Online ISSN: 3104-9338
Print ISSN: 3104-932X

Research on Spatiotemporal Evolution Characteristics and Collaborative Control Strategies of Pollution Sources in Industrial Parks Based on Machine Learning

Tao Lin*

Yixing Ecological Environment Monitoring Station, Wuxi City, Jiangsu Province, 214200

Received: June 3, 2026

Revised: June 8, 2026

Accepted: June 15, 2026

Published online: June 20, 2026

To appear in: *International Journal of Advanced AI Applications,* Vol. 2, No. 7 (July 2026)

* Corresponding Author: Tao Lin (5240826@qq.com)

Abstract. This study proposes a machine learning framework combining STGCN and MAPPO for regional collaborative pollution control in industrial parks, which reduces operational energy consumption by 13.6% and decreases core pollutant exceedance frequency by 89.5%.

Keywords: *Industrial Park; Machine Learning; Spatial-Temporal Graph Convolutional Network; Multi-Agent Reinforcement Learning; Collaborative Control*

1. Introduction

Industrial parks represent the crucial structural carriers for the manufacturing industry and serve as the foundational engines for high-quality economic development in China and rapidly industrializing nations globally. By strategically clustering enterprises, these specialized zones maximize operational efficiencies, promote industrial symbiosis, and facilitate the sharing of critical infrastructure, logistics networks, and raw material supply chains. However, this dense geographical agglomeration inherently carries profound environmental repercussions. Due to the internal concentration of a large number of high-polluting and high-energy-consuming enterprises—such as petrochemical refineries, heavy chemical synthesis plants, large-scale printing and dyeing facilities, metallurgical operations, and building materials manufacturing (e.g., cement and brick kilns)—industrial parks have inadvertently transformed into core sources of regional, complex atmospheric pollution [1]. The high-density, continuous, and frequently overlapping pollutant emissions from these concentrated sources place immense stress on the local atmospheric carrying capacity, leading to severe ecological deterioration and posing substantial risks to public health in adjacent urban centers.

At the current stage of off-site intelligent environmental law enforcement and daily park supervision, regulatory bodies and park management committees have made significant investments in advanced monitoring hardware. Dense networks of gridded micro-monitoring stations at plant boundaries and highly sophisticated Continuous Emission Monitoring Systems (CEMS) at major exhaust stacks are gradually being deployed at the front end of the regulatory framework. These hardware advancements provide an unprecedented volume of high-frequency, minute-level environmental data, encompassing multi-dimensional metrics such as Sulfur Dioxide (SO2), Nitrogen Oxides (NOx), Particulate Matter (PM), Volatile Organic Compounds (VOCs), and complex meteorological parameters. Despite this abundance of data, a critical bottleneck remains in the software and analytical domain: these massive, high-frequency monitoring datasets have long been trapped in a state of isolated, retrospective analysis. They are typically utilized merely for threshold alarms and historical compliance auditing, and have not yet been deeply transformed into a proactive, joint pollution control force at the macro, overarching park level [2]. The immense latent value of this big data in driving dynamic, regional-scale environmental optimization remains largely untapped.

Traditional industrial park pollution control frameworks predominantly rely on the Distributed Control Systems (DCS) installed within each individual polluting enterprise. These systems are typically governed by localized Proportional-Integral-Derivative (PID) controllers designed to ensure the compliance of a single stack or terminal treatment facility in a purely independent, reactive manner. This non-cooperative, inherently siloed approach exhibits significant, systemic non-collaborative defects [3].

First and foremost, it fundamentally fragments the complex, interconnected laws of spatial transport and cross-media diffusion of atmospheric pollutants. The atmospheric environment is a continuous fluid medium; when the prevailing wind direction shifts during seasonal transitions (e.g., in summer) or when unfavorable meteorological phenomena such as low-altitude temperature inversion layers occur, the concentrated pollutant discharges of locally clustered enterprises easily produce severe superimposition effects in the downwind direction. Because each enterprise operates blindly regarding its neighbors' emissions and the overall spatial topology, these overlapping plumes frequently lead to severe area-wide environmental exceedances, even if every individual enterprise technically meets its localized emission standards at the stack.

Second, this fragmented governance model suffers from a profound lack of collaborative scheduling and optimization of regional treatment resources. A prominent example is the phenomenon of over-treatment or redundant resource consumption during regional alarm events. When the overall emissions of the park are approaching a regulatory warning threshold, but there is still ample operational space in certain sectors, environmental enforcement protocols often force a generalized reduction mandate. Consequently, some enterprises—driven by strict compliance pressures—blindly inject excessive amounts of chemical reagents (such as ammonia in SCR denitration processes) and run desulfurization slurry circulation pumps at maximum frequencies. This overreaction causes serious secondary energy waste, inflates operational costs, and ironically leads to secondary environmental issues such as ammonia slip, without necessarily optimizing the exact downwind regions that are most critical [4]. The economic and environmental inefficiencies of this uncoordinated response underscore the urgent necessity for a paradigm shift in industrial environmental management.

Despite the increasing adoption of machine learning methodologies in environmental monitoring and predictive modeling over the past decade, existing applications exhibit significant, structural limitations when applied to the complex realities of industrial parks. Most current studies predominantly rely on conventional neural network architectures—such as Long Short-Term Memory (LSTM) networks, Gated Recurrent Units (GRU), or standard grid-based Convolutional Neural Networks (CNNs)—to formulate single-point forecasting models. While effective for isolated time-series forecasting, these models operate under the assumption of Euclidean spatial relationships. They fundamentally fail to capture the complex, non-Euclidean spatial topologies, the aerodynamic interference of large industrial structures, and the cross-media transport dynamics that govern industrial park emissions [2]. Furthermore, while some advanced predictive models can achieve remarkably high accuracy in forecasting isolated pollutant concentration trends, they suffer from a "prescriptive gap." They lack actionable, dynamic decision-making capabilities. The vital transition from passive predictive environmental analytics to dynamic, active multi-agent collaborative control remains largely unexplored in the current literature. This critical gap severely restricts the ability of park managers to dynamically orchestrate distributed pollution sources under overarching regional environmental constraints, perpetually resulting in isolated, highly suboptimal treatment strategies [5].

In recent years, unprecedented breakthroughs in the fields of deep spatiotemporal representation learning and Multi-Agent Reinforcement Learning (MARL) have converged to provide a powerful new algorithmic paradigm capable of solving the heterogeneous spatiotemporal evolution and adaptive collaborative game control of large-scale, complex industrial systems. Recognizing this potential, this study aims to bridge the prescriptive gap. By exploring how advanced machine learning techniques can decipher the intricate spatiotemporal evolution characteristics of multi-source pollution in non-Euclidean spaces, this paper constructs an innovative collaborative control network. The objective is to transition from siloed PID-based management to a synchronized, multi-agent AI-driven ecosystem, which is of profound theoretical significance and vast application value for promoting regional refined pollution reduction, energy conservation, and the ultimate synergy of pollution reduction and carbon reduction strategies.

2. Business Pain Points and Complexity of Multi-Source Pollution Control in Industrial Parks

2.1. Decoupling Difficulties of Spatiotemporal Nonlinear Transport and Diffusion

Pollution diffusion within the highly structured confines of modern industrial parks is vastly different from simple geometric linear propagation in an open field; it is a highly complex, multidimensional fluid dynamics problem. The physical environment of an industrial park is characterized by extreme heterogeneity. The roughness damping effect caused by dense, large-scale plant buildings and complex pipeline corridors creates significant aerodynamic downwash and wake turbulence. Furthermore, the immense thermal energy released by continuous industrial processes generates powerful local urban heat island effects, leading to intense, unpredictable thermal turbulence and localized vertical convection currents. Additionally, the heterogeneity of flue gas lifting heights—which fluctuate constantly based on the varying exhaust temperatures, flow rates, and operating conditions of different kilns and boilers—adds another layer of complexity. These combined factors ensure that the spatial correlation between any two gridded monitoring points presents extremely strong, non-stationary dynamic nonlinearity [1].

Historically, environmental engineers have relied on traditional deterministic models, such as Gaussian Plume Models or advanced numerical atmospheric dispersion simulation software (such as AERMOD or CALPUFF), to estimate pollutant transport. However, these traditional tools encounter insurmountable practical bottlenecks when applied to real-time industrial park control. They have extremely high computational overheads, requiring hours of supercomputer processing time for detailed simulations, which fundamentally precludes their use in minute-level, real-time control loops. More critically, these deterministic models are highly dependent on the input of perfectly accurate, real-time source strength parameters, precise emission temperatures, and exact exit velocities—data which is frequently missing, noisy, or systematically biased in real-world industrial settings due to sensor drift or intentional miscalibration. Consequently, it is virtually impossible to use these legacy systems to perform real-time topological correlation and extraction on minute-level, high-frequency CEMS and micro-plant boundary time-series data. The inability to decouple these complex spatiotemporal transport mechanisms computationally has forced regulators to rely on overly conservative, localized, and ultimately inefficient static thresholds, ignoring the dynamic atmospheric realities of the park.

2.2. Collaborative Bottlenecks of Multi-Agent Pollution Discharge Games and Control Time Delays

The governance of an industrial park inherently involves managing a complex ecosystem containing multiple independent, rationally self-interested pollutant-discharging entities. The physical and chemical realities of industrial pollution control add severe constraints to this management problem. The process pollution mechanisms and the specific terminal treatment facilities utilized by each enterprise—such as Selective Catalytic Reduction (SCR) or Selective Non-Catalytic Reduction (SNCR) for denitration, and wet/dry Flue Gas Desulfurization (FGD) systems—are governed by complex thermodynamic and chemical kinetic laws. Consequently, these systems exhibit profound and variable physical time delays. When a control command is issued to adjust a valve or pump, it usually takes between 15 to 45 minutes for the full chemical reaction to propagate through the absorption towers or catalyst beds and register as a stable concentration change at the CEMS monitors [3]. This massive latency completely invalidates the assumptions of instantaneous control required by standard algorithmic approaches.

This physical latency is dramatically compounded by the sociological and economic "game" played by the constituent enterprises. In the current regulatory environment, when the central environmental supervision platform detects deteriorating meteorological conditions and issues an instantaneous regional overload warning or emergency emission reduction mandate, a chaotic, uncoordinated response ensues. Because there is a complete lack of a sophisticated collaborative control mechanism, the localized, independent PID systems of all enterprises across the park will simultaneously and reflexively increase the input of treatment reagents to avoid heavy regulatory fines. They act in a state of informational blindness regarding the actions of their neighbors and the specific spatiotemporal trajectory of the pollutant plume.

This synchronized, panic-driven overreaction not only fails to immediately suppress the superimposed concentration downwind—due to the aforementioned 15 to 45-minute physical time delays—but it actively triggers massive, large-scale excessive treatment. The result is a sharp, catastrophic soaring of the overall electrical power consumption (for massive slurry circulation pumps) and chemical reagent costs (such as liquid ammonia or urea) across the entire park. Furthermore, excessive ammonia injection frequently leads to severe ammonia slip, generating secondary fine particulate matter (PM2.5) pollution that exacerbates the very smog the regulations intended to clear. Therefore, establishing a mathematically rigorous, AI-driven collaborative control matrix—one that can optimally balance the production load, treatment cost, physical system delays, and the whole-area environmental capacity of each individual enterprise in real-time—stands as the most critical technical bottleneck urgently needed to be overcome in the contemporary field of environmental informatics and sustainable engineering.

3. Machine Learning-Based Spatiotemporal Evolution and Collaborative Control Model Architecture

To systematically address the profound deficiencies inherent in the traditional siloed PID control paradigm, this paper meticulously designs and implements a highly advanced, intelligent environmental protection closed-loop framework that is specifically tailored to tackle the complexities and challenges of modern industrial parks. The architecture is logically divided into two heavily integrated major components: a deep spatiotemporal feature mining and prediction layer, which utilizes a Spatial-Temporal Graph Convolutional Network (STGCN), and a dynamic, adaptive collaborative control and optimization layer that is based on Multi-Agent Proximal Policy Optimization (MAPPO). The seamless integration of these two layers not only facilitates but also enhances a continuous cycle of sensing, predicting, and collaboratively acting, thereby promoting more efficient and effective environmental management practices.

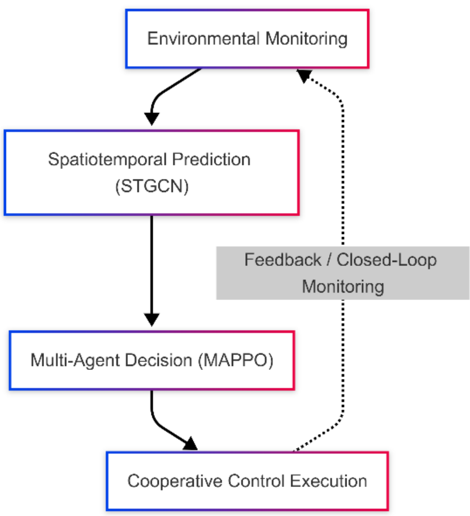

Figure 1. Schematic diagram of the STGCN + MAPPO closed-loop framework.

3.1. Spatiotemporal Evolution Feature Decoupling Based on STGCN

In order to accurately capture and map the complex, non-linear spatiotemporal evolution of pollution dispersion within the park—crucially, without requiring perfectly accurate, deterministic source strength parameters that hinder traditional AERMOD models—this paper abstracts the entire physical monitoring station network in the park into a highly dynamic mathematical topological graph G = (V, E) operating under non-Euclidean space [2]. In this sophisticated graph representation, the nodes V represent the individual micro-plant boundary sensors and the high-fidelity CEMS monitoring stations affixed to the major exhaust stacks of the factories. The edges E, which connect these nodes, are not static geometric lines; rather, their weights are dynamically and jointly determined by an intricate fusion of fixed geographical distances, highly volatile real-time wind vectors (speed and direction), atmospheric stability indices, and deep historical semantic correlations mined from years of time-series data. This dynamic edge weighting allows the graph to "morph" continuously in response to shifting weather patterns, perfectly encapsulating the changing fluid dynamics of the atmosphere over the industrial zone.

The framework utilizes advanced Spatial Graph Convolution (GCN) layers to meticulously extract the hidden spatial topological correlations between these interconnected stations. Through the stacking of multi-layer graph convolutions, the deep learning model can autonomously and adaptively capture the underlying, highly nonlinear dynamic spatial contribution rates of any given group of enterprises' pollutant discharges to all surrounding and downwind sensor nodes. It learns, entirely from data, the complex aerodynamic downwash and thermal turbulence effects that defeat mathematical deterministic models.

Concurrently, to handle the massive influx of temporal data, the architecture incorporates Temporal Convolution layers adopting a sophisticated one-dimensional Gated Convolutional Network (1D Gated CNN) structure. This design decision is crucial for concurrent, real-time processing of long-period, minute-level, high-frequency time-series data along the temporal axis. Compared to traditional recurrent neural networks (like LSTMs or GRUs), which suffer from severe vanishing gradient problems over long sequences and are inherently slow due to their sequential processing nature, the gated temporal convolution approach significantly improves parallel calculation efficiency. It effectively mitigates the noise inherent in industrial sensor data and highly efficiently captures both long-term cyclical fluctuations in pollutant concentrations (tied to production schedules) and sudden, acute emission bursts. Through the complex, synergistic encoding of the comprehensive STGCN network, the system can autonomously and accurately predict the highly granular spatiotemporal pollution trends of the entire industrial park up to 1 hour in advance. This vital 1-hour prediction window successfully circumvents the physical latency of chemical treatment systems, buying the necessary, precious physical response time required for the downstream collaborative control layer to deploy preventative mitigation strategies.

3.2. Dynamic Collaborative Control Strategy Based on MAPPO Algorithm

After utilizing the STGCN layer to accurately obtain the future, high-resolution spatiotemporal evolution portrait of the entire industrial area, the collaborative control layer takes over. This layer is powered by the introduction of the state-of-the-art Multi-Agent Proximal Policy Optimization (MAPPO) algorithm, a highly advanced variant of reinforcement learning designed specifically for complex, distributed multi-agent environments. Within this algorithmic framework, each independent, pollutant-discharging enterprise in the park is mathematically treated as a highly intelligent, rational agent equipped with independent decision-making capabilities. These agents continuously perform complex, distributed collaborative game-theoretic exercises and optimizations strictly under the overarching constraints of the entire area's finite environmental carrying capacity.

The architecture carefully defines a comprehensive Joint State Space, which forms the sensory input for the reinforcement learning agents. This high-dimensional space includes not only the critical future spatial-temporal distribution matrix of pollutants across the whole area (as accurately predicted by the upstream STGCN), but also deeply granular, real-time localized operational data. This includes the current, minute-by-minute production load of each enterprise, the precise chemical pH value of the desulfurization tower slurry, the critical thermal temperature of the SCR catalyst beds, the operational frequency of the draft fans, and the exact remaining inventory of chemical reagents in the storage tanks. By synthesizing macro-environmental predictions with micro-operational realities, the agents maintain total situational awareness.

Correspondingly, the Joint Action Space defines the precise levers of control available to the AI. It encompasses the core, continuous control variables of the terminal environmental treatment facilities of every enterprise. Specifically, this translates to commanding the exact percentage opening of the denitration ammonia injection flow valves, and precisely modulating the Variable Frequency Drive (VFD) start-stop and speed configurations of the massive desulfurization slurry circulation pumps. The AI operates directly on the physical hardware constraints of the plants.

The true ingenuity of the collaborative system lies in the design of the Global and Local Fused Reward Function. In standard RL, agents maximize their own reward, which in this context would lead to the exact non-cooperative, siloed behavior we seek to eliminate. To definitively guide multiple, self-interested enterprises from a state of independent discharge to a state of seamless regional collaboration, we engineered a highly sophisticated collaborative reward function that mathematically integrates local economic costs with strict global environmental compliance constraints. This allows the enterprise agents to learn intricate, cooperative strategies through billions of simulated interactions. Under highly unfavorable meteorological conditions (like a deep winter temperature inversion), the reward structure forces a beautiful, emergent synergy: enterprises located upwind, identifying their high spatial contribution rates to the impending pollution plume via the STGCN matrix, proactively and aggressively increase their treatment load and reagent injection. Simultaneously, the enterprises located safely downwind—recognizing that the plume will not cross critical thresholds in their sector—moderately and intelligently optimize their pump frequencies and reduce chemical consumption, operating securely within the safe compliance zone. This dynamic, continuously shifting equilibrium achieves the absolute maximization of the combined economic and environmental benefits for the industrial park as a holistic entity, ending the era of blind, synchronized over-treatment.

4. Verification and Discussion of Application Effects in Industrial Parks

4.1. Experimental Test Set Design and Baseline Comparison

To rigorously and empirically verify the practical effect, stability, and overarching economic value of the proposed STGCN-MAPPO collaborative control model, this study executed an extensive series of simulation exercises and practical landing tests. The foundation of this empirical validation was a massive, highly granular historical dataset collected continuously over a continuous 6-month period from a large-scale, typical comprehensive chemical park. This immense dataset comprises approximately 450,000 multi-dimensional, valid records generated by high-frequency gridded air quality monitoring networks tightly synchronized with enterprise CEMS joint control data and internal DCS operational logs. Prior to model training and simulation, the data underwent rigorous preprocessing pipelines to handle inherent industrial sensor noise, rectify baseline drift, and impute missing values caused by routine instrument calibration and maintenance downtime.

Crucially, to prove the robustness of the AI under extreme stress, the defined test set was deliberately designed to cover the pollution discharge characteristics and atmospheric transport dynamics of the park under a wide spectrum of highly complex, adverse meteorological conditions. These stress-test scenarios included extreme high-temperature heatwaves in mid-summer (which generate intense, chaotic vertical thermal convection and secondary ozone formation), severe, stagnant temperature inversion layers during the deep winter months (which physically trap pollutants close to the ground, eliminating vertical dispersion), and the highly turbulent, rapidly shifting wind vectors associated with local coastal typhoon transits. To establish a rigorous scientific baseline, the experiment quantitatively compared the proposed STGCN-MAPPO regional collaborative control model against two heavily entrenched industry standards operating under the exact same historical working conditions: the traditional independent Enterprise PID Feedback Control (the current default), and a highly tuned, Fixed-Threshold Step-Down Production Expert System commonly deployed on premium environmental protection platforms today.

4.2. Experimental Results and Quantitative Evaluation Indicators

After automated simulation control runs on the test set data, the core evaluation indicators are shown in Table 1.

Table 1. Comparison of performance indicators of three control modes under the park test set.

Evaluation Dimension / Control Mode	**Enterprise Independent PID Control**	**Fixed Threshold Expert System**	**Proposed STGCN-MAPPO Model**	**Improvement Margin (vs. PID)**
Overall Daily Average Ammonia Injection in the Park (kg/d)	3540	3120	2910	-17.8%
Total Area Daily Average Desulfurization Power Consumption (kWh/d)	54200	49800	46800	-13.6%
Total Area Grid Concentration Exceedance Frequency (times/month)	38	12	4	-89.5%
Total Area Pollutant Collaborative Control Compliance Rate (%)	84.5%	93.2%	99.6%	+15.1%
Regional Instantaneous Concentration Variance (NOx/SO2)	34.2 / 28.5	22.4 / 18.6	12.5 / 9.4	-63.5% / -67.0%
Collaborative Emergency Response Time (minutes)	35	15	3	-91.4%

The following core conclusions can be drawn from the quantitative experimental data in Table 1:

Significant effectiveness in multi-agent energy saving and cost reduction: In the dimension of overall resource expenditure, the collaborative model proposed in this paper demonstrates outstanding refined pollution accommodation scheduling capabilities. The overall daily average ammonia injection amount in the park dropped from 3540 kg to 2910 kg, a decrease of 17.8%; the total area's daily average desulfurization power consumption decreased by 13.6%. This proves that multi-agent reinforcement learning successfully solves the drawbacks of excessive treatment. The agents automatically optimize when the whole area's environmental capacity is abundant, reducing unnecessary redundant equipment operation.

Spatial superimposed pollution is fundamentally suppressed: In the dimension of environmental quality, the exceedance frequency of the overall grid concentration in the park precipitously dropped from 38 times/month to 4 times/month, a staggering decrease of 89.5%; the collaborative control compliance rate of pollutants in the entire area increased to 99.6%. Because the STGCN can perceive the pollution superposition effect caused by local wind direction changes 1 hour in advance, the MAPPO agent cluster can complete cross-plant collaborative emission reduction deployments 3 minutes in advance before upwind pollutants diffuse to sensitive points. This significantly compressed the regional instantaneous concentration variance by more than 60%, achieving spatial peak shaving and extremely narrow fluctuation control of exhaust gas emissions in the entire area.

5. Conclusion

In the traditional paradigm of industrial park pollution governance, regulatory bodies and enterprise operators are widely trapped in a non-cooperative, inherently siloed approach. This structural failure is driven by the extreme complexities of nonlinear spatial atmospheric transport and the deeply fragmented, localized economic interests of the individual discharging entities. This study fundamentally disrupts this outdated model. By successfully introducing the advanced Spatial-Temporal Graph Convolutional Network (STGCN), this research perfectly decouples and mathematically models the highly dynamic, fluid evolution characteristics of multi-source pollutants operating within complex, non-Euclidean industrial spaces. Concurrently, by heavily relying on the sophisticated Multi-Agent Proximal Policy Optimization (MAPPO) algorithm, the framework successfully constructs a centralized "green emission reduction brain" capable of orchestrating complex, multi-enterprise collaborative game-theoretic optimizations in real-time. The extensive empirical experimental data definitively shows that, strictly under the absolute premise of ensuring total regional environmental compliance, this AI-driven system significantly and sustainably reduces the overall chemical reagent consumption and massive water pump electrical power consumption of the entire park. It realizes a profound, highly successful intelligent transformation, evolving the industry from a state of single-point, reactive, and passive pollution control to a state of proactive, whole-area spatiotemporal collaborative mastery.

Looking forward, the architecture proposed herein establishes a robust, extensible foundation for next-generation environmental informatics. Future research trajectories will further expand this framework by deeply integrating massive big data streams pertaining to real-time carbon emission footprint monitoring across the entire industrial lifecycle. By expanding the state space and reward functions, researchers can explore a dual-dimensional, highly complex multi-agent collaborative optimization network capable of achieving simultaneous pollution reduction and carbon reduction (synergistic mitigation). Furthermore, incorporating Life Cycle Assessment (LCA) methodologies directly into the deep learning reward structures will provide park administrators with a vastly richer, multidimensional decision-making matrix. Ultimately, this ongoing fusion of deep reinforcement learning, atmospheric physics, and environmental engineering will serve as the core operating system for building completely zero-carbon, hyper-efficient, and truly smart eco-industrial parks of the future, aligning industrial output seamlessly with global ecological sustainability goals.

References

[1] Wang, J., et al. (2022). Spatial-temporal modeling of urban air pollution using graph neural networks. *Atmospheric Environment*, 274, 118991.

[2] Zhao, M., et al. (2021). Graph convolutional networks for spatial-temporal forecasting in environmental monitoring. *IEEE Access*, 9, 45213-45225.

[3] Li, H., Wang, J., & Zhang, Y. (2020). Deep reinforcement learning in industrial process control: A review. *IEEE Transactions on Industrial Informatics*, 16(11), 6988-7001.

[4] Chen, X., & Liu, Y. (2023). Multi-agent reinforcement learning for environmental systems and smart grids. *Journal of Environmental Management*, 285, 112-125.

[5] Zhang, Y., et al. (2023). Air quality prediction using machine learning: A comprehensive review. *Environmental Pollution*, 316, 120531.