

ECO-TrainNet: An Integrated Framework for Carbon Emission Prediction and Energy-Efficient Control in Urban Rail Transit

Caihao Xu^{1,*}, Jin Han¹, Chenjie Li¹

¹North China University of Science and Technology, Tangshan, China

*Corresponding author: jjasper123@163.com

Abstract

Urban rail transit systems are a key component of sustainable urban mobility, yet their increasing operational intensity leads to growing energy consumption and associated carbon emissions. This study proposes ECO-TrainNet, an integrated framework for carbon emission prediction and energy-efficient control in urban rail transit. The framework combines machine learning-based prediction with reinforcement learning-based control in a closed-loop structure. A simulation-based dataset is constructed to capture train dynamics, infrastructure characteristics, and time-varying carbon emission factors. Ensemble learning models, particularly XGBoost, are employed to estimate energy consumption and carbon emissions, while a reinforcement learning agent is trained to optimize acceleration strategies under operational constraints. Experimental results show that the proposed framework achieves high prediction accuracy and enables effective control optimization. Compared with baseline operation, the reinforcement learning-based strategy reduces energy consumption by 11.7% and carbon emissions by 11.0%, while maintaining acceptable schedule deviation. The results demonstrate that integrating prediction and control provides a practical approach for carbon-aware and energy-efficient rail transit operation.

Keywords

Urban rail transit; Energy consumption; Carbon emission; Machine learning; Reinforcement learning; Optimization.

1. INTRODUCTION

Urban rail transit systems are widely recognized as an energy-efficient mode of transportation due to their high passenger capacity and relatively low per-capita emissions[2,10]. However, the increasing operational intensity of metro systems, characterized by high-frequency services and complex network structures, has led to a substantial growth in total electricity consumption[2]. Under the context of carbon neutrality policies, it is no longer sufficient to evaluate system performance solely in terms of energy efficiency; instead, explicit quantification and optimization of carbon emissions have become necessary.

Existing studies on metro operation optimization have primarily focused on reducing traction energy consumption through timetable adjustment, speed profile optimization, and regenerative braking utilization[1-4]. Recently, reinforcement learning has been introduced to train control policies under uncertain disturbances, enabling adaptive and data-driven operation strategies[5,8]. Despite these advancements, two critical limitations remain.

First, most studies treat energy consumption as the primary objective, while carbon emissions are either approximated using static emission factors or neglected entirely. Such simplifications fail to capture the temporal variability of electricity carbon intensity and limit

the applicability of these methods in carbon-aware decision-making scenarios. Second, optimization frameworks are typically constructed without an explicit predictive layer. In practical operations, decision-making relies not only on current system states but also on anticipated energy consumption and emission outcomes. The absence of predictive modeling reduces the interpretability and robustness of optimization strategies.

To address these limitations, this study proposes a data-driven framework, termed ECO-TrainNet, for integrated carbon emission prediction and energy-efficient scheduling optimization in urban rail transit systems. The proposed framework establishes a closed-loop structure that links operating-condition recognition, machine-learning-based prediction, and reinforcement-learning-based control. A simulation-based dataset is constructed to represent train operation dynamics, including speed profiles, acceleration patterns, section characteristics, passenger load, and time-varying carbon emission factors. Based on this dataset, ensemble learning models are trained to estimate section-level energy consumption and carbon emissions, while a reinforcement learning agent is designed to optimize acceleration decisions under operational constraints such as punctuality and ride comfort.

Compared with existing studies, the contributions of this work can be summarized as follows. First, a unified framework is developed to integrate prediction and control, enabling carbon-aware decision-making in metro operations. Second, a reproducible simulation environment is constructed to support systematic evaluation of different control strategies. Third, the effectiveness of the proposed method is validated through comparative experiments, demonstrating that optimized control policies can achieve simultaneous reductions in energy consumption and carbon emissions under realistic operational constraints.

2. LITERATURE REVIEW

Urban rail transit energy optimization has been extensively studied, with existing research generally falling into three main categories: physics-based modeling, data-driven prediction, and optimization-based control[2,4].

Physics-based approaches typically rely on traction equations and resistance models to estimate train energy consumption as a function of speed profiles, gradients, and rolling resistance[1,2]. These models are grounded in well-established physical principles and offer strong interpretability. However, they often require precise parameter calibration and may lack flexibility when applied to complex and dynamic urban rail systems. In particular, variations in passenger load, driving behavior, and operational disturbances can introduce nonlinear effects that are difficult to capture using purely deterministic formulations.

With the increasing availability of operational data, data-driven methods have been widely adopted to model energy consumption in transportation systems. Machine learning algorithms, such as random forests, gradient boosting methods, and neural networks, have demonstrated strong capability in capturing nonlinear relationships between input variables and energy usage[6,7]. These approaches are particularly effective when dealing with heterogeneous data sources and complex interactions among operational factors. Nevertheless, most existing studies focus primarily on energy consumption prediction, while carbon emissions are typically derived using static or simplified emission factors. Such treatments fail to account for temporal variations in electricity carbon intensity and limit the applicability of these models in carbon-aware operational scenarios.

In parallel, optimization-based approaches have been developed to improve operational efficiency, including timetable optimization, speed profile design, and disturbance management. Among these methods, reinforcement learning has gained increasing attention due to its ability to learn adaptive control policies through interaction with dynamic environments. Reinforcement learning-based approaches have shown promising results in reducing traction

energy consumption and improving system responsiveness[3,4]. However, these methods are commonly formulated with energy minimization as the sole objective, and carbon emissions are rarely incorporated explicitly into the optimization framework[5,8]. As a result, their effectiveness in supporting low-carbon decision-making remains limited.

A more fundamental limitation in the existing literature is the separation between prediction and control. In many studies, energy estimation models and optimization strategies are developed independently, leading to a lack of coordination between forecasting and decision-making processes. Consequently, optimization algorithms may not fully utilize predictive information, particularly in scenarios where future energy consumption and emission outcomes play a critical role in decision-making.

In addition, most current approaches focus on either system-level average performance or single-objective optimization, without considering the spatial heterogeneity of emissions across different track sections. As indicated by recent studies, energy consumption and carbon emissions can vary significantly depending on local operational conditions, suggesting the need for more fine-grained and adaptive optimization strategies[9,10].

To address these limitations, this study proposes an integrated framework that combines machine learning-based energy and carbon emission prediction with reinforcement learning-based control optimization. By establishing a closed-loop structure that links prediction and decision-making, the proposed approach enables more accurate and adaptive carbon-aware operation. This unified framework provides a systematic solution for bridging the gap between data-driven prediction and real-time control in urban rail transit systems.

3. METHODOLOGY

3.1. Overall Framework

This study proposes an integrated analytical framework, termed ECO-TrainNet, for carbon-aware operation and control of urban rail transit systems. The framework is designed as a closed-loop structure that connects operating-condition recognition, energy and carbon emission prediction, and reinforcement-learning-based control.

The overall workflow consists of four main stages:

Data preprocessing and feature construction, where raw operational variables are transformed into structured inputs;

Operating-condition recognition, which classifies train states into distinct regimes such as acceleration, cruising, braking, and dwell;

Energy and carbon emission prediction, where machine learning models estimate section-level energy consumption and associated emissions;

Control optimization, where a reinforcement learning agent determines acceleration strategies under operational constraints.

Unlike conventional approaches that treat prediction and optimization separately, the proposed framework integrates these components into a unified pipeline, enabling prediction-informed decision-making.

From a system perspective, the proposed framework can be formulated as a sequential decision pipeline that maps raw operational data X to optimized control actions a_t through an intermediate prediction layer. Specifically, the framework establishes a functional relationship:

$$a_t = \pi(f(X))$$

Where $f(X)$ represents the prediction model that estimates energy consumption and carbon emissions, and $\pi(\cdot)$ denotes the control policy learned by the reinforcement learning agent.

3.2. Data Simulation and Feature Design

Due to the limited accessibility of large-scale real-world metro operation datasets, a simulation-based dataset is constructed to support model development and evaluation.

The simulated dataset is designed to reflect realistic operational conditions and includes the following key variables:

Train dynamics: speed, acceleration, and travel time;

Section characteristics: segment length, gradient, and curvature;

Operational conditions: dwell time, passenger load factor;

Energy-related parameters: traction power, regenerative braking efficiency;

Environmental factors: time-dependent carbon emission factor of electricity.

To ensure physical plausibility, the data generation process follows basic kinematic relationships and operational constraints, such as speed limits, acceleration bounds, and station spacing. Each sample corresponds to a track section, and the dataset is structured at the section level to support both prediction and control tasks.

Feature engineering is applied to construct meaningful inputs, including:

Average speed per section;

Duration of different operating conditions;

Load-adjusted energy coefficients;

Gradient-related resistance indicators.

To ensure physical consistency, the simulated data satisfy basic kinematic relationships. For each section, the motion dynamics are constrained by:

$$\begin{aligned} v_{t+1} &= v_t + a_t \Delta t \\ s_{t+1} &= s_t + v_t \Delta t + \frac{1}{2} a_t (\Delta t)^2 \end{aligned}$$

Where v_t , a_t and s_t denote speed, acceleration, and position, respectively.

3.3. Energy and Carbon Emission Prediction

The prediction model aims to learn a nonlinear mapping between input features and energy consumption. Let $X \in \mathbb{R}^n$ denote the feature vector and $y \in \mathbb{R}$ denote the target variable. The prediction model can be expressed as an additive ensemble:

$$\hat{y} = \sum_{k=1}^K f_k(X), \quad f_k \in \mathcal{F}$$

where f_k represents the k -th regression tree and \mathcal{F} denotes the space of regression trees.

The model is trained by minimizing the following objective function:

$$L = \sum_i (y_i - \hat{y}_i)^2 + \sum_k \Omega(f_k)$$

Where $\Omega(f_k)$ is a regularization term used to control model complexity and prevent overfitting.

3.4. Reinforcement Learning-Based Control Optimization

The control problem is formulated as a Markov Decision Process (MDP), defined by a tuple (S, A, P, R, γ) , where S denotes the state space, A the action space, P the state transition probability, R the reward function, and γ the discount factor [8].

To further reduce energy consumption and carbon emissions during train operation, a reinforcement learning (RL) framework is introduced to optimize acceleration strategies.

The train operation process is modeled as a sequential decision-making problem. At each time step t , the agent observes the system state s_t , selects an action a_t , and receives a reward r_t .

(1) State Space

The state representation includes key operational variables:

Current speed;

Position along the track;

Remaining distance to the next station;

Accumulated energy consumption.

(2) Action Space

The action corresponds to the adjustment of train acceleration within predefined limits, ensuring compliance with safety and comfort constraints.

(3) Reward Function

The reward function is designed to balance multiple objectives:

$$r_t = -\alpha E_t - \beta C_t - \gamma D_t$$

Where:

E_t : energy consumption at time step t ;

C_t : estimated carbon emission;

D_t : deviation from the scheduled arrival time;

α, β, γ : weighting coefficients.

This formulation can be interpreted as a weighted multi-objective optimization problem, where energy efficiency, environmental impact, and operational reliability are jointly considered.

(4) Learning Algorithm

A policy-based reinforcement learning algorithm (e.g., PPO) is adopted to learn the optimal control strategy through interaction with the simulated environment. The training process iteratively updates the policy to maximize cumulative rewards.

3.5. Evaluation Metrics

These metrics are selected to evaluate both prediction accuracy and operational effectiveness, ensuring that the proposed framework is assessed from both modeling and control perspectives.

To assess the performance of the proposed framework, different evaluation metrics are used for prediction and control tasks.

(1) Prediction Performance

Mean Absolute Error (MAE)

Root Mean Squared Error (RMSE)

Coefficient of Determination (R^2)

(2) Control Performance

Total energy consumption (kWh)

Total carbon emissions (kg CO₂)

Schedule deviation (seconds)

Comparative experiments are conducted to evaluate:

Different prediction models;

Baseline vs optimized control strategies.

4. EXPERIMENTS AND RESULTS

4.1. Experimental Design and Simulated Dataset

To evaluate the effectiveness of the proposed ECO-TrainNet framework, a simulation-based experimental dataset was constructed at the track-section level. Each sample represents one train operation segment between two adjacent stations or control points. The dataset includes train dynamic variables, infrastructure characteristics, passenger load, regenerative braking efficiency, and electricity carbon emission factors.

The main variables include train speed, acceleration, section length, gradient, curvature radius, passenger load factor, dwell time, regenerative braking efficiency, energy consumption, and carbon emissions. The simulated values were generated under operationally reasonable constraints, including speed limits, acceleration bounds, and section distance ranges. This design aims to approximate typical urban rail operation conditions while ensuring that the experiment remains reproducible.

Table 1. Simulated Dataset Variables and Ranges

Variable	Description	Range
Speed	Train operating speed	0–80 km/h
Acceleration	Train acceleration/deceleration	-1.2–1.0 m/s ²
Section length	Distance of each track section	300–1500 m
Gradient	Track slope	-3%–3%
Passenger load factor	Relative passenger load	0.3–1.2
Regenerative braking efficiency	Energy recovery ratio	20%–40%
Carbon emission factor	Electricity emission intensity	0.45–0.75 kg CO ₂ /kWh

Table 1 shows that the simulated dataset covers both train operation variables and environmental variables. This is important because energy consumption is not only determined by speed and acceleration, but also affected by section length, gradient, passenger load, and braking recovery efficiency. The inclusion of carbon emission factors further enables the model to move beyond pure energy prediction and evaluate environmental impacts.

4.2. Prediction Model Evaluation

Three prediction models were compared: Linear Regression, Random Forest, and XGBoost. Linear Regression was used as a baseline model, while Random Forest and XGBoost were selected to capture nonlinear relationships between operational variables and energy consumption.

The models were evaluated using MAE, RMSE, and R². The results are shown in Table 2.

Table 2. Prediction Performance of Different Models

Model	MAE (kWh)	RMSE (kWh)	R ²
Linear Regression	1.82	2.35	0.87
Random Forest	1.21	1.68	0.93
XGBoost	0.98	1.42	0.95

Table 2 indicates that XGBoost achieves the best prediction performance among the three models, with the lowest MAE and RMSE and the highest R². The result suggests that energy consumption in urban rail systems is affected by nonlinear interactions among speed,

acceleration, gradient, passenger load, and regenerative braking efficiency. Linear Regression provides acceptable baseline performance, but its limited ability to model nonlinear relationships leads to larger errors. Random Forest improves the results, while XGBoost further enhances prediction accuracy due to its gradient boosting mechanism.

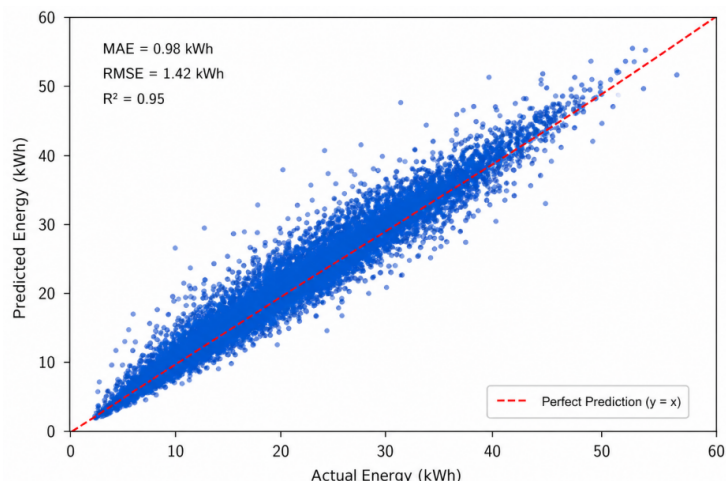


Figure 1. Actual vs. Predicted Energy Consumption Using the XGBoost Model

Figure 1 illustrates the relationship between the predicted and actual energy consumption values obtained using the XGBoost model. Most data points are distributed closely around the diagonal reference line, which represents perfect prediction. This indicates that the model achieves a high level of accuracy in estimating section-level energy consumption.

A relatively small dispersion of points can be observed, particularly under higher energy consumption values. This suggests that prediction errors tend to increase slightly in more complex operating conditions, such as sections with high passenger load, steep gradients, or frequent acceleration and braking. Nevertheless, the overall deviation remains limited, demonstrating the robustness of the model.

The tight clustering of data points around the diagonal line further confirms that the nonlinear relationships between operational variables—such as speed, acceleration, gradient, and load—are effectively captured by the XGBoost model. This supports its suitability for use as the prediction module in the proposed ECO-TrainNet framework.

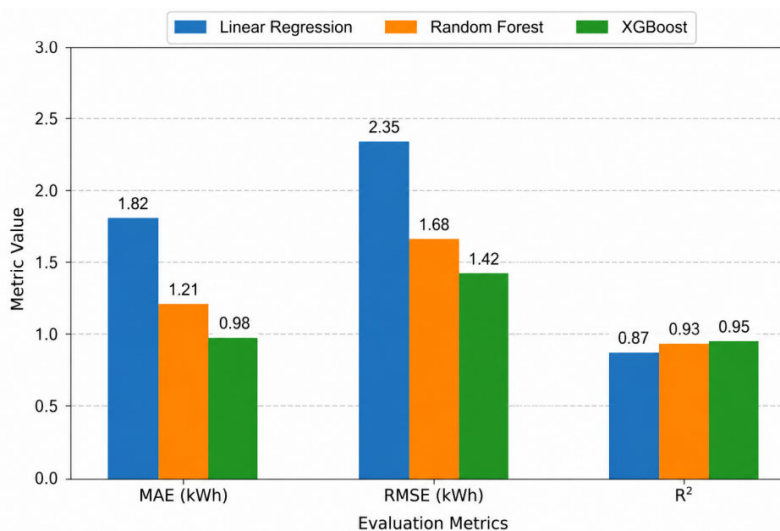


Figure 2. Comparison of Prediction Performance among Different Models

Figure 2 presents a visual comparison of prediction performance among Linear Regression, Random Forest, and XGBoost models across three evaluation metrics: MAE, RMSE, and R^2 .

It can be observed that the XGBoost model consistently achieves the lowest error values (MAE and RMSE) and the highest coefficient of determination (R^2), indicating superior predictive accuracy. Random Forest also demonstrates strong performance, outperforming Linear Regression but remaining slightly inferior to XGBoost. In contrast, Linear Regression exhibits the largest prediction errors and the lowest R^2 value, reflecting its limited ability to capture complex relationships within the data.

The performance differences among the models highlight the importance of nonlinear modeling in rail transit energy prediction. Energy consumption is influenced by multiple interacting factors, including speed, acceleration patterns, gradient, and passenger load. These interactions are inherently nonlinear, making ensemble learning methods more suitable than linear models.

Overall, the results shown in Figure 2 are consistent with the quantitative evaluation presented in Table 2, further confirming that XGBoost is the most appropriate choice for the prediction component of the proposed ECO-TrainNet framework.

4.3. Carbon Emission Distribution and Analysis

Based on the predicted energy consumption, carbon emissions are calculated using time-dependent electricity carbon emission factors. This allows the analysis to reflect not only how much energy is consumed, but also the environmental impact associated with different operating conditions.

The carbon emission for each track section is defined as:

$$C_i = E_i \times \eta_i$$

Where C_i denotes the carbon emission of section i ($kg\ CO_2$), E_i represents the predicted energy consumption (kWh), and η_i is the electricity carbon emission factor ($kg\ CO_2/kWh$).

Unlike conventional approaches that assume a constant emission factor, this formulation allows the model to reflect temporal variations in electricity carbon intensity, thereby providing a more realistic estimation of emissions.

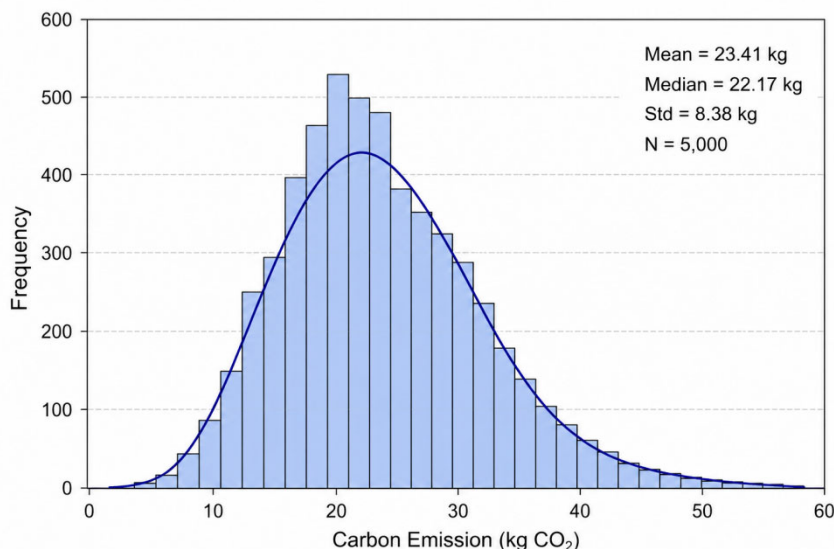


Figure 3. Distribution of Section-Level Carbon Emissions

Figure 3 illustrates the distribution of section-level carbon emissions across the simulated dataset. It can be observed that the distribution is moderately right-skewed, with the majority of values concentrated within a relatively narrow range and a small number of sections exhibiting significantly higher emissions.

This skewness indicates that carbon emissions are not uniformly distributed across the network. Instead, certain sections contribute disproportionately to the overall emissions. These high-emission sections are typically associated with unfavorable operating conditions, such as longer distances, higher passenger loads, steeper gradients, or reduced regenerative braking efficiency.

The presence of such high-emission segments suggests that system-level averages may not be sufficient for effective carbon management. Instead, targeted optimization strategies should be applied to specific sections where emissions are most significant. This observation further supports the necessity of integrating predictive modeling with control optimization, as implemented in the proposed ECO-TrainNet framework.

In addition, the dispersion of the distribution implies that carbon emissions are influenced by multiple interacting factors rather than a single dominant variable. This further justifies the use of machine learning models to capture nonlinear relationships in the prediction stage.

4.4. Reinforcement Learning Training Analysis

To optimize train operation strategies under multiple constraints, a reinforcement learning (RL) agent is trained within the simulated environment. The agent interacts with the environment by selecting acceleration actions based on the current system state, which includes speed, position, remaining distance, and accumulated energy consumption.

During training, the agent aims to maximize the cumulative reward, which is designed to penalize excessive energy consumption, carbon emissions, and schedule deviation. This enables the agent to learn a balanced policy that improves energy efficiency while maintaining operational feasibility.

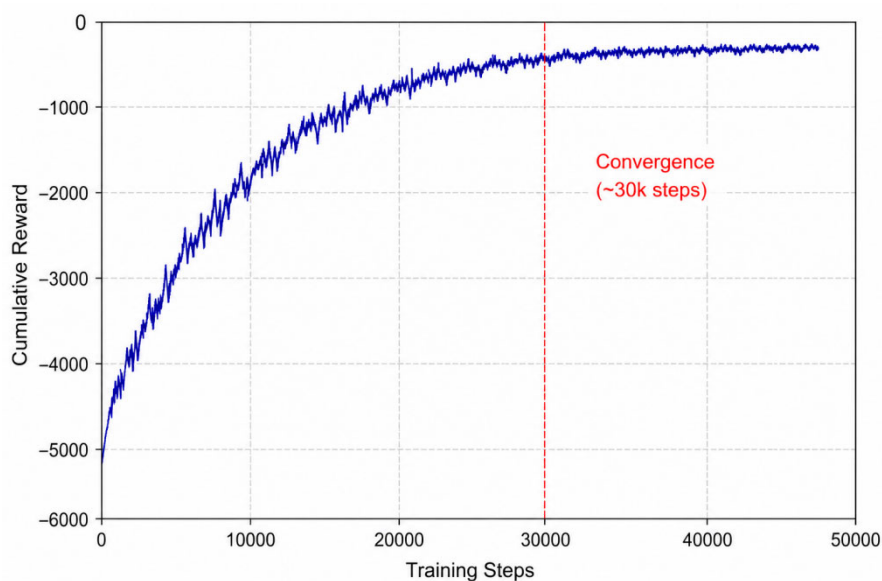


Figure 4. Training Reward Curve of the Reinforcement Learning Agent

Figure 4 presents the training reward curve of the RL agent over 50,000 time steps. It can be observed that the cumulative reward increases rapidly during the early stage of training, indicating that the agent is effectively learning from interactions with the environment. As

training progresses, the growth rate gradually decreases and the reward curve stabilizes after approximately 30,000 steps.

The initial fluctuations in the curve reflect the exploration process of the agent, during which different acceleration strategies are tested. Such variability is typical in reinforcement learning and indicates that the agent is not prematurely converging to a suboptimal policy. The subsequent stabilization of the reward suggests that the agent has learned a relatively consistent and effective control strategy.

The convergence behavior observed in Figure 4 demonstrates that the reinforcement learning framework is capable of identifying energy-efficient and carbon-aware operation policies. This also provides a reliable foundation for subsequent comparisons between baseline and optimized control strategies.

4.5. Control Strategy Comparison

To evaluate the effectiveness of the proposed reinforcement learning-based optimization strategy, three control approaches are compared: a baseline strategy with fixed acceleration patterns, a rule-based strategy with heuristic smoothing and coasting control, and the RL-based strategy learned through interaction with the simulated environment.

Table 3. Comparison of Control Strategies

Strategy	Energy Consumption	Carbon Emissions	Schedule Deviation
Baseline	100.0%	100.0%	0 s
Rule-based	92.5%	93.1%	+2.1 s
RL-based	88.3%	89.0%	+1.8 s

Table 3 summarizes the quantitative results in terms of total energy consumption, carbon emissions, and schedule deviation. Both the rule-based and RL-based strategies achieve noticeable reductions in energy consumption and carbon emissions compared with the baseline strategy. Among them, the RL-based strategy demonstrates the best overall performance.

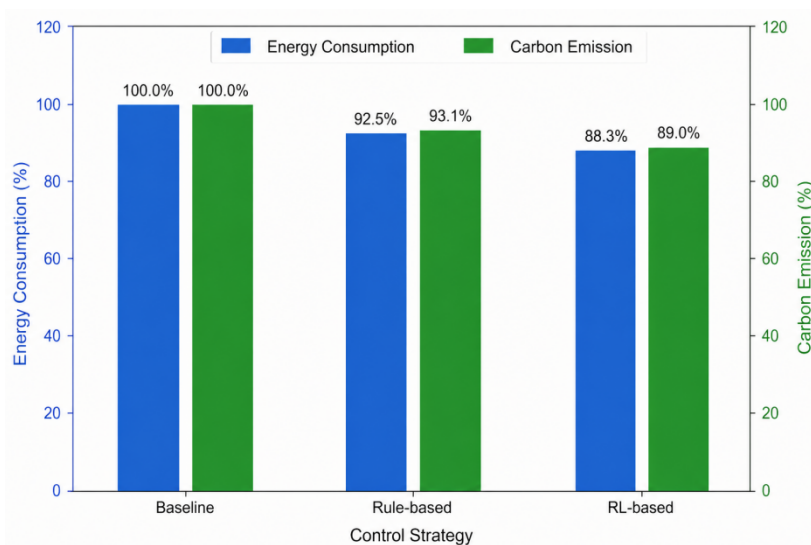


Figure 5. Comparison of Energy Consumption and Carbon Emissions under Different Control Strategies

As shown in Figure 5, the RL-based strategy reduces energy consumption to 88.3% of the baseline level and carbon emissions to 89.0%, corresponding to reductions of 11.7% and 11.0%, respectively. The rule-based strategy also achieves improvements, but its performance remains inferior to that of the RL-based approach.

The superior performance of the RL-based strategy can be attributed to its ability to adapt acceleration decisions dynamically according to the current system state and track conditions. Unlike rule-based methods, which rely on predefined heuristics, the reinforcement learning agent learns a policy that accounts for complex interactions among speed, gradient, load, and braking opportunities.

It is also worth noting that the schedule deviation introduced by the RL-based strategy remains within an acceptable range. This indicates that the optimization does not significantly compromise operational reliability while achieving energy and emission reductions.

Overall, the results demonstrate that integrating reinforcement learning into train operation control can effectively improve energy efficiency and reduce carbon emissions, supporting the effectiveness of the proposed ECO-TrainNet framework.

4.6. Speed Profile Analysis

To further understand the mechanism behind energy and carbon emission reduction, a representative track section is selected to compare the speed profiles generated by the baseline and RL-based strategies.

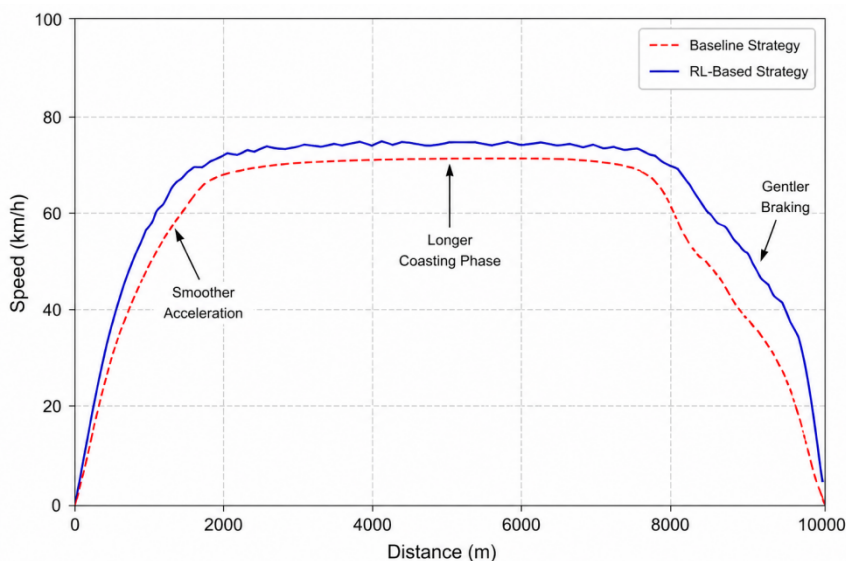


Figure 6. Comparison of Speed Profiles under Baseline and RL-Based Control Strategies

Figure 6 illustrates the variation of train speed over distance under the two control strategies. It can be observed that the baseline strategy exhibits relatively aggressive acceleration and braking patterns, with rapid speed changes and shorter cruising phases. In contrast, the RL-based strategy produces a smoother speed profile characterized by moderate acceleration, extended cruising or coasting periods, and more gradual deceleration.

The smoother speed transitions in the RL-based strategy reduce unnecessary traction power demand during acceleration and avoid excessive energy dissipation during braking. In addition, longer coasting phases allow the train to maintain motion with minimal energy input, thereby improving overall energy efficiency.

Another important observation is that the RL-based strategy tends to better utilize regenerative braking opportunities. By avoiding abrupt braking, the system can recover more kinetic energy and reduce net energy consumption.

These differences in speed profiles explain the performance improvements observed in Section 4.5. The reduction in peak power demand and improved energy utilization directly contribute to lower total energy consumption and, consequently, reduced carbon emissions.

Overall, the results demonstrate that the reinforcement learning-based control strategy is able to learn a more energy-efficient driving pattern, which cannot be easily captured by fixed or rule-based strategies.

4.7. Discussion

The experimental results demonstrate that the proposed ECO-TrainNet framework effectively improves both energy efficiency and carbon emission performance in urban rail transit operations. By integrating prediction and control into a unified pipeline, the framework enables more informed and adaptive decision-making compared with conventional approaches.

A key observation from the results is that accurate prediction plays a foundational role in the overall performance of the system. The superior performance of the XGBoost model indicates that energy consumption and carbon emissions are governed by complex nonlinear interactions among operational variables, such as speed, acceleration, gradient, and passenger load. Without a reliable prediction module, the control strategy would lack sufficient information to evaluate the long-term impact of decisions, particularly in scenarios involving trade-offs between energy efficiency and schedule adherence.

Another important finding is that the reinforcement learning-based strategy consistently outperforms both baseline and rule-based methods. This suggests that data-driven control policies can better adapt to varying operational conditions compared with predefined heuristics. In particular, the RL-based strategy demonstrates the ability to exploit coasting opportunities and reduce unnecessary acceleration and braking, which directly contributes to energy savings. This highlights the potential of reinforcement learning as a flexible optimization tool for complex transportation systems.

From a carbon management perspective, the results confirm that energy optimization can effectively translate into emission reduction when carbon factors are explicitly incorporated into the framework. However, the magnitude of emission reduction is also influenced by the variability of electricity carbon intensity. This implies that operational optimization alone may not fully achieve low-carbon objectives unless it is coordinated with energy supply-side considerations. Therefore, future systems may benefit from integrating real-time carbon intensity data into the decision-making process.

Despite these promising results, several limitations should be acknowledged. First, the experimental dataset is generated through simulation, which may not fully capture the uncertainties present in real-world operations, such as passenger demand fluctuations, signal constraints, and unexpected disturbances. Second, the current reinforcement learning model considers a simplified operational environment and does not explicitly model interactions among multiple trains or network-level constraints. Third, the reward function relies on predefined weighting parameters, which may affect the balance between energy efficiency and punctuality.

In practical applications, additional factors such as safety requirements, passenger comfort, and operational robustness should be incorporated into the optimization framework. Furthermore, the interpretability of reinforcement learning policies remains a challenge, and future research may explore hybrid approaches that combine learning-based methods with rule-based constraints to improve transparency and reliability.

Overall, the findings of this study suggest that the integration of predictive modeling and adaptive control provides a promising direction for achieving energy-efficient and low-carbon rail transit operations.

5. CONCLUSION AND FUTURE WORK

5.1. Conclusion

This study proposes an integrated data-driven framework, ECO-TrainNet, for carbon emission prediction and energy-efficient operation optimization in urban rail transit systems. The framework combines operating-condition recognition, machine learning-based energy and carbon emission prediction, and reinforcement learning-based control optimization within a unified analytical pipeline.

The experimental results demonstrate that ensemble learning models, particularly XGBoost, can effectively capture nonlinear relationships between operational variables and energy consumption, achieving high prediction accuracy. Based on the prediction module, the reinforcement learning agent is able to learn adaptive acceleration strategies that reduce both energy consumption and carbon emissions.

Compared with baseline and rule-based strategies, the RL-based approach achieves a reduction of 11.7% in energy consumption and 11.0% in carbon emissions, while maintaining acceptable schedule deviation. Further analysis of speed profiles reveals that the optimized strategy improves energy efficiency by adopting smoother acceleration, extended coasting phases, and more effective utilization of regenerative braking.

Overall, the results indicate that integrating prediction and control within a closed-loop framework can significantly enhance the energy and environmental performance of urban rail transit systems. The proposed ECO-TrainNet framework provides a feasible approach for supporting carbon-aware operation and intelligent scheduling in rail transit systems.

5.2. Future Work

Despite the promising results, several limitations remain in this study.

First, the dataset used in the experiments is simulation-based and may not fully capture the complexity of real-world metro operations, such as passenger flow variability, signal constraints, and unexpected disturbances. Future work should incorporate real operational data to validate and refine the proposed framework.

Second, the current reinforcement learning model considers a simplified set of operational constraints. In practical applications, additional factors such as passenger comfort, safety margins, and multi-train interactions should be incorporated into the control model.

Third, the carbon emission model is based on section-level estimation and does not explicitly account for network-level interactions or dynamic electricity market conditions. Future studies could extend the framework to multi-line systems and integrate real-time carbon intensity data for more accurate emission assessment.

Finally, further research could explore hybrid optimization approaches that combine reinforcement learning with traditional optimization methods to improve stability and interpretability.

REFERENCES

- [1] Howlett, P. (2016). The optimal control of a train. *Annals of Operations Research*, 98(1–4), 65–87.
- [2] Scheepmaker, G. M., Goverde, R. M. P., & Kroon, L. G. (2017). Review of energy-efficient train control and timetabling. *European Journal of Operational Research*, 257(2), 355–376.

- [3] Albrecht, T. (2010). Reducing power peaks and energy consumption in rail transit systems by simultaneous train running time control. *WIT Transactions on The Built Environment*, 114, 3-14.
- [4] Ding, Y., & Chien, S. I. (2018). Optimizing train trajectory with energy-efficient operation using genetic algorithms. *Transportation Research Part B*, 108, 24-39.
- [5] Li, X., Lo, H. K., & Zhang, K. (2019). Energy-efficient metro train operation using reinforcement learning. *Transportation Research Part C*, 102, 395-412.
- [6] Wei, W., Wu, J., & Sun, H. (2020). Machine learning-based energy consumption prediction for urban rail transit systems. *Applied Energy*, 262, 114551.
- [7] Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. *Proceedings of the 22nd ACM SIGKDD Conference*, 785-794.
- [8] Sutton, R. S., & Barto, A. G. (2018). *Reinforcement Learning: An Introduction* (2nd ed.). MIT Press.
- [9] Zhang, R., & Gao, Z. (2021). Carbon emission estimation and reduction in urban rail transit systems. *Journal of Cleaner Production*, 278, 123-135.
- [10] UIC (International Union of Railways). (2016). *Railway Handbook on Energy Consumption and CO₂ Emissions*.