strategically aligned with the sustainability and digitalisation objectives of the European water sector. However, for a large-scale implementation, given the sociotechnical challenges, carefully planned policies and transition pathways are required.

9.2. Reflections on Engineering & Policy Integration

The thesis encompasses the vision of the *Engineering and Policy Analysis* (EPA) programme: integrating analytical rigour from engineering with strategic reasoning from policy sciences. It brings together deep reinforcement learning for process optimisation with transition frameworks and stakeholder interviews to show that complex sustainability challenges cannot be solved from one disciplinary perspective alone.

From an engineering perspective, this research moved the frontier on the use of AI-based control in real wastewater operations, demonstrating how digital twins and data-driven controllers can yield benefits in terms of efficiency and compliance. The policy analysis showed that regulatory frameworks, institutional capacity, and actor perceptions will determine whether such innovations can be trusted and adopted at scale.

This interdisciplinary synthesis represents precisely what the graduates from EPA are trained for: to understand both the mechanics of the system and the dynamics of its governance. Several interviewees explicitly noted the value of this dual competence, underlining the fact that few people can ”speak both engineering and policy,” which is an essential feature in bridging innovation and regulation. In that respect, this thesis is not only an academic contribution but also a demonstration of the EPA mindset: connecting models to missions, algorithms to accountability, and optimisation to legitimacy.

9.3. Limitations

While this work provided significant insights, several limitations must be acknowledged.

Models and architectures tested. Given the fact that one of the main, high-level goals of this research is to bridge the gap between technical experts and policy makers (which coincides with the scope of the EPA programme) and not to excel in both parts, the evaluation of algorithms was limited. Only SAC and TD3 were tested and compared in two different configurations. Hyperparameter tuning and tests with different reward functions are not addressed in this work and were not conducted systematically.

Limited training and deployment scope. The developed controllers were pre-trained on historical and simulated data from a real station, ensuring safety and reproducibility, but they have not been implemented online yet.

Data and interpretability constraints. Despite using real datasets, access to plant-level metadata and long-term operational histories is constrained by non-disclosure agreements (NDAs), potentially impacting the transparency of the research.

Limited interview sample and socio-economic evaluation. Time constraints made it impossible to conduct more interviews; therefore, the results are limited to the 15 interviewee sample described in Chapter 6. This is a typical constraint of qualitative research methods, where depth of insight is prioritised over breadth of statistical representation. This means that the economic and social implications of large-scale AI adoption were only qualitatively addressed. This also results in the fact that certain opinions on lower-level stakeholders (e.g. operators) were indirectly collected from interviews with higher-level actors.

Limited geographic coverage. The policy analysis focused on three European countries: Italy, France, and the Netherlands. Even though these offer diverse backgrounds, this work cannot be taken as an exhaustive representation of the European water sector.

Evolving regulatory context. The AI Act (2024) and the recast Urban Wastewater Treatment Directive (2024/3019) are still in their implementation phase. Their practical interpretation, especially concerning conformity assessment, liability, and certification, may significantly alter the institutional dynamics described in this research.

9.4. Future Research

Building on the results and limitations discussed above, several directions emerge for future research on AI-based control in wastewater treatment and, more broadly, in water-related critical infrastructures.

First, future work should prioritise closed-loop testing under real plant conditions. Ideally, this would be conducted within *European regulatory sandboxes* specifically designed for critical infrastructures, so that the practical limitations, safety issues, and organisational frictions of AI-based control systems can be assessed without jeopardising compliance. In parallel, controllers should be equipped with explainable AI components (e.g. post-hoc explanation methods and policy visualisations) and be trained and evaluated on open or at least shareable datasets, to improve transparency, replicability, and cross-utility learning.

On the engineering side, future work should deepen and broaden the controller design:

- Reward function engineering: explore multi-objective and risk-sensitive reward formulations that explicitly capture trade-offs between energy use, effluent quality, and operational robustness, and better reflect operator preferences and regulatory constraints.
- Hyperparameter and architecture optimisation: replace manual tuning with systematic strategies (e.g. Bayesian optimisation or population-based training) to improve sample efficiency and robustness of DRL agents.
- Multi-phase, multi-model training: design training pipelines where agents are first trained on simplified surrogate models, then progressively refined on higher-fidelity digital twins and soft sensors, before any (safe) online adaptation.
- ODE-based MPC on a rolling horizon: use the mechanistic ODE model developed in this thesis, combined with parameter identification from plant data, as the prediction model in a nonlinear Model Predictive Control (MPC) scheme operating on a rolling horizon. This would enable a systematic comparison between DRL, MPC, and more classical controllers (relay, PID, ABAC), and possibly hybrid architectures where RL and MPC interact.

Beyond pure control performance, future research should quantify the broader impacts of AI-based controllers through integrated techno-economic and societal assessments. This includes linking process-level KPIs (energy use, aeration volume, effluent quality, GHG proxies) to: (i) economic indicators (operational expenditures, avoided fines, investment needs), and (ii) societal metrics (reliability, service quality, perceived risk). Such an assessment would support more informed cost-benefit discussions among regulators, utilities, and technology providers.

From a governance and policy perspective, the comparative framework developed in this thesis could be expanded in several ways. Geographically, including Eastern and Northern European countries would enhance the generalisability of findings and reveal new institutional configurations. Sectorally, extending interviews and case studies beyond wastewater (e.g. to dams, drinking-water networks, desalination plants, and flood protection systems) would clarify to what extent the identified barriers and leverage points are *sector-specific* or *cross-cutting*. In doing so, future work should also collect more direct stakeholder perspectives from underrepresented groups and sectors.

Finally, further research is needed to elaborate concrete policy instruments and institutional designs that can operationalise safe AI adoption in critical water infrastructures. This includes detailed proposals for regulatory sandboxes, certification and conformity-assessment pathways for adaptive controllers, cross-utility and cross-border data spaces for training and validating AI models, and capability-building programmes that equip operators, managers, and regulators with the skills needed for effective human oversight.

By combining these engineering, assessment, and governance-oriented directions, the next generation of studies can move beyond demonstrating technical potential towards designing and implementing fully operational, trustworthy, and sustainable AI governance for Europe's critical water infrastructures.

References

1. European Parliament and Council of the European Union. Directive (EU) 2022/2557 on the resilience of critical entities. Official Journal of the European Union, L 333, 27 December 2022. 2022. Available from: https://eur-lex.europa.eu/eli/dir/2022/2557/oj [Accessed on: 2025 Aug 27]
2. Drewnowski J, Remiszewska-Skwarek A, Duda S, and Łagód G. Aeration Process in Bioreactors as the Main Energy Consumer in a Wastewater Treatment Plant. Review of Solutions and Methods of Process Optimization. Processes 2019; 7:311. DOI: 10.3390/pr7050311
3. Jamaludin M, Tsai YC, Lin HT, Huang CY, Choi W, Chen JG, and Sean WY. Modeling and Control Strategies for Energy Management in a Wastewater Center: A Review on Aeration. Energies 2024; 17:3162. DOI: 10.3390/en17133162
4. Gu Y, Li Y, Yuan F, and Yang Q. Optimization and Control Strategies of Aeration in Wastewater Treatment Plants: A Review. Journal of Cleaner Production 2023; 418:138008. DOI: 10.1016/j.jclepro.2023.138008
5. Croll HC, Ikuma K, Ong SK, and Sarkar S. Systematic Performance Evaluation of Reinforcement Learning Algorithms Applied to Wastewater Treatment Control Optimization. Environmental Science & Technology 2023; 57:18382–90. DOI: 10.1021/acs.est.3c00353
6. Chen K, Wang H, Valverde-Pérez B, Zhai S, Vezzaro L, and Wang A. Optimal control towards sustainable wastewater treatment plants based on multi-agent reinforcement learning. Chemosphere 2021; 279:130498. DOI: 10.1016/j.chemosphere.2021.130498
7. Zhu Z, Dong S, Zhang H, Parker W, Yin R, Bai X, Yu Z, Wang J, Gao Y, and Ren H. Bayesian optimization-enhanced reinforcement learning for self-adaptive and multi-objective control of wastewater treatment. Bioresource Technology 2025; 421:132210. DOI: 10.1016/j.biortech.2025.132210
8. United Nations. Sustainable Development Goals. 2015. Available from: https://sdgs.un.org/goals [Accessed on: 2025 May 13]
9. Obaideen K, Shehata N, Sayed ET, Abdelkareem MA, Mahmoud MS, and Olabi AG. The role of wastewater treatment in achieving sustainable development goals (SDGs) and sustainability guideline. Energy Nexus 2022; 7:100112. DOI: 10.1016/j.nexus.2022.100112
10. Roothans N, Pabst M, Diemen M van, Herrera Mexicano C, Zandvoort M, Abeel T, Loosdrecht MCM van, and Laureni M. Long-term multi-meta-omics resolves the ecophysiological controls of seasonal N2O emissions during wastewater treatment. Nature Water 2025. DOI: 10.1038/s44221-025-00430-x
11. Cairone S, Hasan SW, Choo KH, Lekkas DF, Fortunato L, Zorpas AA, Korshin G, Zarra T, Belgiorno V, and Naddeo V. Revolutionizing wastewater treatment toward circular economy and carbon neutrality goals: pioneering sustainable and efficient solutions for automation and advanced process control with smart and cutting-edge technologies. Journal of Water Process Engineering 2024; 63:105486. DOI: 10.1016/j.jwpe.2024.105486
12. SUEZ. Improve environmental performance through digital solutions. 2025. Available from: https://www.suez.com/en/group/about-us/digital-solutions [Accessed on: 2025 May 13]
13. European Parliament and Council of the European Union. Directive (EU) 2024/3019 of the European Parliament and of the Council of 13 November 2024 on urban wastewater treatment and repealing Council Directive 91/271/EEC. Official Journal of the European Union, L 2024/3019, 12 December 2024. 2024. Available from: https://eur-lex.europa.eu/eli/dir/2024/3019/oj [Accessed on: 2025 Aug 27]

14. European Commission. Communication from the Commission: A Water Resilience Initiative for the Union. COM(2025) 280 final. 2025. Available from: https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:52025DC0280 [Accessed on: 2025 Aug 27]
15. OECD. Water Governance in the Netherlands: Fit for the Future? Tech. rep. Organisation for Economic Co-operation and Development (OECD), 2014. DOI: 10.1787/9789264102637-en
16. European Parliament and Council of the European Union. Regulation (EU) 2024/1689 of the European Parliament and of the Council of 12 July 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). Official Journal of the European Union, L 1689, 12 July 2024. 2024. Available from: https://eur-lex.europa.eu/eli/reg/2024/1689/oj [Accessed on: 2025 Aug 27]
17. Jørgensen BN and Ma ZG. Impact of EU Regulations on AI Adoption in Smart City Solutions: A Review of Regulatory Barriers, Technological Challenges, and Societal Benefits. Information 2025; 16. DOI: 10.3390/info16070568
18. Schmude T, Yurrita M, Alfrink K, Goff TL, Tschiatschek S, and Viard T. Two Means to an End Goal: Connecting Explainability and Contestability in the Regulation of Public Sector AI. 2025. arXiv: 2504.18236 [cs.CY]. Available from: https://arxiv.org/abs/2504.18236
19. Anglmayer I. European Critical Infrastructure: Revision of Directive 2008/114/EC. Tech. rep. PE 662.604. Implementation Appraisal. European Parliamentary Research Service (EPRS), 2021. Available from: https://www.europarl.europa.eu/thinktank/en/document/EPRS_BRI(2021)662604 [Accessed on: 2025 Sep 26]
20. WWAP (UNESCO World Water Assessment Programme). The United Nations World Water Development Report 2017: Wastewater: The Untapped Resource. Paris, France: UNESCO, 2017. Available from: https://unesdoc.unesco.org/ark:/48223/pf0000247153?posInSet=1&queryId=719a7c03-6247-4548-801c-218d5ff430ab [Accessed on: 2025 May 13]
21. United Nations Human Settlements Programme (UN-Habitat) and World Health Organization (WHO). Progress on the proportion of domestic and industrial wastewater flows safely treated – Mid-term status of SDG Indicator 6.3.1 and acceleration needs, with a special focus on climate change, wastewater reuse and health. Licence: CC BY-NC-SA 3.0 IGO. Geneva, Switzerland: United Nations Human Settlements Programme (UN-Habitat) and World Health Organization (WHO), 2024. Available from: https://www.unwater.org/sites/default/files/2024-08/SDG6_Indicator_Report_631_Progress-on-Wastewater-Treatment_2024_EN_0.pdf [Accessed on: 2025 Sep 2]
22. Corcoran E, Nellemann C, Baker E, Bos R, Osborn D, and Savelli H, eds. Sick Water? The central role of wastewater management in sustainable development: A Rapid Response Assessment. Arendal, Norway: United Nations Environment Programme (UNEP), UN-HABITAT, GRID-Arendal, 2010. Available from: https://digitallibrary.un.org/record/760768
23. Ritchie H, Rosado P, and Roser M. Breakdown of carbon dioxide, methane and nitrous oxide emissions by sector. Our World in Data 2020. Available from: https://ourworldindata.org/emissions-by-sector [Accessed on: 2025 Jul 4]
24. Liu R, Qiao Y, Zhang H, Lv M, and Hao X. Contribution of solar photovoltaic to the decarbonization of wastewater treatment plants in China. Desalination 2025; 611:118935. DOI: https://doi.org/10.1016/j.desal.2025.118935
25. Heia JT, Isah ME, Masuda S, Matsubae K, Chen R, Li YY, and Sano D. Quantifying health risks from wastewater reclamation and agricultural reuse using life cycle assessment and quantitative microbial risk assessment. Science of The Total Environment 2025; 977:179363. DOI: https://doi.org/10.1016/j.scitotenv.2025.179363
26. Aparna K, Swarnalatha R, and Changmai M. Optimizing wastewater treatment plant operational efficiency through integrating machine learning predictive models and advanced control strategies. Process Safety and Environmental Protection 2024; 188:995–1008. DOI: https://doi.org/10.1016/j.psep.2024.05.148

27. Zhang Y, Zhang Q, Xu X, Wang W, Wang A, Lee DJ, Ren N, and Chen C. High-Efficiency mixotrophic biological nitrogen removal process at low Temperature: Sulfur Disproportionation-Mediated greenhouse gas mitigation. Chemical Engineering Journal 2025; 514:163028. DOI: https://doi.org/10.1016/j.cej.2025.163028
28. Eggen RIL, Hollender J, Joss A, Schärer M, and Stamm C. Reducing the Discharge of Micropollutants in the Aquatic Environment: The Benefits of Upgrading Wastewater Treatment Plants. Environmental Science & Technology 2014; 48:7683–9. DOI: 10.1021/es500907n
29. Preisner M, Neverova-Dziopak E, and Kowalewski Z. Analysis of eutrophication potential of municipal wastewater. Water Science and Technology 2020; 81:1994–2003. DOI: 10.2166/wst.2020.254. Available from: https://doi.org/10.2166/wst.2020.254
30. Rijnsoever FJ van and Leendertse J. A practical tool for analyzing socio-technical transitions. Environmental Innovation and Societal Transitions 2020; 37:225–37. DOI: 10.1016/j.eist.2020.08.004
31. Prabu SL, Suriyaprakash TNK, and Kumar JA. Wastewater treatment technologies: A review. Pharma Times 2011; 43:9–14. Available from: https://www.researchgate.net/publication/296188508_Wastewater_treatment_technologies_A_review [Accessed on: 2025 Aug 23]
32. Wei Y, Van Houten RT, Borger AR, Eikelboom DH, and Fan Y. Minimization of excess sludge production for biological wastewater treatment. Water Research 2003; 37:4453–67. DOI: 10.1016/S0043-1354(03)00441-X
33. Watt S, Nelson M, Hai F, and Sidhu H. Modelling of the Activated Sludge process with a stratified settling unit. *Proceedings of the 25th International Congress on Modelling and Simulation (MODSIM2023)*. Darwin, Australia: Modelling, Simulation Society of Australia, and New Zealand, 2023 :1044–9. Available from: https://mssanz.org.au/modsim2023
34. Akpor O, Otohinoyi DA, Olaolu TD, and Aderiye JB. Pollutants in wastewater effluents: Impacts and remediation processes. International Journal of Environmental Research and Public Health 2014. Available from: https://www.researchgate.net/publication/261834688
35. (EEB) EEB and NGOs allied. Joint NGO analysis of the UWWTD / Recast proposal. Report / policy analysis. 2022. Available from: https://eeb.org/wp-content/uploads/2022/12/Joint-NGO-analysis-UWWTD.pdf
36. Ardern E and Lockett WT. Experiments on the oxidation of sewage without the aid of filters. Journal of the Society of Chemical Industry 1914; 33:523–39. DOI: 10.1002/jctb.5000331005
37. Shah MP and Rodriguez-Couto S, eds. Wastewater Treatment Reactors: Microbial Community Structure. ISBN: 978-0-12-823991-9. Amsterdam, Netherlands: Elsevier, 2021. Available from: https://www-sciencedirect-com.tudelft.idm.oclc.org/book/edited-volume/9780128239919/wastewater-treatment-reactors [Accessed on: 2025 Sep 5]
38. Wagner M and Loy A. Microbial community composition and function in wastewater treatment plants. Current Opinion in Biotechnology 2002; 13:218–27. DOI: 10.1016/S0958-1669(02)00315-4
39. Schmidt I, Sliekers O, Schmid M, Cirpus I, Strous M, Bock E, Kuenen J, and Jetten MS. Aerobic and anaerobic ammonia oxidizing bacteria: competitors or natural partners? FEMS Microbiology Ecology 2002 Mar; 39:175–81. DOI: 10.1111/j.1574-6941.2002.tb00920.x
40. Odize V. Diffuser Fouling Mitigation, Wastewater Characteristics And Treatment Technology Impact on Aeration Efficiency. PhD thesis. Virginia Tech, 2018. Available from: https://vtechworks.lib.vt.edu/server/api/core/bitstreams/a2435245-86dd-4329-8e9f-9cf9087044f0/content [Accessed on: 2025 Aug 26]
41. Azis NF, Ramli N, and Abdul Hamid MF. Energy efficiency of wastewater treatment plant through aeration system. Desalination and Water Treatment 2019. DOI: 10.5004/DWT.2019.24134
42. Pöyry L, Mikola A, and Blomberg M. Modelling wastewater aeration system and high-speed turbo blower for energy efficiency improvements. Technical Report 2020. Available from: https://www-semanticscholar-org.tudelft.idm.oclc.org/paper/Modelling-wastewater-aeration-system-and-high-speed-P%C3%B6yry-Mikola/c8623d01f67ad468ffa2d54c343daef0df82d30e [Accessed on: 2025 Sep 10]

43. Herrmann-Heber R, Reinecke S, and Hampel U. Dynamic aeration for improved oxygen mass transfer in the wastewater treatment process. Chemical Engineering Journal 2020; 386:122068. DOI: 10.1016/j.cej.2019.122068
44. Lozano Avilés AB, Velázquez FdC, and Pascual del Riquelme ML. Methodology for Energy Optimization in Wastewater Treatment Plants. Phase III: Implementation of an Integral Control System for the Aeration Stage. Sensors 2020; 20:4342. DOI: 10.3390/s20154342
45. Wei Y, Jiao Y, Wang Q, Du Y, and Wei Q. Review of Dissolved Oxygen Detection Technology: From Laboratory Analysis to Online Intelligent Detection. Sensors 2019; 19:3995. Available from: https://www.mdpi.com/1424-8220/19/18/3995 [Accessed on: 2025 Nov 15]
46. McDonagh C, Kolle C, McEvoy AK, Dowling DL, et al. Phase fluorometric dissolved oxygen sensor. Sensors and Actuators B: Chemical 2001; 74:124–30. DOI: 10.1016/S0925-4005(00)00721-8
47. Zhao Y et al. Ratiometric Optical Fiber Dissolved Oxygen Sensor Based on Fluorescence Quenching Principle. Sensors 2022; 22:4811. Available from: https://www.mdpi.com/1424-8220/22/13/4811 [Accessed on: 2025 Nov 15]
48. Hamilton Company. Optical Dissolved Oxygen Sensors & Principles of Operation. 2023. Available from: https://www.hamiltoncompany.com/knowledge-base/article/optical-dissolved-oxygen-sensors-principles-of-operation [Accessed on: 2025 Nov 15]
49. Wang X et al. Characterization of Oxidation–Reduction Potential (ORP) Sensor Behavior for Water Quality Monitoring. Processes 2022; 10:2607. Available from: https://www.mdpi.com/2227-9717/10/12/2607 [Accessed on: 2025 Nov 15]
50. Weißbach M et al. Application of the oxidation–reduction potential (ORP) for process control in biological wastewater treatment. Chemical Engineering Journal 2018; 343:452–61. DOI: 10.1016/j.cej.2018.02.001
51. YSI/Xylem. Oxidation–Reduction Potential: An Often Misunderstood Parameter. 2023. Available from: https://www.ysi.com/ysi-blog/water-blogged-blog/2023/02/oxidation-reduction-potential-an-often-misunderstood-parameter [Accessed on: 2025 Nov 15]
52. Cuartero M, Colozza N, Fernández-Pérez BM, and Crespo GA. Why ammonium detection is particularly challenging but insightful with ionophore-based potentiometric sensors – an overview of the progress in the last 20 years. The Analyst 2020 Jan; 145:3188–210. DOI: 10.1039/d0an00327a
53. Papias S et al. In situ continuous monitoring of nitrogen with ion-selective electrodes in a constructed wetland receiving treated wastewater: an operating protocol to obtain reliable data. Water Science and Technology 2018; 77:1706–13. DOI: 10.2166/wst.2018.052
54. Amazon Italia. YOGAYE Portable Dissolved Oxygen Meter. Available from: https://www.amazon.it/dp/B0CDKFN8HL [Accessed on: 2025 Sep 16]
55. Wei Y, Jiao Y, An D, Li D, Li W, and Wei Q. Review of Dissolved Oxygen Detection Technology: From Laboratory Analysis to Online Intelligent Detection. Sensors 2019; 19:3995. DOI: 10.3390/s19183995
56. L-Com. Industrial Dissolved Oxygen Sensor (4–20 mA / RS485, SRWQ100-DO104). Available from: https://www.l-com.com/dissolved-oxygen-sensor-4-20ma-rs485-m39x1.5-5m-cable-srwq100-do103-8901 [Accessed on: 2025 Sep 16]
57. Amazon Italia. Generic ORP Pen Tester. Available from: https://www.amazon.it/dp/B0C14PV3VK [Accessed on: 2025 Sep 16]
58. Zainurin SN, Wan Ismail WZ, Mahamud SNI, Ismail I, Jamaludin J, Ariffin KNZ, and Wan Ahmad Kamil WM. Advancements in Monitoring Water Quality Based on Various Sensing Methods: A Systematic Review. International Journal of Environmental Research and Public Health 2022 Oct; 19:14080. DOI: 10.3390/ijerph192114080
59. Omega Engineering. TX-ORP ORP Transmitter. Available from: https://it.omega.com/ [Accessed on: 2025 Sep 16]

60. Vernier Science Education. Ammonium Ion-Selective Electrode (S16214ND). Available from: https://www.fishersci.com/shop/products/vernier-ammonium-ion-selective-electrode/S16214ND [Accessed on: 2025 Sep 16]

61. Hach. Amtax sc Ammonium Analyser Outdoor 0.05–20 mg/L NH4-N, 2-Ch LXV421.99.13002. Internet. 2023 Sep. Available from: https://www.ebay.com/itm/334819073767 [Accessed on: 2025 Sep 3]

62. SUEZ. Our Group. https://www.suez.com/en/group. Accessed: 2025-05-13. 2025

63. Zatarain Salazar J, Kwakkel JH, and Witvliet M. Evaluating the choice of radial basis functions in multiobjective optimal control applications. Environmental Modelling & Software 2024; 171:105889. DOI: 10.1016/j.envsoft.2023.105889

64. Osika Z, Zatarain Salazar J, Oliehoek FA, and Murukannaiah PK. Navigating Trade-offs: Policy Summarization for Multi-Objective Reinforcement Learning. *ECAI 2024: 27th European Conference on Artificial Intelligence*. Vol. 392. Frontiers in Artificial Intelligence and Applications. IOS Press, 2024 :2919–26. DOI: 10.3233/FAIA240830

65. Osika Z, Rădulescu R, Zatarain Salazar J, Oliehoek FA, and Murukannaiah PK. Multi-Objective Reinforcement Learning for Water Management. arXiv preprint arXiv:2505.01094 2025. Available from: https://arxiv.org/abs/2505.01094

66. Tian W, Liao Z, Zhi G, Zhang Z, and Wang X. Combined Sewer Overflow and Flooding Mitigation Through a Reliable Real-Time Control Based on Multi-Reinforcement Learning and Model Predictive Control. Water Resources Research 2022; 58:e2021WR030703. DOI: 10.1029/2021WR030703

67. Zhang Z, Tian W, and Liao Z. Towards Coordinated and Robust Real-Time Control: A Decentralized Approach for Combined Sewer Overflow and Urban Flooding Reduction Based on Multi-Agent Reinforcement Learning. Water Research 2023; 229:119498. DOI: 10.1016/j.watres.2022.119498

68. Hajgató G, Paál G, and Gyires-Tóth B. Deep Reinforcement Learning for Real-Time Optimization of Pumps in Water Distribution Systems. Journal of Water Resources Planning and Management 2020; 146:04020079. DOI: 10.1061/(ASCE)WR.1943-5452.0001287

69. Hu S, Gao J, Zhong D, Wu R, and Liu L. Real-Time Scheduling of Pumps in Water Distribution Systems Based on Exploration-Enhanced Deep Reinforcement Learning. Systems 2023; 11:56. DOI: 10.3390/systems11020056

70. Åmand L, Olsson G, and Carlsson B. Aeration control - A review. Water science and technology : a journal of the International Association on Water Pollution Research 2013 Jun; 67:2374–98. DOI: 10.2166/wst.2013.139

71. Aponte-Rengifo O, Vilanova R, Francisco M, Vega P, and Revollar S. Optimal set point generation based on deep reinforcement learning and transfer knowledge for Wastewater Treatment Plants. *2023 IEEE 28th International Conference on Emerging Technologies and Factory Automation (ETFA)*. 2023 :1–7. DOI: 10.1109/ETFA54631.2023.10275522

72. Ching PM, So RH, and Morck T. Advances in soft sensors for wastewater treatment plants: A systematic review. Journal of Water Process Engineering 2021; 44:102367. DOI: https://doi.org/10.1016/j.jwpe.2021.102367

73. Schneider MY, Furrer V, Sprenger E, Carbajal JP, Villez K, and Maurer M. Benchmarking Soft Sensors for Remote Monitoring of On-Site Wastewater Treatment Plants. Environmental Science & Technology 2020 Sep; 54:10840–9. DOI: 10.1021/acs.est.9b07760

74. Klawikowska Z and Grochowski M. Optimizing Control of Wastewater Treatment Plant With Reinforcement Learning: Technical Evaluation of Twin-Delayed Deep Deterministic Policy Gradient Agent. IEEE Access 2024; 12:130318–33. DOI: 10.1109/access.2024.3458186

75. Kåge L, Milić V, Andersson M, and Wallén M. Reinforcement learning applications in water resource management: a systematic literature review. Frontiers in Water 2025; 7:1537868. DOI: 10.3389/frwa.2025.1537868

76. European Parliament and Council of the European Union. Regulation (EU) 2022/868 on European Data Governance (Data Governance Act). Official Journal of the European Union, L 152, 3 June 2022. Establishes mechanisms for the re-use of certain public sector data, data intermediaries, and data altruism frameworks. 2022. Available from: https://eur-lex.europa.eu/eli/reg/2022/868/oj [Accessed on: 2025 Sep 26]
77. European Parliament and Council of the European Union. Regulation (EU) 2023/2854 on harmonised rules on fair access to and use of data (Data Act). Official Journal of the European Union, L 2023/2854, 22 December 2023. Defines rights for users to access and share data generated by connected devices and services, including safeguards for B2G and B2B data sharing. 2023. Available from: https://eur-lex.europa.eu/eli/reg/2023/2854/oj [Accessed on: 2025 Sep 26]
78. European Parliament and Council of the European Union. Directive 2000/60/EC of the European Parliament and of the Council of 23 October 2000 establishing a framework for Community action in the field of water policy. Official Journal of the European Communities, L 327, 22 December 2000. Known as the Water Framework Directive; introduces the river basin management paradigm and the objective of good ecological status. 2000. Available from: https://eur-lex.europa.eu/eli/dir/2000/60/oj [Accessed on: 2025 Aug 27]
79. Capodaglio AG. Energy Audits and Energy Efficiency of Urban Wastewater Systems, Following UWWTP Directive 2024/3019. Water 2025; 17. DOI: 10.3390/w17142049
80. White House. America's AI Action Plan: Remove Red Tape and Onerous Regulation. PDF, Office of Science and Technology Policy. See section "Remove Red Tape and Onerous Regulation". 2025. Available from: https://www.whitehouse.gov/wp-content/uploads/2025/07/Americas-AI-Action-Plan.pdf [Accessed on: 2025 Sep 26]
81. Sheehan M. China's AI Regulations and How They Get Made. Carnegie Endowment for International Peace. 2023. Available from: https://carnegieendowment.org/research/2023/07/chinas-ai-regulations-and-how-they-get-made?lang=en [Accessed on: 2025 Sep 26]
82. CSET / Georgetown University. Translation: Artificial Intelligence Law of the People's Republic of China (Draft). Translation draft PDF. 2024. Available from: https://cset.georgetown.edu/wp-content/uploads/t0592_china_ai_law_draft_EN.pdf [Accessed on: 2025 Sep 26]
83. Pahl-Wostl C, Lebel L, Knieper C, and Nikitina E. From applying panaceas to mastering complexity: Toward adaptive water governance in river basins. Environmental Science & Policy 2012; 23:24–34. DOI: https://doi.org/10.1016/j.envsci.2012.07.014
84. Gupta J and Pahl-Wostl C. Global Water Governance in the Context of Global and Multilevel Governance: Its Need, Form, and Challenges. Ecology and Society 2013; 18:53. DOI: 10.5751/ES-05952-180453
85. Özerol Gül Vinke-de Kruijf J, Brisbois MC, Casiano Flores C, Deekshit P, Girard C, Knieper C, Mirnezami SJ, Ortega-Reig M, Ranjan P, Schröder NJS, and Schröter B. Comparative Studies of Water Governance: A Systematic Review. Ecology and Society 2018 Jan; 23. DOI: 10.5751/es-10548-230443
86. García MM, Hileman J, and Bodin Ö. Collaboration and conflict in complex water governance systems across a development gradient: addressing common challenges and solutions. Ecology and Society 2019; 24. Available from: https://www.jstor.org/stable/26796968 [Accessed on: 2025 Oct 19]
87. McMillan L and Varga L. A review of the use of artificial intelligence methods in infrastructure systems. Engineering Applications of Artificial Intelligence 2022; 116:105472. DOI: 10.1016/j.engappai.2022.105472
88. Kim BJ, Jeong S, Cho BK, and Chung JB. AI Governance in the Context of the EU AI Act: A Bibliometric and Literature Review Approach. arXiv preprint arXiv:2502.03468 2025. DOI: 10.48550/arXiv.2502.03468
89. Croll HC, Ikuma K, Ong SK, and Sarkar S. Reinforcement learning applied to wastewater treatment process control optimization: Approaches, challenges, and path forward. Critical Reviews in Environmental Science and Technology 2023; 53:1775–94. DOI: 10.1080/10643389.2023.2183699

90. Bedué P and Fritzsche A. Can We Trust AI? An Empirical Investigation of Trust Requirements and Guide to Successful AI Adoption. Journal of Enterprise Information Management 2021; 35. DOI: 10.1108/JEIM-06-2020-0233
91. Zhou T and Lu H. The effect of trust on user adoption of AI-generated content. The Electronic Library 2024 Dec; 43:61–76. DOI: 10.1108/el-08-2024-0244
92. O'Reilly CA and Tushman ML. Organizational Ambidexterity: Past, Present, and Future. The Academy of Management Perspectives 2013; 27:324–38. DOI: 10.5465/amp.2013.0025
93. OECD. Regulatory Sandboxes in Artificial Intelligence. OECD Digital Economy Papers, No. 356. 2023. DOI: 10.1787/8f80a0e6-en. Available from: https://www.oecd.org/en/publications/regulatory-sandboxes-in-artificial-intelligence_8f80a0e6-en.html
94. Beaudouin V, Bloch I, Bounie D, et al. Flexible and Context-Specific AI Explainability: A Multi-Disciplinary Approach. arXiv preprint 2020. eprint: 2003.07703. Available from: https://arxiv.org/abs/2003.07703
95. Yordanova F. Regulating AI: Challenges and the Way Forward Through Regulatory Sandboxes. *Artificial Intelligence Law and Regulation*. Springer, 2023 :379–98. DOI: 10.1007/978-3-031-41264-6_23
96. Ahern M. Operationalising AI Regulatory Sandboxes under the EU AI Act: Capacity, Coordination and Attractiveness. arXiv preprint 2025. eprint: 2509.05985. Available from: https://arxiv.org/abs/2509.05985
97. Queensland Government. Using an Outcome-Based Procurement Approach. Guidance Note. 2019. Available from: https://www.hpw.qld.gov.au/__data/assets/pdf_file/0018/3456/usingoutcomebasedprocurementapproach.pdf
98. European Commission. Water Research and Innovation Projects 2014–2020. EU Publications Office. 2022. Available from: https://research-and-innovation.ec.europa.eu/system/files/2022-03/water_research_and_innovation_projects_2014-2020.pdf
99. European Institute of Innovation and Technology. Call for EIT Water: Transforming Europe's Water, Marine and Maritime Sectors. 2025. Available from: https://www.eit.europa.eu/news-events/news/call-eit-water-transforming-europes-water-marine-and-maritime-sectors-and
100. Consortium I. Advancing Water Governance: Reference Guide. 2025. Available from: https://www.innwater.eu/sites/default/files/delivrables/InnWater_D2.2_Reference_Guide_Programming%20VF%20PU.pdf
101. Balla KM, Bendtsen JD, Schou C, Kallesøe CS, and Ocampo-Martinez C. A learning-based approach towards the data-driven predictive control of combined wastewater networks – An experimental study. Water Research 2022; 221:118782. DOI: https://doi.org/10.1016/j.watres.2022.118782
102. Sutton RS and Barto AG. Reinforcement learning: An introduction. 2nd ed. MIT press, 2018. Available from: https://www.andrew.cmu.edu/course/10-703/textbook/BartoSutton.pdf [Accessed on: 2025 Aug 10]
103. Goodfellow I, Bengio Y, and Courville A. Deep learning. MIT press, 2016. Available from: http://www.deeplearningbook.org [Accessed on: 2025 Aug 10]
104. Haarnoja T, Zhou A, Abbeel P, and Levine S. Soft Actor-Critic: Off-Policy Maximum Entropy Deep Reinforcement Learning with a Stochastic Actor. *Proceedings of the 35th International Conference on Machine Learning*. Ed. by Dy J and Krause A. Vol. 80. Proceedings of Machine Learning Research. PMLR, 2018 Jul :1861–70. Available from: https://proceedings.mlr.press/v80/haarnoja18b.html
105. Haarnoja T, Zhou A, Hartikainen K, Tucker G, Ha S, Tan J, Kumar V, Zhu H, Gupta A, Abbeel P, and Levine S. Soft Actor-Critic Algorithms and Applications. 2019. arXiv: 1812.05905 [cs.LG]. Available from: https://arxiv.org/abs/1812.05905
106. Raffin A, Hill A, Gleave A, Kanervisto A, Ernestus M, and Dormann N. Stable-Baselines3: Reliable Reinforcement Learning Implementations. Journal of Machine Learning Research 2021; 22:1–8. Available from: http://jmlr.org/papers/v22/20-1364.html

107. Fujimoto S, Hoof H van, and Meger D. Addressing Function Approximation Error in Actor-Critic Methods. *Proceedings of the 35th International Conference on Machine Learning*. Ed. by Dy J and Krause A. Vol. 80. Proceedings of Machine Learning Research. PMLR, 2018 Jul :1587–96. Available from: https://proceedings.mlr.press/v80/fujimoto18a.html [Accessed on: 2025 Aug 14]

108. Lillicrap TP, Hunt JJ, Pritzel A, Heess N, Erez T, Tassa Y, Silver D, and Wierstra D. Continuous control with deep reinforcement learning. arXiv preprint arXiv:1509.02971 2019. Available from: https://arxiv.org/abs/1509.02971 [Accessed on: 2025 Aug 15]

109. S.p.A. R. Water treatment. https://www.rovatti.com/market-area/water-utilities/1135/water-treatment. Accessed: 6 November 2025. 2025

110. Chang CN, Cheng HB, and Chao AC. Applying the Nernst Equation to Simulate Redox Potential Variations for Biological Nitrification and Denitrification Processes. Environmental Science & Technology 2004; 38:1807–12. DOI: 10.1021/es021088e

111. Venturini M, Rossen A, Bucci P, and Silva Paulo P. Applying the Nernst Equation to Control ORP in Denitrification Process for Uranium-Containing Nuclear Effluent with High Loads of Nitrogen and COD. Water 2022; 14:2227. DOI: 10.3390/w14142227

112. Henze M, Gujer W, Mino T, and Loosedrecht M van. Activated Sludge Models ASM1, ASM2, ASM2d and ASM3. Vol. 5. IWA Publishing, 2015 Dec. DOI: 10.2166/9781780402369

113. Fan H, Qi L, Liu G, Zhang Y, Fan Q, and Wang H. Aeration optimization through operation at low dissolved oxygen concentrations: Evaluation of oxygen mass transfer dynamics in different activated sludge systems. Journal of Environmental Sciences 2017; 55:224–35. DOI: 10.1016/j.jes.2016.08.008

114. Geels FW. Technological transitions as evolutionary reconfiguration processes: a multi-level perspective and a case-study. Research Policy 2002; 31:1257–74. DOI: 10.1016/S0048-7333(02)00062-8

115. Robinson OC. Probing in qualitative research interviews: Theory and practice. Qualitative Research in Psychology 2023; 20:1–17. DOI: 10.1080/14780887.2023.2238625

116. Anthony RN. Planning and Control Systems: A Framework for Analysis. Accessed via TU Delft Library. Boston: Division of Research, Graduate School of Business Administration, Harvard University, 1965

117. Rasmussen J. Risk management in a dynamic society: a modelling problem. Safety Science 1997; 27:183–213. DOI: 10.1016/S0925-7535(97)00052-0

118. Kallio H, Pietilä AM, Johnson M, and Kangasniemi M. Systematic methodological review: Developing a framework for a qualitative semistructured interview guide. Journal of Advanced Nursing 2016; 72:2954–65. DOI: 10.1111/jan.13031

119. Braun V and Clarke V. Using thematic analysis in psychology. Qualitative Research in Psychology 2006; 3:77–101. DOI: 10.1191/1478088706qp063oa

120. Mendelow AL. Stakeholder Mapping. *Proceedings of the 2nd International Conference on Information Systems*. Cambridge, MA, USA, 1991 :113–25. Available from: https://search.worldcat.org/title/10001959 [Accessed on: 2025 Sep 21]

121. Brugha R and Varvasovszky Z. Stakeholder analysis: a review. Health Policy & Planning 2000; 15:239–46. DOI: 10.1093/heapol/15.3.239

122. Citroni G and Lippi A. The permanent (de-)institutionalisation of multi-level governance of water services in Italy. *Water Law, Policy and Economics in Italy: Between National Autonomy and EU Law Constraints*. Cham: Springer International Publishing, 2021 :289–307. DOI: 10.1007/978-3-030-69075-5_12

123. ISTAT. Water Services in Italy – Progress and Challenges (World Water Day 2025 Report). Tech. rep. in Italian. Italian National Institute of Statistics, 2025. Available from: https://www.istat.it/wp-content/uploads/2025/05/WWD2025.pdf/ [Accessed on: 2025 Oct 26]

124. Balzarolo D, Lazzara P, Colonna P, Becciu G, and Rana G. The Implementation of the Water Framework Directive in Italy. Options Méditerranéennes. A 2011; 98:29–38. Available from: https://crabgrass.riseup.net/assets/112611/dce-it.pdf [Accessed on: 2025 Oct 24]

125. Domorenok E. Traps of Multi-Level Governance: Lessons from the Implementation of the Water Framework Directive in Italy. Journal of European Integration 2017; 39:657–71. DOI: 10.1080/07036337.2017.1322076

126. ARERA. Annual Report on the State of Water Services 2023. https://www.arera.it/eng/. Italian Regulatory Authority for Energy, Networks and Environment. 2023

127. Acea S.p.A. La nostra storia – Gruppo Acea. https://www.gruppoacea.it/nostra-storia. Accessed: 2025-11-04. 2023

128. European Commission. The Commission decides to refer Italy to the Court of Justice of the European Union for failure to fully comply with the Urban Waste Water Treatment Directive. Press release, European Commission Press Corner. Infringements package – key decisions, Case INFR(2017)2181. 2024 Mar. Available from: https://ec.europa.eu/commission/presscorner/detail/en/inf_24_663

129. Italy in the dock over wastewater failures. Water News Europe. 2023. Available from: https://www.waternewseurope.com/italy-in-the-dock-over-wastewater-failures/

130. European Commission takes Italy to court over lack of urban wastewater treatment. The Water Diplomat. 2023. Available from: https://www.waterdiplomat.org/story/2023/06/european-commission-takes-italy-court-over-lack-urban-wastewater-treatment

131. Dipartimento della Protezione Civile. Civil Protection in Italy. 2022. Available from: https://www.protezionecivile.gov.it/static/f72430e2c68a27106bfa228b63010aa5/volume-pc-eng-14-09-2022-low_0.pdf [Accessed on: 2025 Oct 24]

132. European Investment Bank. Italy: EIB provides €435 million to help Acea improve water service quality and resilience. 2023. Available from: https://www.eib.org/en/press/all/2023-259-italy-eib-provides-eur435-million-to-help-acea-improve-water-service-quality-and-resilience [Accessed on: 2025 Oct 24]

133. European Investment Bank. Italy: First InvestEU project in Umbria — Valle Umbra Servizi receives €35 million EIB loan to make water service more resilient and efficient. 2024. Available from: https://www.eib.org/en/press/all/2024-173-first-investeu-project-in-umbria-valle-umbra-servizi-receives-35-million-eib-loan-to-make-water-service-more-resilient-and-efficient [Accessed on: 2025 Oct 24]

134. European Investment Bank. EIB, with SACE and InvestEU guarantees, provides €150 million for Piedmont water services (Acqua Novara.VCO). 2025. Available from: https://www.eib.org/en/press/all/2025-229-eib-with-sace-and-investeu-guarantees-provides-eur150-million-for-piedmont-water-services [Accessed on: 2025 Oct 24]

135. ISTAT. Municipal and Regional Water Service Indicators 2025. Istituto Nazionale di Statistica, 2025. Available from: https://www.istat.it [Accessed on: 2025 Oct 24]

136. Water4All Partnership. Water-Oriented Living Labs: Good Practices from Europe. https://www.water4all-partnership.eu/. Case examples from EU-funded research pilots. 2023

137. Bettin A. Digital transformation for the water industry: how a data-driven business intelligence platform can improve operations. Water Practice and Technology. 2023; 18:1599–607. DOI: 10.2166/wpt.2023.091

138. ISTAT. Environmental and Resource Indicators 2025: Water Use and Public Perception. Istituto Nazionale di Statistica, 2025. Available from: https://www.istat.it [Accessed on: 2025 Oct 25]

139. ISPRA – Istituto Superiore per la Protezione e la Ricerca Ambientale. Rapporto sullo Stato dell'Ambiente 2023: Risorse Idriche e Qualità delle Acque. ISPRA, 2023. Available from: https://www.isprambiente.gov.it/it/pubblicazioni/rapporti [Accessed on: 2025 Oct 25]

140. Organisation for Economic Co-operation and Development. Regional, rural and urban development. OECD, 2023. Available from: https://www.oecd.org/regional/water-governance-italy.htm [Accessed on: 2025 Oct 24]

141. Organisation for Economic Co-operation and Development. Managing Water Risks for Resilient Economies: Country Review Italy. OECD, 2023. Available from: https://www.oecd.org/environment/water/ [Accessed on: 2025 Oct 25]

142. European Commission. Special Eurobarometer 538: Attitudes of European Citizens Towards the Environment. Directorate-General for Communication, 2024. Available from: https://europa.eu/eurobarometer [Accessed on: 2025 Oct 25]

143. Colon M, Richard S, and Roche PA. The evolution of water governance in France from the 1960s: disputes as major drivers for radical changes within a consensual framework. Water International 2018; 43:1–24. DOI: 10.1080/02508060.2018.1403013

144. Amblard L and Mann C. Understanding collective action for the achievement of EU water policy objectives in agricultural landscapes: Insights from the Institutional Design Principles and Integrated Landscape Management approaches. Environmental Science & Policy 2021; 125:76–86. DOI: 10.1016/j.envsci.2021.08.015

145. Turri VM. Remunicipalisation of Water Services in Europe: Comparative Study of the Neapolitan and Parisian Cases. Water Policy 2022; 24:1842–58. DOI: 10.2166/wp.2022.090

146. Vewin – Vereniging van Waterbedrijven in Nederland. Drinking Water Fact Sheet 2024: The Dutch Water Supply Sector. Vewin, 2024. Available from: https://www.vewin.nl/wp-content/uploads/2024/09/Vewin-Drinking-Water-fact-Sheet.pdf [Accessed on: 2025 Oct 16]

147. Rijkswaterstaat. National Water Program 2022–2027: Managing Water for a Safe and Sustainable Future. Ministry of Infrastructure and Water Management, 2024. Available from: https://iplo.nl/publish/pages/199381/reaction-paper-to-opinions-and-recommendations-national-water-programme-2022-2027.pdf [Accessed on: 2025 Nov 5]

148. Organisation for Economic Co-operation and Development (OECD). Implementing the OECD Principles on Water Governance: Indicator Framework and Evolving Practices. OECD Publishing, Paris, 2018. Available from: https://www.oecd.org/content/dam/oecd/en/publications/reports/2018/03/implementing-the-oecd-principles-on-water-governance_g1g8b14c/9789264292659-en.pdf [Accessed on: 2025 Oct 11]

149. Dijk L van. Improving Collaboration in Regional Mobility Hub Projects: Strategies to Improve Stakeholder Collaboration in the Development and Implementation Phases of Regional Mobility Hubs. TIL5040 Thesis, Faculty of Civil Engineering and Geosciences. MA thesis. Delft, The Netherlands: Delft University of Technology, 2025. Available from: https://repository.tudelft.nl/ [Accessed on: 2025 Oct 16]

150. Huijskens M. Farmers' Responses to Transformations in Dutch Water Management Policies: Navigating the Shift from Intensive Agriculture to Nature Conservation in West-Brabant. Master thesis, Knowledge, Technology and Innovation; Water Resources Management Chair Groups. MA thesis. Wageningen, The Netherlands: Wageningen University & Research, 2025. Available from: https://edepot.wur.nl/693905 [Accessed on: 2025 Oct 16]

151. Augustijn MFW, Alblas E, and Richter A. Catching up with Drought: Law and Policy Responses in the Netherlands. Regional Environmental Change 2025; 25:1–19. DOI: 10.1007/s10113-025-02449-y

152. Esteban TAOE, Heijmeskamp T, and Sjamsoeoed Sadjad M. Where We Stand: Risk Perceptions on Changing Neighborhoods in South Rotterdam. Atlantis 2024; 35:45–52. Available from: https://pure.eur.nl/en/publications/where-we-stand-risk-perceptions-on-changing-neighborhoods-in-sout [Accessed on: 2025 Oct 16]

153. Samuel AL. Some Studies in Machine Learning Using the Game of Checkers. IBM Journal of Research and Development 1959; 3:210–29. DOI: 10.1147/rd.33.0210

154. Douglas DG. The Evolution of Large Technological Systems. *The Social Construction of Technological Systems: New Directions in the Sociology and History of Technology*. Ed. by Bijker WE, Hughes TP, and Pinch T. Cambridge, MA: MIT Press, 1987 :51–82. Available from: http://www.jstor.org/stable/j.ctt5vjrsq [Accessed on: 2025 Oct 19]

155. Ostrom E. Understanding Institutional Diversity. Princeton University Press, 2005. DOI: 10.1007/s11127-007-9157-x

156. Donge W, Janssen M, and Bharosa N. Preparing Public Agencies for Harnessing AI: A Study on Variables Shaping Multi Actor Information Infrastructures. Government Information Quarterly 2024. DOI: 10.1007/978-3-031-70274-7_16
157. European Commission. Common European Data Spaces. https://digital-strategy.ec.europa.eu/en/policies/data-spaces. Accessed: 2025-10-06. 2025
158. Yadav S and Rao S. Treated sewage effluents as a source of microbiological contamination on receiving watersheds. *Environmental Biotechnology for Soil and Wastewater: Implications on Ecosystems*. Springer, 2019. DOI: 10.1007/978-981-13-6846-2_10
159. Chen Y, Shen W, Wang B, Zhao X, Su L, Kong M, Li H, Zhang S, and Li J. Occurrence and fate of antibiotics, antimicrobial resistance determinants and potential human pathogens in a wastewater treatment plant and their effects on receiving waters in Nanjing, China. Ecotoxicology and Environmental Safety 2020; 206:111371. DOI: 10.1016/j.ecoenv.2020.111371
160. Abramova AA, Isakov VG, Dyagelev MY, Nepogodin AM, and Grakhova EV. Specific pollutants in urban wastewater and their impact on the environment. *IOP Conference Series: Earth and Environmental Science*. Vol. 981. IOP Publishing, 2022 :032083. DOI: 10.1088/1755-1315/981/3/032083
161. Rastogi M and Nandal M. Toxic metals in industrial wastewaters and phytoremediation using aquatic macrophytes for environmental pollution control: An eco-remedial approach. *Bioremediation of Industrial Waste for Environmental Safety*. Springer, 2019 :257–81. DOI: 10.1007/978-981-13-3426-9_11
162. Bharagava RN and Saxena G. Bioremediation of industrial waste for environmental safety: Biological agents and methods for industrial waste management. *Bioremediation of Industrial Waste for Environmental Safety*. Springer, 2020. DOI: 10.1007/978-981-13-3426-9
163. He R, Wu X, Mu H, Chen L, Hu H, Wang J, Ren H, and Wu B. Priority control sequence of 34 typical pollutants in effluents of Chinese wastewater treatment plants. Water Research 2023; 243:120338. DOI: 10.1016/j.watres.2023.120338
164. SUEZ Water Technologies & Solutions. Activated Sludge Reactor: Type and Configurations. Available from: https://www.suezwaterhandbook.com/processes-and-technologies/biological-processes/suspended-growth-cultures/activated-sludge-reactor-type-and-configurations [Accessed on: 2025 Sep 9]
165. SUEZ Water Technologies & Solutions. Sequencing Batch Reactor Process. Available from: https://www.suezwaterhandbook.com/processes-and-technologies/biological-processes/suspended-growth-cultures/sequencing-batch-reactor-process [Accessed on: 2025 Sep 9]
166. Rahman TU et al. The Advancement in Membrane Bioreactor (MBR) Technologies for Wastewater Treatment: A Review. International Journal of Environmental Research and Public Health 2023; 20:652. Available from: https://www.mdpi.com/2077-0375/13/2/181 [Accessed on: 2025 Jul 21]
167. Li J, Jiang J, Li J, He C, Luo Y, and Wei L. Anaerobic membrane bioreactors for wastewater treatment: mechanisms, fouling control, novel configurations, and future perspectives. Environmental Engineering Research 2023; 28. Available from: http://eeer.org/upload/eer-2021-575.pdf [Accessed on: 2025 Jul 21]
168. Hamza RA et al. A review of the state of development of aerobic granular sludge technology over the last decade: Progress made and remaining challenges. Environmental Challenges 2022; 7:100495. DOI: 10.1016/j.cscee.2021.100173
169. Ogata K and Yang J. Modern control engineering. Prentice Hall, 2010. Available from: https://dl.acm.org/doi/book/10.5555/241683 [Accessed on: 2025 Jul 5]
170. Åström KJ and Hägglund T. Advanced PID control. ISA-The Instrumentation, Systems, and Automation Society, 2008. Available from: https://www.isa.org/getmedia/fb0e41bc-e4f3-422a-9f67-b9bd31340e16/Advanced-PID-Control_AstromHagglund_Chapter1-Introduction.pdf [Accessed on: 2025 Jul 5]

171. Rawlings JB, Mayne DQ, and Diehl M. Model predictive control: theory, computation, and design. Nob Hill Publishing, 2017. Available from: https://sites.engineering.ucsb.edu/~jbraw/mpc/MPC-book-2nd-edition-1st-printing.pdf [Accessed on: 2025 Jul 5]

172. Bijker WE. Of Bicycles, Bakelites, and Bulbs: Toward a Theory of Sociotechnical Change. MIT Press, 1995. Available from: https://sciencepolicy.colorado.edu/students/envs_5110/bijker.pdf [Accessed on: 2025 Aug 14]

173. Howlett M, Ramesh M, and Perl A. Studying Public Policy: Policy Cycles and Policy Subsystems. 3rd. Oxford University Press, 2009. Available from: https://global.oup.com/academic/product/studying-public-policy-9780199026142 [Accessed on: 2025 Jul 13]

174. Lasswell HD. The Policy Orientation. Ed. by Lerner D and Lasswell HD. Stanford University Press, 1951 :3–15. Available from: https://www.scirp.org/reference/referencespapers?referenceid=2659395 [Accessed on: 2025 Jul 13]

A

Appendix to Chapter 1: Additional and Fundamental Concepts

A.1. Additional Concepts of Environmental Engineering

Main Water Pollutants and Their Effects - Extended Version

As anticipated in Chapter 1, the main water contaminants, their effects on human health, the environment, and what is typically done for their removal are all explained in this Appendix. Extensive research on the topic is easily available in the literature, as seen for example by the work of Akpor and colleagues [34], which is the main reference used for this subsection of the introduction.

A first group of pollutants consists of excess nutrients, particularly nitrogen (N) and phosphorus (P), which remain only partially removed in many WWTPs. These components, when released in aquatic habitats, trigger algal blooms, causing eutrophication, oxygen reduction, fish deaths, and long-term biodiversity loss [34]. This ecological damage directly impacts water supply, recreational use, and fisheries. Eutrophication remains an EU-wide issue, affecting more than 30% of rivers, lakes and coastal waters and 80% of EU marine waters. As a matter of fact, in the Urban Waste Water Treatment Directive recast (2024), provisions have been made up to date and harmonised to further stress nutrient removal as a priority in WWTPs [35]. It is then clear how nutrient management remains one of the main long-standing priorities in wastewater research and practice.

Another threat is microbiological contamination. Domestic wastewater contains a diverse array of pathogenic microorganisms, including bacteria (e.g. *Escherichia coli* and *Salmonella*), viruses (e.g. Rotavirus and Hepatitis A), and protozoan parasites (e.g. *Giardia* and *Cryptosporidium*). If not properly treated, these pathogens can reach water bodies, leading to outbreaks of diseases like cholera, typhoid, and hepatitis, especially in lower-income regions [158]. Although disinfection processes are usually applied, their persistence in effluents shows that traditional treatment can fail to perform under conditions of high load or poor operational control and more advanced strategies should be developed.

In recent years, WWTPs have also been identified as hotspots for the spread of antibiotics and antimicrobial resistance genes (ARGs). This is because biological processes can only partially eliminate antibiotics from wastewater. As a matter of fact, fluoroquinolones, macrolides, tetracyclines, and other commonly utilised classes remain detectable in effluents even after secondary or tertiary treatment [159, 160]. If their emission in the environment is consistent, this not only poses risks for aquatic organisms but also facilitates ARGs diffusion through horizontal gene transfer, presenting a significant threat to global public health.

Industrial effluents introduce another set of pollutants, typically associated with carcinogenicity, neurological impairment, and organ toxicity [161, 162]. These elements are heavy metals, such as cadmium, mercury, arsenic, and chromium, which are non-biodegradable and tend to bioaccumulate in aquatic food chains. While precipitation and ion exchange are mainly used for their removal, research is increasingly focusing on eco-friendly strategies such as phytoremediation.

Another group of contaminants still associated with industrial discharges is represented by organic micropollutants such as phenols, hydrocarbons, surfactants, and synthetic dyes. These compounds are often toxic, carcinogenic, and persistent in the environment, reducing light penetration and inhibiting microbial activity in WWTPs themselves [34, 162]. Their presence further complicates treatment and has motivated research into biodegradation, constructed wetlands, and nanomaterials.

Finally, endocrine-disrupting compounds (EDCs) and pharmaceuticals are another group of pollutants with persistence in treated effluents. Diclofenac, acetaminophen, and synthetic estrogens act at minimal concentrations, disrupting hormonal regulation and reproduction in fish and aquatic invertebrates [163]. Chronic exposure is linked to population decline in wildlife and potential carcinogenic and reproductive diseases in humans. Conventional systems often fail to achieve sufficient removal, requiring advanced oxidation processes, membrane bioreactors, and activated carbon adsorption.

Literature generally shows that wastewater effluents contain a complex mixture of pollutants, each with distinct toxicological profiles and removal challenges. Nutrients and pathogens remain the most widespread threats to ecosystems and public health, while antibiotics, pharmaceuticals, EDCs, and heavy metals represent an evolving frontier in water pollution control. The coexistence and potential interactions of these pollutants demand advanced monitoring and treatment approaches. In this context, data-driven and AI-based methods offer promising paths to improve conventional treatment, enabling the real-time detection of contaminants and the optimisation of removal processes.

Main Reactor Configurations for the Activated Sludge Process

As mentioned in Chapter 1, the activated sludge process is implemented through different reactor configurations that influence treatment performance and operational efficiency. According to SUEZ's Degremont Water Handbook, several reactor types are commonly used, each offering specific advantages depending on plant scale, influent characteristics, and energy constraints [164]. The most common configurations are described below; however, a more complete view of the process will be gained after reading the subsection dedicated to aeration Section 1.4.4, which completes this one.

In *complete-mix reactors*, all components, including microorganisms, oxygen, and substrate (which is the organic matter to be decomposed), are uniformly distributed, making these systems robust against shock loads but more susceptible to filamentous bulking under low food-to-microorganism (F/M) ratios [164]. Below, a visual scheme for this configuration is reported Figure A.1.

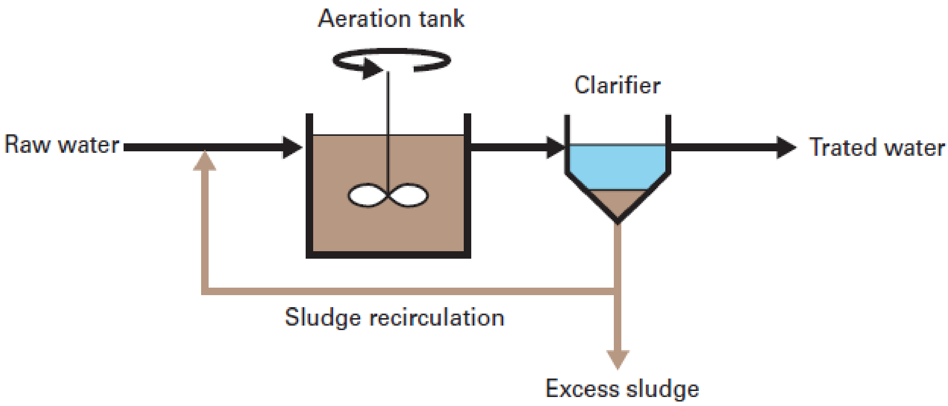

Figure A.1: Complete-mix reactor configuration (© SUEZ, adapted from Degremont Water Handbook [164]).

Plug-flow reactors, by contrast, consist of long channels that create gradients in substrate and oxygen concentration along their length, as can be seen in Figure A.2. These systems are typically used in large plants and benefit from tapered aeration, which means that the highest amount of air is injected at the beginning and gradually decreases as the process progresses [164].

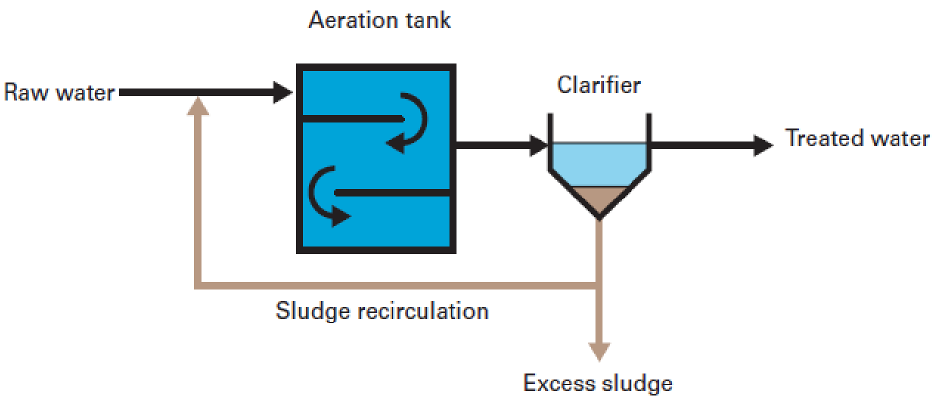

Figure A.2: Plug-flow reactor configuration (© SUEZ, adapted from Degremont Water Handbook [164]).

Step-feed reactors represent a variation of plug flow systems, where influent is introduced at multiple points to improve oxygen and substrate distribution, as can be seen in Figure A.3. This enhances overall performance [164].

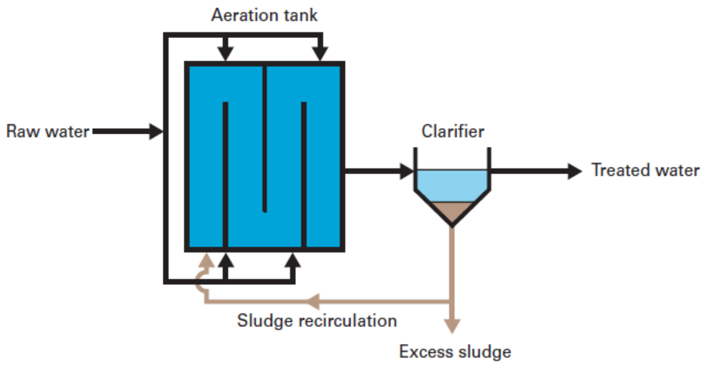

Figure A.3: Step-feed reactor configuration (© SUEZ, adapted from Degremont Water Handbook [164]).

Ditch-type reactors (e.g., oxidation ditches), as the one in Figure A.4 promote circulation in closed loops, enabling alternating aerobic and anoxic conditions, while selector tanks, usually placed upstream, help suppress filamentous bacteria and encourage the growth of floc-forming organisms [164].

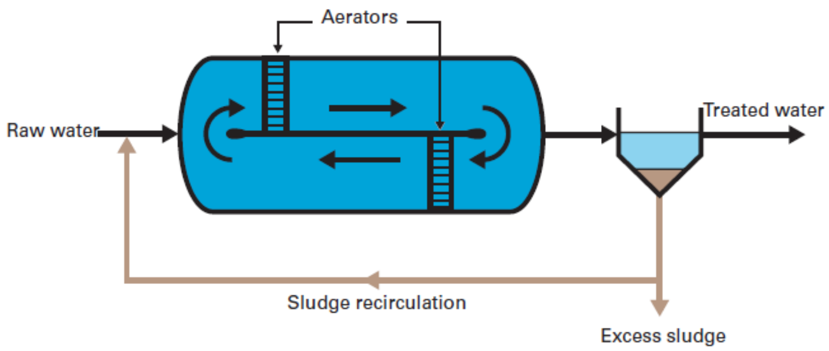

Figure A.4: Ditch-type configuration (© SUEZ, adapted from Degremont Water Handbook [164]).

Additionally, for enhanced nutrient removal, SUEZ also implements advanced configurations. For example, the *Modified Ludzack-Ettlinger (MLE) process* (Figure A.5) incorporates an upstream anoxic zone followed by an aerobic zone, with internal recirculation to support denitrification [164].

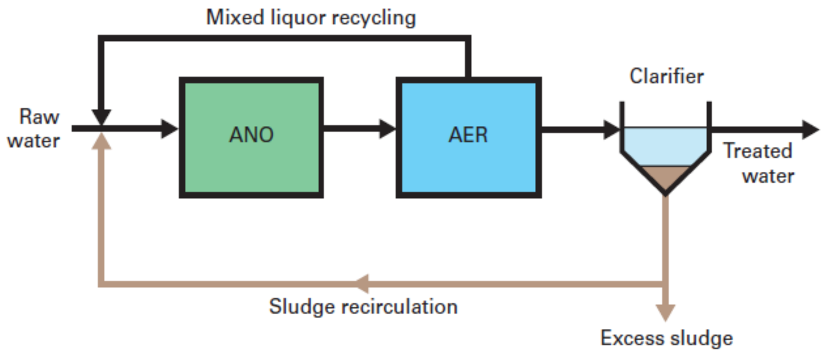

Figure A.5: Modified Ludzack–Ettlinger (MLE) process (© SUEZ, adapted from Degremont Water Handbook [164]).

A more refined approach, the *three-zone process*, introduces a final anoxic compartment designed for endogenous denitrification (Figure A.6), improving nitrogen removal under low BOD/N conditions [164].

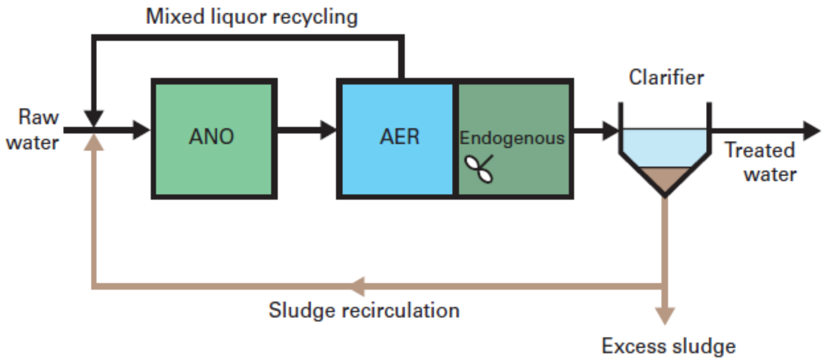

Figure A.6: Three-zone process (© SUEZ, adapted from Degremont Water Handbook [164]).

Nitrification-denitrification ditches (Figure A.7) use sequenced aeration, controlled by sensors such as Greenbass™, to alternate between aerobic and anoxic phases [164].

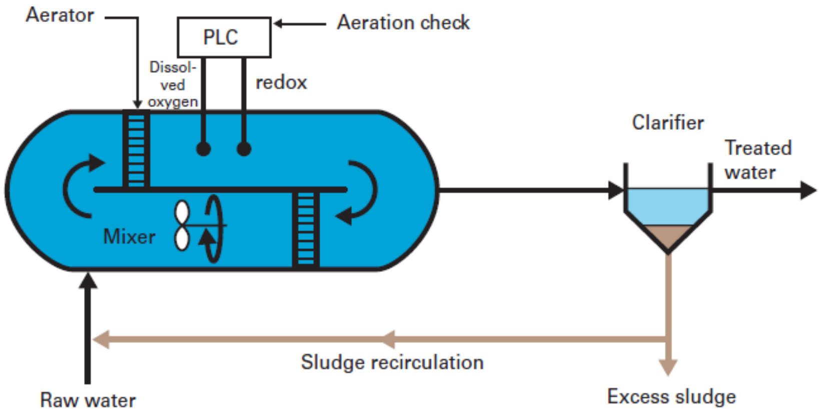

Figure A.7: Nitrification–denitrification ditch (© SUEZ, adapted from Degremont Water Handbook [164]).

In high-capacity plants, *multiple-stage processes* arrange successive anoxic and aerobic zones in series (as shown in Figure A.8), allowing for efficient nitrogen removal without excessive internal recirculation [164].

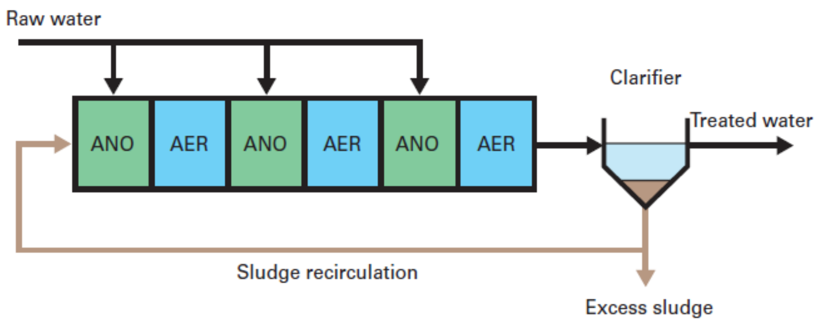

Figure A.8: Multiple-stage reactor configuration (© SUEZ, adapted from Degremont Water Handbook [164]).

An alternative to continuous-flow systems, the *Sequencing Batch Reactor (SBR)* integrates all treatment phases - filling, aeration, settling, and decanting -within a single tank operated in cycles. This time-based operation offers flexibility and space efficiency, particularly advantageous for small and medium-sized facilities or where influent flow varies widely. As described by SUEZ [165], SBRs eliminate the need for separate clarifiers and allow for precise control over process timing, leading to enhanced nutrient removal efficiency, reduced footprint, and simplified automation. However, although simple and inexpensive, this method has limited efficiency and is unsuitable for high-load activated sludge systems [2]. Below, in Figure A.9, is shown a visualization of the main phases that happen in an SBR.

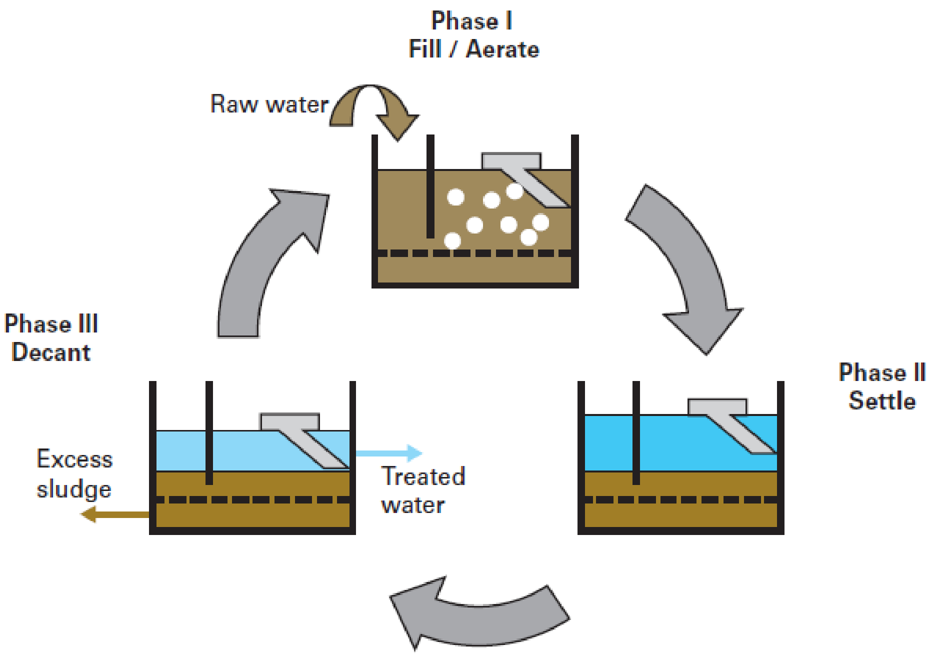

Figure A.9: Sequencing batch reactor (SBR) process (© SUEZ, adapted from Degremont Water Handbook [165]).

Other modern solutions not presented in SUEZ's Degremont Water Handbook, but still worth mentioning, are *membrane bioreactors (MBRs)* and more recently, *aerobic granular sludge (AGS) systems*.

MBRs simply combine biological treatment with membrane filtration, enabling operation at higher sludge concentrations, smaller reactor volumes, and better effluent quality, even if with higher capital and energy demands [166, 167].

On the other hand, AGS systems rely on the formation of dense microbial granules, allowing for excellent settling, simultaneous nutrient removal, and lower energy requirements compared to conventional systems [168].

A.2. Foundations of Control Engineering for a Broad Audience

In this part of the Appendix, foundational concepts of Systems and Control Engineering are explained at a high level for readers with a non-technical background. Please note that this is not to be intended as a supplement to the literature review in Section 2.1.2, but as a necessary effort to make the research self-standing and understandable from different perspectives, given its scope of bridging the gap between engineers and policy-makers. Therefore, the following definitions and explanations are given.

Open-Loop and Closed-Loop Control

Controllers are systems that, by the use of actuators, can modify the state of another system to achieve a desired behaviour. A very simple example is a light switch (controller) that, by activating a lamp (actuator), can modify the state of the system (luminosity of a room). Controllers can be divided into open-loop and closed-loop or feedback controllers [169].

In an open-loop control system, actions are executed according to a predefined plan, generally time-based, without checking whether the desired outcome has been achieved. For example, open-loop control is used in irrigation systems that are activated at the same hours, regardless of whether it has just rained, or it is raining (or even if it will soon rain) or not.

On the other hand, in a closed-loop control system, sensors continuously measure the actual system behaviour (state), compare it to the desired target, and adjust actions accordingly. A simple example is a home thermostat that switches heating on or off depending on the measured room temperature. This is the foundation of most modern engineering applications, including WWTPs.

Relay Controllers

Relay controllers are the simplest forms of feedback control. They are reactive, as they act only in response to current deviations from a desired behaviour measured by sensors. The controller switches the actuator between two discrete states (e.g., ON or OFF) depending on whether the measured variable is above or below a desired threshold. To avoid continuous switching on/off, a hysteresis band is introduced: the system only changes state when the variable crosses an upper or lower limit.

For instance, in wastewater aeration, when the dissolved oxygen (DO) in the tank gradually decreases due to biological activity, the aeration pump is switched ON once DO falls below a certain level (e.g., 0.2 mg/L). Then, the pump remains active until DO concentration exceeds the upper bound (e.g., 3 mg/L), at which point the pump is switched OFF. In this way, the hysteresis band (0.2–3 mg/L) prevents the pump from turning on and off too frequently when DO fluctuates around a single threshold.

Extensions of this idea include multi-level relay controllers, which operate at different discrete levels by one actuator (e.g., one pump), or the simultaneous switching of multiple actuators (e.g., more pumps).

Proportional-Integral-Derivative (PID) Controllers

PID controllers are also feedback controllers, and as relays they are reactive in nature [170]. They continuously adjust the control action based on three terms:

- The Proportional (P) term reacts to the current error, that is the current difference between desired and measured value.
- The Integral (I) term considers the accumulation of past errors, to face steady-state offsets.
- The Derivative (D) term predicts short-term changes by considering the error's rate of change, but still as a reaction to measured deviations. This term is sometimes omitted as it is very sensitive to burst errors.

By combining these components, PID controllers can achieve robust and stable regulation in many processes, including those found in water treatment. With respect to a relay, the control can assume values that are tuned according to the current need, while in the relay, the control, when activated, is fixed at given values.

Optimisation-Based Controllers: From Linear Optimisation to MPC

Optimisation techniques are increasingly used in modern control. While linear or non-linear optimisation alone, for example, based on mathematical programming, is not a controller, it may represent the mathematical foundation of advanced control schemes.

Model Predictive Control (MPC) [171] is a genuine feedback controller, but unlike PID or relay, it is predictive rather than purely reactive. MPC forecasts the future evolution of the system using a simplified model, mostly showing the trend of state variables of interest, and then solves an optimisation problem to select the best control actions over a finite horizon as a model-based open-loop control. Only the control at the first time interval is implemented, and in the next step, the optimisation is repeated using updated measurements. This ”rolling horizon” strategy makes MPC effectively a closed-loop controller. An obvious requirement is that the time to solve the optimisation problem is sensibly lower than the duration of the time interval.

Example: Suppose that a wastewater aeration tank is subject to a known upcoming increase in inflow, or that, by a model, it is possible to forecast an important future bacterial oxygen demand. A reactive controller (relay or PID) would only respond later, when the dissolved oxygen starts to drop, possibly causing a temporary violation of quality standards. MPC, by contrast, can anticipate the increased demand for oxygen based on the forecast and activate aeration before the concentration falls, thereby avoiding the deviation altogether. This anticipatory ability is the hallmark of predictive control.

AI-Based Controllers: NN, RL, and DRL

Recent advances in Artificial Intelligence (AI) have introduced new paradigms in control. These can be used either as stand-alone controllers or in combination with traditional methods:

- Neural Network (NN)-based controllers are typically implemented as feedback controllers. They are predictive in the sense that they approximate non-linear system behaviour from historical data, resuming previous control experiences, and their decisions remain reactive to sensor inputs.
- Reinforcement Learning (RL) [102] explicitly learns control policies by trial-and-error interaction with the system or its simulation model. The resulting controller is a feedback policy, which can be reactive or predictive depending on the reward formulation and the available information.
- Deep Reinforcement Learning (DRL) extends RL by using deep neural networks [103], allowing the controller to learn from complex, high-dimensional data. In principle, DRL can combine reactivity with long-term prediction, but challenges remain regarding safety, stability, and interpretability in critical infrastructure applications.

Although not yet as widely deployed in practice as PID or MPC, AI-based controllers are actively researched for systems like wastewater treatment, where dynamics are highly non-linear and multiple objectives must be considered simultaneously.

A.3. Foundational Policy Concepts for a Broad Audience

In the same way, it has been decided to provide a brief high-level explanation of the main concepts in policy analysis and decision science. The goal is only to equip technical readers with a high-level conceptual map to follow the socio-technical discussion that complements the engineering analysis.

Socio-Technical System

A second key concept is that of a *socio-technical system*. Originating in systems theory and technology studies, this term describes systems in which technical components (machines, infrastructures, algorithms) and social components (people, institutions, norms) are deeply interdependent [154, 172]. Wastewater treatment plants exemplify such systems: they consist not only of reactors, sensors, and control units, but also of operators, regulators, and policies that govern their operation. This framing explains why the introduction of AI controllers cannot be analysed solely as an engineering challenge, but must also consider institutional and governance dimensions.

Actor

Another term that will be widely used in this research is *actor*. In the context of governance and socio-technical systems, an actor is any individual, group, or organization that influences decision-making and outcomes [155]. This includes regulators, utilities, private firms, and civil society organizations. Later in the thesis, an explicit multi-actor perspective will be applied to map these roles and their interactions.

Policy: Control Engineering and Public Governance Meanings

Another important clarification concerns the notion of *policy*, which is used with different meanings in this thesis. In control engineering, a *control policy* refers to a set of rules or algorithms that maps observed states to actions. For example, in a reinforcement learning controller, the *control policy* has to be viewed in this way. In public governance, however, *policy* refers to a course of action adopted by governments or institutions to achieve societal goals [173]. To avoid confusion, the text will always make explicit which meaning is intended. The definition of policy still corresponds to a set of rules in both cases; however, these are applied to different systems and at a very different scale.

Policy vs Politics

In this governance sense, it is also helpful to distinguish between *policy* and *politics*: whereas policy concerns the substantive content of collective decisions, politics refers to the processes of contestation, negotiation, and power through which those decisions are made [174].

B

Appendix to Chapter 4: Technical Tools and Further Models Tested

B.1. Technical Tools

This thesis employed computational and modelling tools to develop and evaluate deep reinforcement learning (RL) control strategies for wastewater treatment optimisation. The implementation combined Python, as the main development environment, and MATLAB, used in the early stages for model calibration and parameter fitting. Together, these environments supported the development of digital twins, data-driven models, and RL controllers for aeration processes.

Programming Languages

- Python: used for data processing, simulation, and learning-based control, integrating mechanistic models, neural networks, and RL algorithms.
- MATLAB: implemented in initial phases for model calibration, system identification, and testing simplified control laws through optimisation routines and rolling-horizon simulations.

Python Libraries and Frameworks

Table B.1 summarises the main Python libraries and their roles in the project.

Category	**Library / Tool**	**Purpose / Application**
Numerical Computing	`numpy`, `scipy.integrate`, `pandas`	Solving ODE systems, handling time-series data, and computing process variables.
Machine Learning	`torch`, `sklearn`, `xgboost`	Training data-driven models, feature selection, and predictive modelling.
Reinforcement Learning	`stable_baselines3`, `gymnasium`	Implementation of deep RL algorithms interacting with custom Gym environments.
Data Engineering	`pandas`, `sklearn.preprocessing`, `joblib`	Preprocessing, scaling, and management of historical data and replay buffers.
Visualisation	`plotly`, `matplotlib`, `seaborn`	Visualisation of sensor data, model predictions, training curves, and policy performance.
Optimisation	`scipy.optimize`	Parameter calibration of kinetic models and hyperparameter tuning.
Project Management	`pathlib`, `json`, `yaml`, `argparse`	Configuration handling and reproducible experiment setup.

Table B.1: Main Python libraries and their roles in the project.

Computational Environment and Reproducibility

All experiments were executed in Jupyter Notebooks within a dedicated `conda` environment (Python 3.10). Version control was managed through a `Git` repository with a modular folder hierarchy separating data preprocessing, model calibration, and RL training pipelines. Configurations and metadata were stored in `JSON` and `YAML` files to ensure full reproducibility and traceability across experiments. Visual analytics were produced using `Plotly`, `Matplotlib`, and `Seaborn`. This toolchain provided a unified and transparent computational ecosystem supporting the subsequent modelling and control analyses described in the following sections.

B.2. Additional Models Tested: ExtraTreesRegressor and LSTM-ATT

As explained in Chapter 4, more models have been developed to train the deep RL agent. Since these were not yet implemented, but might be used in future steps, once the agent has learned the key dynamics of the system and how to respect the operational constraints of the pumps, they are reported here.

ExtraTrees Regressor Models for DO and Redox Prediction (V2)

This section describes the development of two data-driven regression models based on the *ExtraTrees* algorithm, designed to predict the next-step dissolved oxygen (DO) and redox potential (ORP) in the aeration tank of the wastewater treatment plant.

Objective

The aim was to predict, for each 15-minute sampling interval:

$$\text{DO}_{t+1} \quad \text{and} \quad \text{Redox}_{t+1}$$

given current process conditions (airflow rate, inflow, rainfall, and system status). Such short-term forecasts allow rapid estimation of aeration dynamics without requiring full biochemical simulation.

Data and variables

The training data (`train_data.csv`) and test data (`test_data.csv`) are the same from Chapter 4. Table B.2, below, summarises the key operational data utilised.

Variable	Description	Unit
DO sensor	Dissolved oxygen concentration	mg/L
Redox sensor	Oxidation–reduction potential	mV
Influent flow rate	Flow entering the biological tank	m^3/h
Airflow rate	Aeration flow supplied to the tank	Nm^3/h
Rainfall	Accumulated precipitation over 5 minutes	mm
Blower status	Binary indicator of aeration system activation	–

Table B.2: Key variables used for ExtraTrees model training.

Feature Engineering

The feature-builder function generated both temporal and process-related predictors:

- Calendar features: day, day of month, day of week, and hour (fractional);
- Process inputs: inflow rate, rainfall, airflow, and blower status;
- Airflow lags at 15 min, 30 min, 60 min, and 90 min to capture short-term actuator inertia;
- Current sensor state: either DO or Redox (depending on target).

These final 13 features were selected through forward feature selection.

Model Configuration

Each target was modelled separately using a scikit-learn `ExtraTreesRegressor` with the following hyper-parameters:

```
n_estimators      = 600
max_features      = "sqrt"
max_depth         = 18
min_samples_split = 5
min_samples_leaf  = 3
random_state      = 42
```

ExtraTrees is an ensemble of randomized decision trees that averages multiple de-correlated trees for robust nonlinear regression. It is well suited for heterogeneous environmental datasets due to its low variance, automatic feature selection, and ability to capture nonlinear dependencies between aeration inputs and output variables.

Training and Artifacts

Models were trained on the training dataset. Each run generated:

- two `.pkl` model files (`model_do`, `model_redox`);
- JSON files listing feature names and target columns;
- and a summary metadata file (`latest_v2.json`) recording metrics and timestamps.

The complete workflow is automated for reproducibility and future retraining.

Evaluation Procedure

The evaluation notebook automatically detected the project root, loaded the latest artefacts, and computed test-set metrics shown in Table B.3.

Target	R^2	MSE	Samples
DO (next 5 min)	0.34	0.378 mg^2/L^2	6 521
Redox (next 5 min)	0.55	1 376.86 mV^2	6 521

Table B.3: Test-set performance of ExtraTrees Regressor models.

Time-series and scatter plots (True vs. Predicted) were produced using Plotly to visually assess model fidelity. The moderate R^2 values reflect the intrinsic noise and short time horizon of the aeration process.

ExtraTree Performance Plots

Below Figure B.1 and Figure B.2 respectively show predictions for Redox and DO.

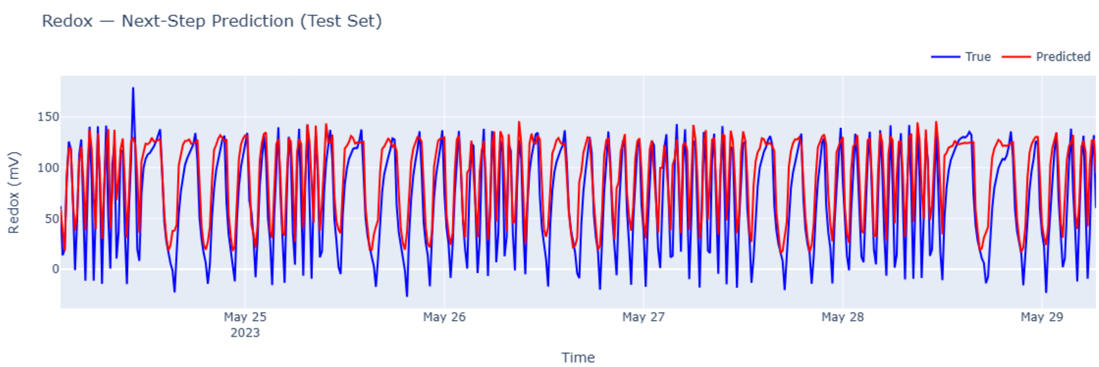

Figure B.1: ExtraTree Regressor as predictor for Redox (ORP)

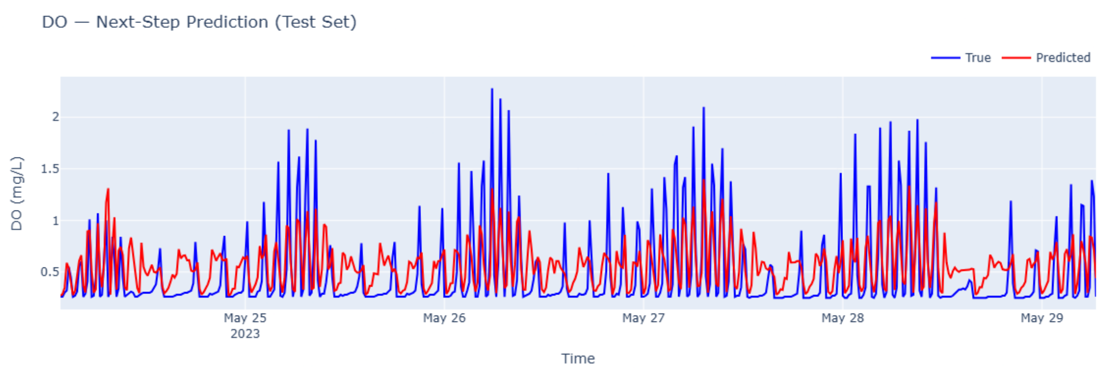

Figure B.2: ExtraTree Regressor as predictor for Dissolved Oxygen (DO)

LSTM-ATT Models for Aeration (Version 2r)

This section presents the development of three refined deep learning models based on the *LSTM-Attention* architecture (v2r) trained to predict the next-step concentrations of dissolved oxygen (DO), redox potential (ORP), and ammonium (NH_4^+) in the aeration tank. These models extend the ExtraTrees baseline by capturing temporal dependencies and nonlinear process dynamics, serving as advanced data-driven surrogates for the digital twin and reinforcement learning environment.

Objective

The goal was to predict, for each 15-minute interval:

$$\text{DO}_{t+1}, \quad \text{Redox}_{t+1}, \quad \text{NH}_{4,t+1}$$

based on a sequence of the previous twelve time steps (equivalent to three hours of plant operation), including process inputs, environmental variables, and current sensor states.

Data and Variables

The models were trained and tested on the same dataset used in Chapter 4. The key input variables are summarised in Table B.4.

Variable	**Description**	**Unit**
DO sensor	Dissolved oxygen concentration	mg/L
Redox sensor	Oxidation–reduction potential	mV
Ammonium sensor	Ammonium concentration in mixed liquor	mg/L
Influent flow rate	Flow entering the biological tank	m^3/h
Airflow rate	Aeration flow supplied to the tank	Nm^3/h
Rainfall	Accumulated precipitation over 5 minutes	mm
Blower status	Binary indicator of aeration system activation	–

Table B.4: Main variables used as inputs for the ATT-LSTM v2r models.

Feature Engineering and Sequence Construction

Features include both temporal and process descriptors:

- Calendar features: day, day of month, day of week, and hour (fractional);
- Process inputs: inflow rate, rainfall, airflow, and blower status;
- Airflow lags at 15, 30, 60, and 90 minutes to capture actuator inertia;
- Current sensor state of the respective target variable (DO, Redox, or NH_4^+).

Each model uses input sequences of length $L = 12$ (representing 3 hours). All features and targets were scaled using the `MinMaxScaler` before training, and aligned to remove missing values. Outliers beyond the 1st and 99th percentile were clipped to improve stability.

Model Architecture

The predictive model combines a single-layer **Long Short-Term Memory (LSTM)** network with an **attention mechanism** that adaptively weights recent time steps according to their relevance for forecasting the next output value.

The model, implemented in PyTorch, is defined as:

$$\hat{y}_{t+1} = f_\theta\big(X_{t-L+1:t}\big),$$

where f_θ denotes the LSTM-Attention mapping function and $X_{t-L+1:t}$ is the input sequence.

The main hyperparameters are listed below:

```
SEQ_LEN        = 12
BATCH_SIZE     = 128
HIDDEN_DIM     = 64
NUM_LAYERS     = 1
LEARNING_RATE  = 5e-4
EPOCHS         = 120
PATIENCE       = 15
```

The training used the `AdamW` optimiser, a `SmoothL1Loss` criterion, and gradient clipping to prevent exploding gradients. Early stopping with a patience of 15 epochs was applied to avoid overfitting.

Training Procedure and Artifacts

Each target (DO, Redox, NH_4^+) was trained independently using the same architecture and training loop. During training, model checkpoints were saved whenever a new lowest validation loss was achieved. Artifacts stored under `ML_Model/artifacts_lstm_att/` include:

- `.pth` files containing the best model weights for each target;
- Corresponding `.json` metadata files with model parameters and test metrics;
- Serialized feature and target scalers (`.pkl`);
- A summary file (`latest_v2r.json`) referencing all latest checkpoints.

The process was fully automated to ensure reproducibility and version tracking.

Evaluation Procedure

After training, the models were evaluated on a hold-out test set. Predictions were inverse-transformed to physical units before computing the performance metrics (Table B.5). Visual comparisons were generated using time-series and scatter plots (True vs. Predicted) for each target.

Target	**R^2**	**MSE**	**Samples**
DO (next 15 min)	0.78	0.142 mg^2/L^2	6 500
Redox (next 15 min)	0.85	845.32 mV^2	6 500
NH_4^+ (next 15 min)	0.74	0.215 mg^2/L^2	6 500

Table B.5: Performance of ATT-LSTM v2r models on the test set. Values are indicative of the best validation checkpoints.

Interpretation

All three models achieved strong predictive accuracy, with R^2 values between 0.74 and 0.85. The Redox model performed best, indicating that the redox potential dynamics exhibit more consistent temporal dependencies than DO or NH_4^+. The LSTM-Attention architecture successfully captured delayed nonlinear relationships between aeration flow, sensor feedback, and process response.

These results confirm that the ATT-LSTM framework can serve as a reliable surrogate for short-term predictions of key biochemical states, enabling:

1. Model-based policy evaluation and synthetic data generation for RL training;
2. Surrogate simulation for real-time optimisation in digital twins;
3. Robust feature extraction and transfer learning toward multi-variable control.

ATT-LSTM Performance Plots

Figure B.3, Figure B.4, Figure B.5, Figure B.6, Figure B.7, and Figure B.8 illustrate the model predictions and scatter plots for Redox (ORP), DO, and NH_4^+ respectively.

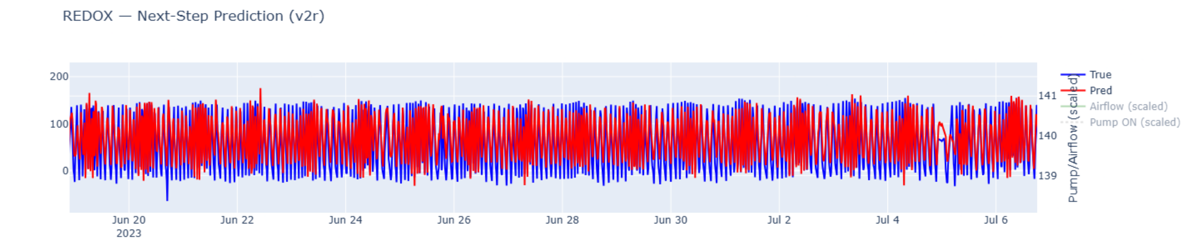

Figure B.3: ATT-LSTM v2r predictions for Redox (ORP).

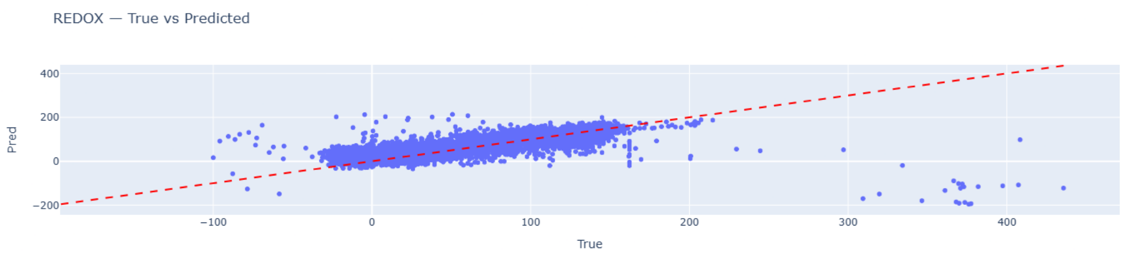

Figure B.4: ATT-LSTM v2r scatter plot for Redox (ORP).

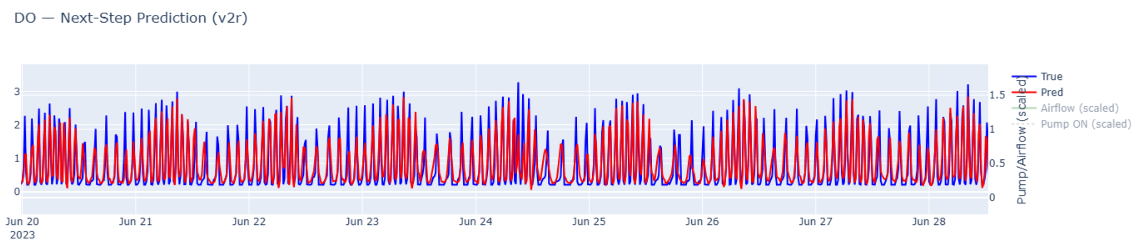

Figure B.5: ATT-LSTM v2r predictions for Dissolved Oxygen (DO).

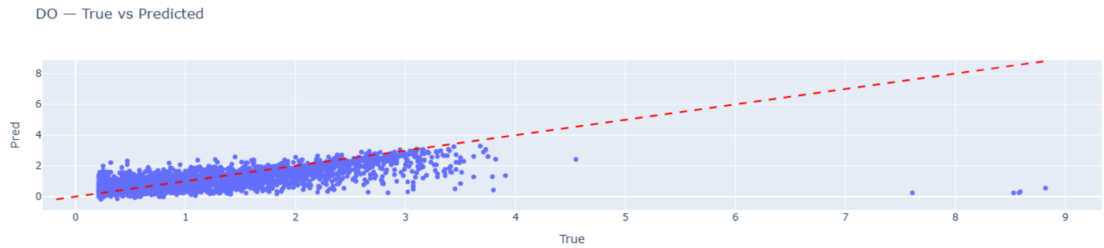

Figure B.6: ATT-LSTM v2r scatter plot for Dissolved Oxygen (DO).

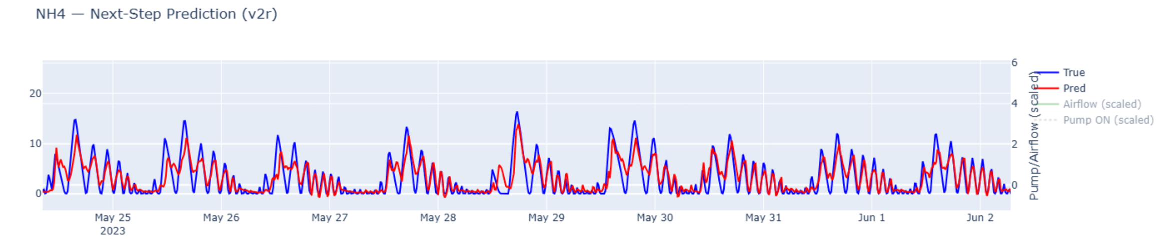

Figure B.7: ATT-LSTM v2r predictions for Ammonium (NH_4^+).

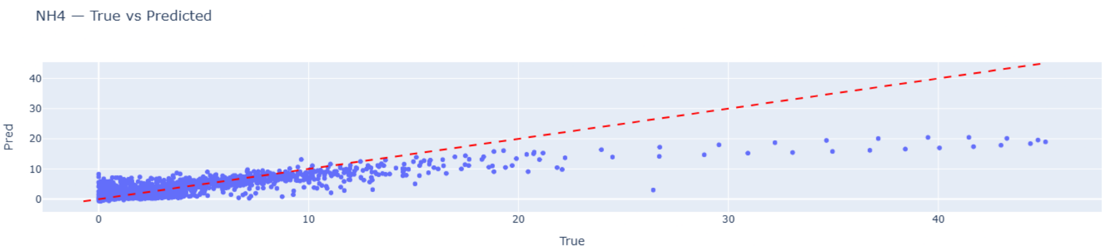

Figure B.8: ATT-LSTM v2r scatter plot for Ammonium (NH_4^+).

Reward Function Variables

Symbol	Description	Units	Used in
r_t, $rV2_t$	Instantaneous reward returned by the environment.	–	RF1, RF2
DO_t	Dissolved oxygen concentration at time step t.	mg/L	RF1, RF2
DO^{set}	DO target (setpoint) above which soft penalties begin.	mg/L	RF1, RF2
DO^{high}	High-risk DO threshold triggering stronger penalty (quadratic).	mg/L	RF2 only
ORP_t	Oxidation–reduction potential at time t.	mV	RF1, RF2
ORP^*	Optimal redox target for balanced nitrification–denitrification.	mV	RF1 only
ORP^{lo}	Lower ORP bound indicating insufficient nitrification.	mV	RF2 only
ORP^{hi}	Upper ORP bound indicating excessive aeration / over-oxidation.	mV	RF2 only
$Q_{\text{air},t}$	Physical airflow delivered by the blower.	Nm^3/h	RF1, RF2
$Q_{\min}$	Minimum stable airflow ensuring aeration lift (ON threshold).	Nm^3/h	RF1, RF2
Q_{off}	OFF tolerance threshold below which blower is considered inactive.	Nm^3/h	RF1 only
$Q_{\max}$	Maximum nominal blower capacity.	Nm^3/h	RF1 only
p_t	Electricity price at time t.	€/kWh	RF1, RF2
p_{ref}	Reference price for normalisation of energy term.	€/kWh	RF1 only
n_{on}	Consecutive ON steps.	steps	RF1, RF2
n_{off}	Consecutive OFF steps.	steps	RF1, RF2
$N^{\text{on}}_{\max}$	Maximum permissible ON duration before penalties.	steps	RF1, RF2
$N^{\text{off}}_{\max}$	Maximum permissible OFF duration before penalties.	steps	RF1, RF2
$\mathbb{1}_{\text{pump just switched}}$	Indicator of pump switching event.	–	RF1, RF2
$\mathbb{1}_{\text{safe region}}$	Indicator for DO/ORP in safe operating window.	–	RF1, RF2
$\mathbb{1}_{\text{safe off}}$	Indicator for rewarding long OFF periods only when safe.	–	RF2 only

Table B.6: Unified description of the variables and parameters used in the reward functions of Equations 4.8 and 4.9, including whether each appears in RF1, RF2, or both.

C

Interview Guide

Interview Questions

1. What is your background, current role, and experience?
2. Can you describe your current role and involvement in water/wastewater management?
3. What is your experience with digitalisation, automation, or AI tools in infrastructure?
4. How feasible is AI-based control (e.g., RL, predictive control) in real water critical infrastructures (e.g. WWTPs, dams, etc.)?
5. What are the main technical challenges? (e.g., sensors, integration, data quality)
6. Where are the main opportunities (energy efficiency, compliance, cost reduction) of applying AI-based control in critical water infrastructure?
7. What concerns do you have about the safety, robustness, or explainability of AI in critical infrastructure?
8. How important is human oversight and the ability to override AI?
9. Do you personally trust AI to manage water-related processes? Why or why not?
10. Do you think other stakeholders (operators, regulators, industry) trust AI? Why?
11. Are you aware of the key risks? How would you mitigate them?
12. What is your view on AI regulation in the US and China? Should Europe adopt similar strategies?
13. What organisational changes (skills, training, culture) are needed for AI adoption in water critical infrastructure?
14. What economic or infrastructure constraints limit adoption?
15. How do investment strategies shape AI adoption?
16. How do you interpret the obligations of the EU AI Act for water infrastructure?
17. Do current regulations support or restrict AI innovation?
18. Which policy instruments (sandboxes, funding, certifications, documentation) would help to support AI adoption?
19. Where do you see AI-based control systems in the water sector in 5–10 years?
20. What role should public–private partnerships (e.g., SUEZ + governments) play in facilitating (or not facilitating)?
21. What strategies could bridge prototypes and deployment?
22. What is the single most important factor for the adoption of AI in water systems?
23. Is there anything else you would like to add?
24. Who else would you recommend I speak with on this topic?