AI and Climate Change

How can AI Address the Energy Efficiency Objectives in the Industry?

Mathias Abitbol
Collège de France

Philippe Aghion
Collège de France, INSEAD and the London School of Economics

Céline Antonin
OFCE

Lint Barrage
ETH Zurich

The rise of generative AI has reignited the debate about the energy cost of AI. However, so far, the potential for AI to help optimize resource use, or to facilitate firms' transition from dirty to green production technologies, has been largely overlooked. No serious attempt has been made at computing the overall energy cost and/or benefit of AI. Traditional AI models require significantly less energy-intensive training and inference and could enable companies to improve the energy efficiency of their technologies and help them meet their environmental objectives for the coming years. Machine and deep learning models used in industry and deployed using digital twins can control complex processes and optimize their resource use, in energy-intensive industries, and achieve significant energy-saving potential. Wastewater treatment is emblematic in this respect: aeration during secondary treatment typically accounts for roughly half of a plant's electricity consumption. Our focus in this article is on the comparison between the direct energy cost of AI operations and the energy saving AI induces when implemented by Veolia.

Veolia, in partnership with PureControl, has rolled out one of the first large-scale deployments of AI for climate-relevant efficiency, covering about 200 plants. PureControl's system ingests high-frequency data (≈15 minutes intervals) on electricity prices, weather, sensor streams, and laboratory quality samples to maintain a live digital replica of the plant. The AI then schedules and doses aeration to minimize cost and consumption while assuring effluent quality and regulatory thresholds.

We evaluate AI's net effect using plant-level operational data, natural experiments (unplanned interruptions), and a full accounting of the AI layer's own electricity use. Preliminary results from ~15 plants indicate a nearly 10% reduction in electricity consumption and GHG emissions, while AI's direct electricity use accounts for less than 1% of the gross energy savings. Even under conservative assumptions about additional required hardware installations, the maximum lifecycle carbon cost of AI remains well below the emissions abatement, thereby pointing to a robust net-positive climate contribution.

Introduction

AI is often portrayed as an energy drain with a significant environmental impact, but this view overlooks compact, domain-specific systems embedded in industrial control. In such settings, AI acts as a continuous optimizer, aligning operations with real-time constraints and prices. Wastewater treatment offers a decisive test: secondary treatment's aeration is energy-intensive and tightly regulated.

Digital twins (virtual replicas fed by live data) offer a promising route for increased efficiency.

This study has been led by the Collège de France in partnership with Veolia and PureControl, a French start-up that develops AI-driven digital twins to optimize the electricity consumption of this process while ensuring effluent quality and regulatory compliance. Using a novel, high-granularity dataset across Veolia facilities, we quantify gross energy and GHG emissions savings, and net out the AI layer's own electricity use. We also examine load shifting toward off-peak hours, while ensuring effluent quality and regulatory compliance. Early evidence points to a nearly 10% reduction in electricity use in achievable at scale.

The Role of AI in Fostering the Energy Transition

A heated Public Debate

The rapid diffusion of AI has spurred concern that digitalization may raise global electricity demand. Media narratives often extrapolate from the power required to train frontier, Large Language Models specific to generative AI to all AI uses. The Artificial Intelligence Policy Institute ran a survey, revealing that 72% of American voters are concerned about the increasing energy consumption of AI data centers.[1] Whether these concerns are warranted depends less on aggregate energy consumption than on the pace and geographic concentration of deployment, which can strain local grid capacity. Moreover, treating AI as a single, homogeneous technology obscures crucial heterogeneity between cloud versus edge deployment, general-purpose versus specialized control systems, and regional infrastructure constraints.

The Ambiguous Net Impact

Scenario analyses suggest that data centers' electricity demand could grow markedly this decade.[2] Yet the same digital technologies can also curb energy consumption by optimizing energy-intensive processes across industry, buildings, transport, and agriculture. The key question is what is the net energy impact of using AI in specific real-world settings? Are the savings of the use of AI greater than the energy costs of running AI, and under what conditions? Digital twins (virtual replicas fed by live data) offer a promising route for increased efficiency,[3] enabling closed-loop control, predictive maintenance, and counterfactual simulation enabling optimized processes without disrupting physical assets.

Why Digital Twins Matter

A digital twin continuously synchronizes with the physical process, learns about its dynamics, and tests "what-if" adjustments before implementation. In energy-intensive operations, the use of AI can help to process efficiency through a reduction in wasteful overshoot, increase flexibility by shifting consumption to off-peak or low-carbon hours while preserving output quality.

These levers are particularly relevant in wastewater treatment, where quality constraints are strict, the physical processes are complex, and energy use is concentrated in a few steps.

Using AI to Reduce Energy in the Water Sector: The Veolia Case

Context: Where the Energy Goes

In conventional wastewater treatment, secondary treatment (biological oxidation) is typically the energy hotspot. Continuous aeration supplies dissolved oxygen for microbial degradation of pollutants; blowers and compressors often represent roughly half of a plant's electricity bill.[4] Historically, control relies on rule-based mechanisms tied to thresholds to guarantee compliance under uncertainty. These rules are robust but conservative, often over-aerating to avoid breaching limits.

Figure 1: *Representation of a wastewater treatment plant with activated sludge*

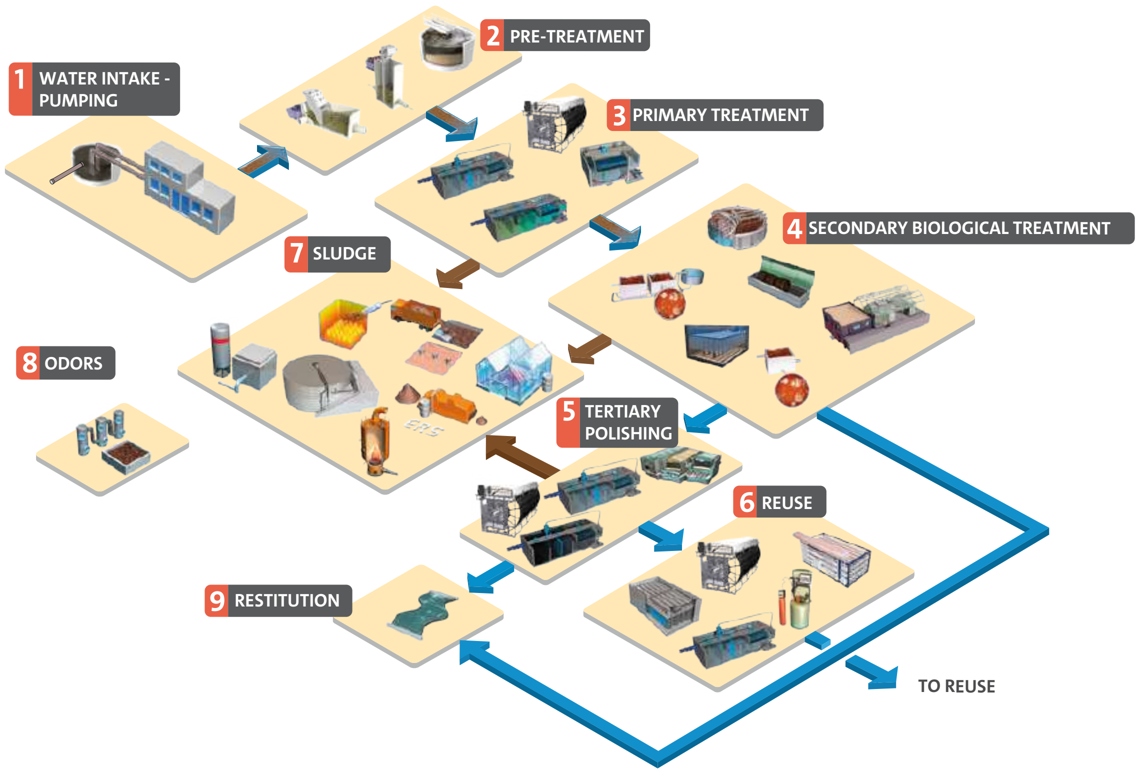

Veolia × PureControl: From Automation to Learning Control

Veolia, a global leader in water and wastewater services, partnered with PureControl to deploy AI-enabled digital twins across its portfolio. Nearly 300 Veolia plants operate with PureControl's device. The system aggregates high-frequency data (generally at quarter-of-hour intervals) from:

- **Internal sensors:** water flow rate, redox level,[5] aeration time, electricity power.
- **Laboratory samples:** regulatory quality metrics (e.g., nitrate, ammonium).
- **External signals:** weather forecasts and ambient conditions (snow, rain, humidity, temperature)
- **Electricity markets:** tariffs, time-of-use prices, and demand charges.

These data are integrated into several modelling layers using both traditional statistical methods to predict short-term outcomes (flow rate and pollution level) and deep learning neural networks to optimize the use of more advanced processes such as the aeration timing and intensity. AI is used to implement the aeration of the basin that minimizes cost and consumption subject to hard effluent-quality constraints and safety margins.

Figure 2: *Geographic distribution of Veolia facilities operating with PureControl's digital-twins*

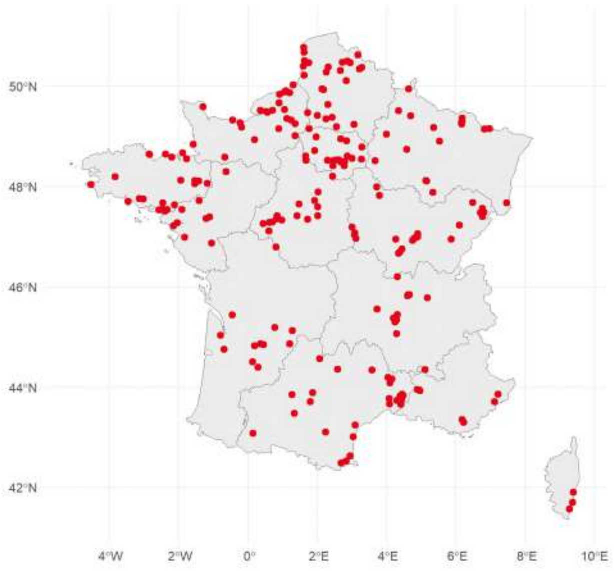

Operational Principle

The AI does not relax compliance, but it refines how compliance is achieved. In practice, that means avoiding systematic oversupply of oxygen, shifting aeration toward off-peak hours when feasible, without jeopardizing the quality of effluents and an adaptive response to real-time data inflows. Yet the human remains in the loop, and operators can retain authority when necessary.

Assessing the Impact: Methods and Early Evidence

Data Architecture and Granularity

The study uses plant-level telemetry, lab results, and opensource weather data assembled by PureControl. Variables include granular electricity draw, aeration blower runtimes and set-points, dissolved oxygen and nutrient concentrations, influent characteristics, ambient conditions, and time-of-use pricing. This resolution allows plants to track both levels (kWh, kg CO_2e) and profiles (load shifting across the day) of consumption.

Figure 3: *Profile of the energy consumption, median aeration time and AI use (in proportion of time) around the implementation of AI for a given plant*

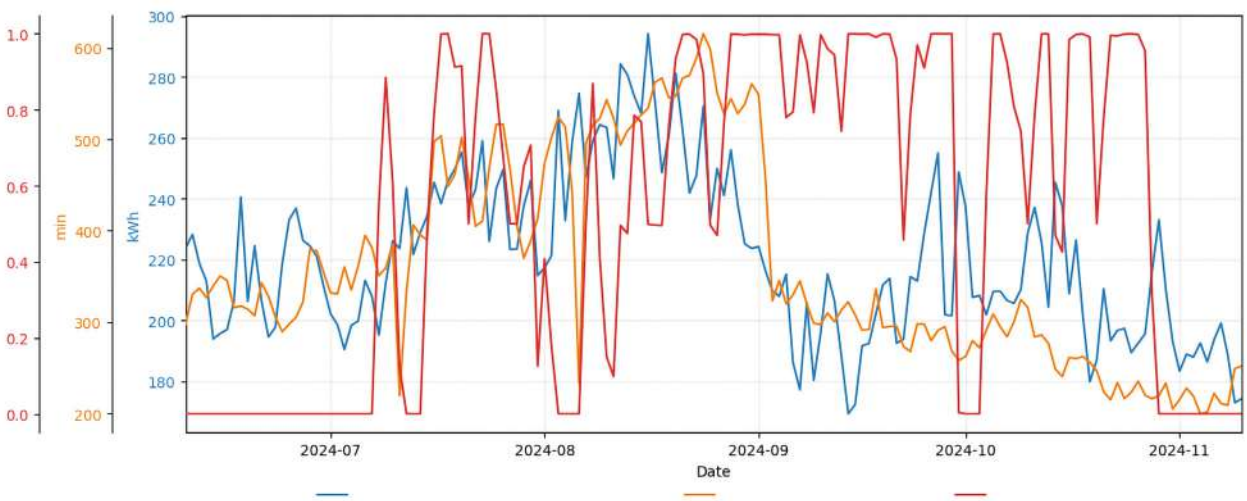

On September 15, the AI was active 36% of time. The plant consumed 1108 kWh of electricity and the aeration system operated for 551 minutes.

Empirical Strategy Agenda to Measure AI Impact

Our approach uses complementary strategies to isolate AI's effect from potential confounding factors.

First, AI implementation and failures offer a natural experimental setup to identify the causal effect of AI on electricity consumption and isolate it from potential confounding factors. In particular, we exploit exogenous interruptions to AI-assisted control (that temporarily revert plants to their baseline behavior). Comparing adjacent windows before/after such shocks within the same plant yields a within-asset estimate of AI's causal effect on consumption and quality, controlling for seasonality and load differences across stations. Staggered adoption across plants is useful to run event-study models comparing treated to not-yet-treated plants, with granular fixed effects at the plant-level, for each calendar time, and weather bins. It captures persistent shifts attributable to AI, while the fixed effects absorb the common shocks that affect all units, preventing them from biasing the estimated impact.

Second, we can net out the AI layers' electricity use, by estimating the electricity consumption induced by AI servers. Net savings then equal gross plant-side reductions minus this overhead.

Finally, we turn to the Profile analysis to disentangle the effect on the plant energy use from the reallocation of energy consumption over time. It allows testing to determine whether AI reallocates consumption toward off-peak hours without harming effluent quality, using price-exposure interactions and intraday load curves.

Preliminary Results

In an initial subsample, we observe, pending further analysis of this data:

- **Energy and emissions:** a nearly 10% average reduction in electricity consumption and its induced level of GHG emissions on days when AI is used at full extent compared to days when

AI does not control the process. This result has only been tested under restrictive assumptions so far and its robustness will be extended in a more detailed analysis in later stages of this project.

- **AI overhead:** the PureControl layer's own electricity draw is estimated to represent less than 1% of the gross energy savings, leaving a strong net reduction. Even in the most conservative cases, in the rare cases when hardware adjustments are required onsite to adopt AI locally, induced lifecycle GHG emissions never exceed 30-45% of the GHG emissions savings.

- **Load shifting:** evidence of partial reallocation toward off-peak tariffs, with potential to enhance grid friendliness as price signals strengthen.

We further observe significant variations in AI's impact on energy consumption and GHG emissions depending on the duration of AI usage on a given day. Specifically, when AI is used over a longer time interval within a daily operational cycle, its ability to reduce both energy consumption and greenhouse gas emissions increases. This effect likely arises from the AI's capacity to optimize plant operations over extended periods, allowing it to adjust aeration cycles more effectively and shift energy consumption to off-peak times. Longer exposure to AI-driven control provides the system with more data and time to fine-tune operational parameters, leading to more significant and sustained reductions in energy use. Conversely, on days when AI is applied for shorter durations, the potential for such optimizations is limited, and the resulting savings in energy and emissions are less substantial.

***Figure 4:** Mean daily consumption and GHG emissions depending on the share AI-control time during the day*

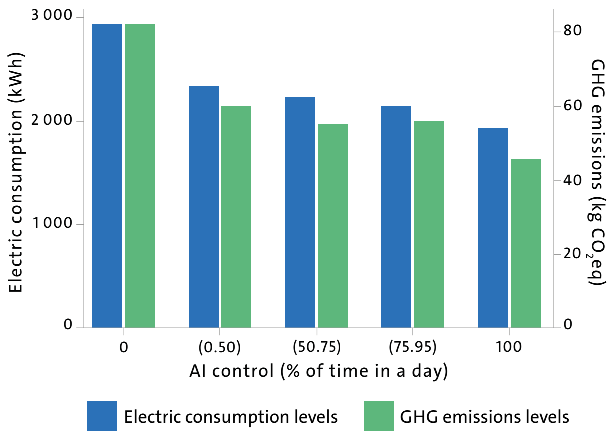

Scaling and Governance

To scale AI across multiple plants effectively, several key prerequisites must be met:

1. **Data infrastructure:** There must be a reliable system in place for collecting and sharing data. The larger data the AI gets, the better the quality of the AI predictions and optimizations. It requires proper connectivity protocols to ensure the AI can access the data it needs in real-time.

2. **Access and control:** For AI to manage plant operations, it needs the ability to interact with the system's controls. This means that operators must allow the AI to adjust settings on machines or automate specific processes.

3. **Up-to-date systems:** The technology running the AI needs to be kept current and well-maintained. This includes ensuring that all devices are compatible with the AI system. Updating hardware helps the system stay efficient and ensures it can handle more data as the plant's operations grow.

These foundational elements ensure that AI can be integrated effectively into plant operations and scaled successfully, while maintaining control and security.

Conclusion

AI's energy story is not one-sided. In wastewater treatment, AI-driven digital twins can deliver verifiable reductions in electricity use and GHG emissions precisely where consumption is concentrated – secondary aeration – without compromising regulatory performance. In Veolia's partnership with PureControl, early evidence across a growing set of facilities points to a nearly 10% energy and GHG cuts, negligible operational overhead from the AI layer, and promising load-shifting benefits.

As the research project broadens, rigorous, portfolio-wide evaluation will refine these estimates. For utilities and policymakers, the message is clear: targeted, specialized AI – embedded in robust governance – can be a practical lever for decarbonization today.

About the authors

Mathias Abitbol is a graduate student from the Ecole Normale Supérieure and HEC Paris, and affiliated to the Innovation Lab at Collège de France.

Philippe Aghion is a Professor at Collège de France, INSEAD, and the London School of Economics. In 2025, he was awarded the Nobel Prize in Economics.

Céline Antonin is a researcher and an affiliate professor at Sciences Po (OFCE) and an associate researcher at the Innovation Lab at Collège de France.

Lint Barrage is an Associate Professor at ETH Zurich.