

Intelligent Control Strategy for Automotive Regenerative Braking Based on Deep Learning

Hongbo Wang

*School of Mechanical and Electrical Engineering, Northeast Forestry University, Harbin, China
WHB0226@gmail.com*

Abstract. The global new energy vehicle industry is developing rapidly, but range anxiety remains a key bottleneck hindering its widespread adoption. Among all methods of improving energy efficiency, the control of regenerative braking is by far the one with the lowest cost—hence its central importance for vehicle performance. Control strategies used in the past do not adjust well to actual working conditions and have difficulty in treating complex nonlinearities or multiple coupled constraints, which rules out the possibility of a fully coordinated multi-objective solution. The present work gives a structured account of basic principles—both of automotive braking energy recovery and of deep learning algorithms. It looks at a two-layer hierarchical control setup based on deep reinforcement learning and its main elements, and considers separately the performance of various approaches under uniform criteria, as well as the engineering obstacles now limiting such systems. Moreover, it offers an organized forecast for the future of the field, with special attention to four possible developments: multi-objective coordinated optimization, transfer learning, distributed vehicle-road cooperation, and multi-mode hybrid control.

Keywords: New energy vehicles, braking energy recovery, deep reinforcement learning, two-layer hierarchical control

1. Introduction

With the global energy crisis and environmental problems now becoming more serious, pushing the automotive sector toward electrification is regarded as a central means of helping countries to meet their carbon peak and carbon neutrality aims. The International Energy Agency's (IEA) *Global EV Outlook 2026* report gives figures indicating that in 2025 the world produced almost 22 million electric vehicles, the figures which represent a year-over-year rise of more than 25%. China's contribution to world production amounts to 75%, to world trade to 40%, with 16 million vehicles manufactured annually (of which 2.5 million are exported) [1]. Yet range anxiety continues to be the central obstacle to mass acceptance of new energy vehicles. At present the major known approaches all focus on raising battery capacity, this development greatly increases both vehicle cost and weight and brings about secondary problems—shortages of such essential items as lithium or cobalt, as well as difficulties with battery recycling or safe disposal. For that reason, attempts to raise energy efficiency (without necessarily expanding battery capacity) now figure centrally among the research goals and technical challenges of the new energy car industry.

As a highly cost-efficient technical pathway for boosting energy utilization efficiency in new energy vehicles, braking energy recovery technology transforms the translational kinetic energy produced during vehicle deceleration into electrical energy, which is subsequently stored in the on-board traction battery. Wager et al. performed chassis dynamometer testing on a Mitsubishi i – MiEV to verify the practical effects of regenerative braking. The test results show that, under diverse parameter configurations and driving cycles, the system can extend the vehicle's driving range by 11% to 22%. Furthermore, it notably cuts down the working duration of friction braking, which further reduces brake pad abrasion and lowers maintenance expenditure throughout the full service life of the vehicle [2]. As a core technical solution, braking energy recovery plays a vital role in improving the energy conversion efficiency of battery electric vehicles and breaking through the limitation of short cruising range. With a well-optimized control strategy, considerable energy-saving gains can be achieved under standard driving conditions. However, the operation of brake energy recovery is subject to multiple constraints, including ECE braking safety regulations, the motor's maximum generator torque, the battery's maximum charging power, and road surface adhesion conditions [3]. How to maximize energy recovery efficiency while ensuring braking safety and ride comfort is the central challenge in control strategy design. Existing braking energy recovery control strategies can be divided into four major categories: traditional control, fuzzy control, neural network control, and intelligent optimization algorithms. Traditional control, represented by rule-based control and PID control, is simple and easy to implement, but relies on preset logic and ideal operating conditions, resulting in insufficient adaptability under complex, time-varying driving conditions; Fuzzy control incorporates expert knowledge to achieve nonlinear multi-objective coordination and offers greater robustness; however, the construction of the rule base involves subjectivity, lacks a self-learning mechanism, and has limited generalization capabilities; neural networks and intelligent optimization algorithms can achieve multi-objective optimization through data-driven approaches or global search, but their high computational complexity and insufficient real-vehicle validation make it difficult to meet the real-time requirements of in-vehicle systems. Overall, the parameters of existing strategies largely rely on manual calibration and are unable to achieve the coordinated optimization of braking safety, ride comfort, and energy recovery under complex dynamic conditions [4].

In recent years, deep learning has provided new approaches to solving multi-constraint dynamic control problems. It does not require an exact system model and can autonomously learn complex control laws from data. Deep reinforcement learning demonstrates particularly significant advantages in continuous action spaces and problems involving dynamic uncertainty, and has been widely applied to regenerative braking systems. However, existing research lacks systematic synthesis and comparison, and experimental conditions and evaluation criteria are inconsistent, making it difficult to reach a consensus and hindering its engineering applications.

To address this, this paper presents a comprehensive review of deep learning-based intelligent control strategies for braking energy recovery. First, the system principles and fundamental algorithms are introduced, with a comparison of the DDPG, SAC, and TD3 algorithms; second, two-layer control architectures and Markov decision process modeling methods are summarized; then, based on unified criteria, the multidimensional performance of different strategies is compared, and the application progress and development potential of deep learning technologies in the field of braking energy recovery control are systematically reviewed, providing theoretical support for promoting technological innovation and engineering implementation in this area.

2. Theoretical foundations of brake energy recovery systems and deep learning

2.1. Principles of automotive brake energy recovery systems

2.1.1. Configuration and operating principles

Brake energy recovery systems in new energy vehicles typically employ a motor-hybrid braking architecture, with the core components consisting of a drive motor, a power battery, bidirectional power electronics, and a mechanical braking unit. During braking, the control system calculates the total braking torque based on parameters such as braking demand, vehicle speed, and the power battery's state of charge (SOC), and dynamically coordinates the distribution of braking force between regenerative braking and mechanical braking. The drive motor switches to generator mode to generate reverse regenerative braking torque, converting the vehicle's kinetic energy into electrical energy. This energy is regulated by a bidirectional converter and stored in the power battery. When the regenerative braking torque is insufficient to meet the total braking demand, the mechanical braking system intervenes to provide supplementary braking torque, jointly ensuring the safety and stability of the vehicle's braking process [5].

Among the brake energy recovery systems used today in new energy vehicles is a motor-hybrid braking type, whose basic parts are a drive motor, a power battery, bidirectional power electronics, and a mechanical brake assembly. At the time of braking, the controller finds out what total braking torque is needed by means of data relating to braking force, to speed, and to the present state of charge (SOC) of the power battery, it also deals dynamically with the allocation of braking force — between regenerative and mechanical braking. The drive motor is switched to generator operation for reverse regenerative braking, thus turning the vehicle's kinetic state into an electric current. That current (or its equivalent) is handled by a two-way converter and is stored in the power battery. Should regenerative braking fail to supply enough torque for full stopping, the mechanical brake is brought in to supplement it, with both systems jointly guaranteeing safe, stable braking of the vehicle [5].

2.1.2. Key constraints

There are a number of restrictions concerning the working of the braking energy recovery system—mainly including [3]:

1. Braking safety regulations: It is necessary that the braking force—distributed between front and rear axles—be in agreement with ECE rules and lie in the safety zone bounded by the ideal braking-force curve and the front-wheel lockup boundary, so as to keep directional control intact for any braking situation and to avoid locking up of wheels.

2. Maximum generator torque constraint: With regard to the motor's rated power and to its operating speed, the regenerative braking torque cannot go beyond the maximum value of generator torque obtainable from it; one has to adjust the upper bound of regenerative braking in light of the motor's external features.

3. Power constraint on battery charging: Because of restrictions — either on the battery's current SOC or on its maximum permissible charging power — the energy returned by regenerative braking as electricity cannot be allowed to go beyond the safe limit of the battery (to avoid overcharging and consequent damage).

4. Road surface adhesion limit constraint: The total braking force must not exceed the maximum adhesion between the tires and the road surface. When braking intensity is high, priority must be

given to ensuring braking stability; if necessary, regenerative braking must be disengaged, and braking force must be provided entirely by mechanical brakes.

2.1.3. Evaluation metrics

The performance evaluation of a braking energy recovery system primarily includes the following four dimensions [6]: energy recovery rate—the ratio of recovered electrical energy to the total kinetic energy lost during vehicle braking—is the core indicator for measuring the system's energy-saving effectiveness; braking distance—the distance traveled by the vehicle from the start of braking to a complete stop—is directly related to braking safety; braking impact—the rate of change in braking deceleration—reflects the smoothness of the braking process and ride comfort; Battery temperature rise rate—the rate at which the battery temperature increases during the braking energy recovery process—affects the battery's cycle life and safety.

2.2. Core algorithms of deep learning

Deep reinforcement learning is a fusion of deep learning and reinforcement learning. It is particularly well-suited to the control requirements of braking energy recovery systems, which involve multiple coupled constraints, continuous action outputs, and dynamic uncertainties in operating conditions. Through continuous interaction between the agent and the vehicle's driving environment, it autonomously learns the optimal regenerative braking force allocation strategy with the goal of maximizing cumulative reward. Its mathematical foundation is the Markov Decision Process (MDP), which can effectively address the challenges of strong nonlinearity and multi-objective cooperative optimization that traditional control methods struggle to handle [7]. Currently, the three most widely used continuous action control algorithms in this field are based on the Actor-Critic framework: Deep Deterministic Policy Gradient (DDPG) was the first algorithm applied to regenerative braking. It consists of a policy network that outputs the front-to-rear brake force distribution coefficients and a value network that evaluates the value of the distribution strategy. It was the first to solve the challenge of continuously and smoothly adjusting brake force; however, it suffers from the drawback of the value network overestimating action values, which can lead to unstable updates in the braking strategy and sudden changes in brake force that affect ride comfort [8]; Double-Delay Deep Deterministic Policy Gradient (TD3) is an improved version of Deep Deterministic Policy Gradient (DDPG). Through three core mechanisms—truncated double-Q learning, delayed policy updates, and target policy smoothing — it significantly mitigates the overestimation of action values by the value network and enhances the stability and convergence performance of policy training in continuous action spaces. It is currently the mainstream improvement for deterministic policy gradient algorithms within the Actor-Critic framework [9].; Soft Actor-Critic (SAC) is a maximum-entropy reinforcement learning algorithm that introduces an entropy regularization term into the reward objective to enhance policy exploration. It employs a stochastic Gaussian policy to adapt to the dynamic uncertainty of the braking process. Under extreme braking conditions — such as high SOC and low coefficient of friction — its robustness is significantly superior to the previous two algorithms, though it converges slightly slower and has a slightly higher computational cost; All three algorithms can meet the millisecond-level control requirements for in-vehicle systems and can be flexibly selected based on the vehicle's application scenario [4].

3. Deep learning-based intelligent control architecture and key technologies

At present, braking energy recovery control — when based on deep learning — is understood to be part of a technical framework centered on a two-layer hierarchical model, one grounded in Markov Decision Process (MDP) logic, and one trained with big data. Such an approach circumvents many of the problems of poor adaptability or of complex multi-objective coordination that have long plagued standard control methodologies [10]. The present chapter sets out the central design and main technical ideas of the deep reinforcement learning-guided two-tier hierarchy and evaluates various control strategies by means of a unified simulation benchmark for performance comparison.

3.1. Design of a general-purpose two-layer hierarchical control architecture

Among the known solutions, the two-layer hierarchical control architecture is standard. It achieves a useful balance between the optimization power of deep learning and the safety or rigidity of the braking mechanism by separating decision-making from execution. The higher layer corresponds to intelligent decision-making, and is based on a Deep Deterministic Policy Gradient (DDPG) network. Among the inputs considered are vehicle dynamics and certain constraint quantities—vehicle speed, motor speed, battery State of Charge (SOC), necessary total torque, and features of dual-motor efficiency. Subject to restrictions involving maximum motor power output, maximum battery power, and best - possible dual - motor efficiency, it yields (in real time) an optimal division of front-to-rear axle motor torques, considering the simultaneous optimization of energy recovery and drive efficiency (multi-objective). The lower layer refers to the braking execution layer. On the basis of orders from the upper layer, it uses data about motor efficiency and about torque sharing to determine whether and how much regenerative braking should be done at the front or rear axles, and also computes the necessary friction braking torque gap that the hydraulic braking system has to compensate for. The design also includes an on-board safety monitor which automatically modifies the torque distribution (in relation to battery SOC or motor load) when battery SOC is excessively high or motor speed is too low [11]. Such an arrangement preserves the benefits of deep learning—its adaptability to complicated situations—without exposing the model to uncontrolled "black-box" behavior thanks to auxiliary hard constraints, an arrangement of great robustness and practical value.

With regard to distributed-drive electric vehicles, the two-layer architecture mentioned above may be extended. The upper layer deals with such data as individual wheel slip rates, speed of the whole vehicle, and braking effort to determine (in light of ideal braking force allocation) the total torque to be applied by the front and by the rear axle(s). It switches to anti-lock braking mode when any wheel slip rate exceeds the optimal threshold, with stability in braking. The lower layer deals with a secondary assignment of regenerative braking torque (from both axle motors) and of hydraulic braking torque—taking into account such limits as battery SOC, maximal motor torque, speed of the car, intensity of braking, and battery charging capability—as a means of attaining joint optimization of energy recovery and of precise control over the braking action [12]. At the current research stage, the deep deterministic policy gradient (DDPG) algorithm can be further extended to the scenarios of speed planning and global energy management. With traffic operation efficiency, energy efficiency and riding comfort set as multi-objective optimization indicators, the collaborative optimal control of vehicle speed planning and regenerative braking can be achieved.

3.2. Core key technologies

3.2.1. MDP-based modeling of the braking process

Converting braking control into a Markov Decision Process (MDP) is a prerequisite for the application of deep reinforcement learning. The key lies in designing a reasonable state space, action space, and reward function. The state space includes core variables such as vehicle speed, brake pedal position, battery SOC, battery temperature, and motor speed, balancing representational completeness with computational efficiency; The action space is defined as the regenerative braking force distribution coefficient between the front and rear axles (with values in the range [0,1]), enabling flexible and smooth adjustment of braking force to avoid the sudden, jarring impacts associated with traditional control methods; the reward function adopts a multi-objective weighted form, expressed as:

$$R = \omega_1 R_{energy} + \omega_2 R_{comfort} + \omega_3 R_{safety}.$$

where R denotes the comprehensive immediate reward fed back to the reinforcement learning agent at each decision step; R_{energy} , $R_{comfort}$ and R_{safety} represent the sub-reward terms corresponding to regenerative energy recovery efficiency, braking ride smoothness, and braking safety performance, respectively; ω_1 , ω_2 and ω_3 are the non-negative weighting coefficients of each objective. These weight parameters can be calibrated through optimization algorithms to realize synergistic optimization of energy recovery, ride comfort and braking safety, while mitigating the subjectivity arising from manual parameter tuning.

This approach delivers a comprehensive balance among energy recovery efficiency, braking jerk and braking distance deviation; the weight coefficients can be determined using optimization algorithms, which cuts down the subjectivity associated with manual parameter calibration [5].

3.2.2. Dataset construction and preprocessing

High-quality data are basic necessities for model training and should include both standard test cycles and real-world road examples. Standard test cycles (WLTC, CLTC, or UDDS) help with offline learning and with benchmarking, and they yield curves relating speed to time under average driving situations; datasets taken from real roads (for example, the EV Kinematic Dataset) capture extended-duration driving data for many different vehicle models are considered—whether urban, suburban, or highway-based—and such data help to assess the model's generality. Before starting training, all input quantities are normalized (to lie in the interval[-1,1]);missing values go through linear interpolation; extreme values (outliers) are deleted using the 3σ rule; and the whole dataset is divided in a 7:3 fashion between training and testing to achieve efficient and stable convergence of the model.

3.3. Performance comparison of different control strategies

Wahid et al. compared three regenerative braking strategies: baseline control, Type-2 fuzzy control, and DDPG-RL. Under WLTC conditions, the DDPG-RL strategy achieved optimal performance, reducing energy consumption to 1.99 kWh, with a regenerative energy contribution rate of 18.15%. The efficiency of energy usage reaches 12.59 km/kWh, and driving range goes up by 15.4%relative

to that of the baseline case. Power distribution is—within the usual city-related speed range of 20 – 40 km/h—more stable [13]. When the baseline model is tested experimentally, its error is limited to 0.17%– 0.3%, which shows how reliable the comparison is likely to be. From this analysis it is clear that the DDPG-RL method gives good results for energy recovery and for extending LEV range; hence it may be regarded as an especially efficient tool for sustainable optimization.

4. Research challenges and future prospects

4.1. Existing challenges

Even though deep learning-oriented braking energy recovery control has now advanced considerably, many problems still exist in real-world engineering contexts. Deep reinforcement learning calls for a huge volume of interaction data during training, and this is neither fast nor economical with respect to the use of samples; it also complicates the task of adjusting promptly to various vehicles or to complex driving situations.

In addition, deep neural networks are complicated and expensive to run computationally, so obtaining millisecond-class real-time performance with low-cost in-vehicle controllers is hard; yet deliberate simplification of the network causes poor control accuracy. At present, practically all studies of this sort are confined to ordinary or average operating conditions, without systematic analysis of extremes—such as rain, snow, and icy roads—do not fully justify reliance on this technology. Moreover, the so-called black-box character of deep learning models (with no obvious explanation of their decisions) may lead to safety problems. There is still no universally accepted set of tests or of evaluation criteria for such models. Because of the many differing simulation setups and test scenarios used by various authors, one cannot easily compare their results, nor does that help to push forward systematically with real-world applications of the method.

4.2. Future development trends

With regard to the above-mentioned difficulties, future developments of deep learning-oriented intelligent control systems for braking energy recovery will follow these lines:

1. Multi-objective Collaborative Optimization: Besides the usual concerns of energy recovery, of safety and of comfort, the cycle life of the battery and of the motor are to be taken into account for full optimization throughout the vehicle's lifetime.

2. Transfer learning and Federated Learning: With transfer learning it is possible to pass quickly from a model trained for one car type to another, with considerable saving of training effort; federated learning enables collaborative model training across multiple vehicles without sharing raw driving data, which protects data privacy while improving algorithm generalization

3. V2X-enabled vehicle-road cooperative control: If roadside perception is used to anticipate certain future traffic situations or braking events, then predictive braking energy recovery becomes feasible, and the energy recovery rate is better still.

4. Hybrid control architecture: An attempt is made here to combine the advantages of deep learning with those of conventional control, so that deep learning control may be used normally (to yield maximum performance) and yet—when needed—traditional control can be applied to avoid hazards and to improve the system's dependability.

5. Conclusion

The present paper gives a systematic review of intelligent control methods using deep learning to recover energy from automotive braking. First, it identifies the two-tier hierarchical control setup as the main technical approach now being pursued here. With its division of the intelligent decision process from the safe implementation part, it is able to harness the power of deep reinforcement learning (in complex situations) to achieve dynamic optimization under conditions as well as mitigation of the "black-box" risks tied to the model by means of hard constraints at the lower level, thus reconciling robustness with engineering feasibility; Second, of the three algorithms derived from the Actor-Critic idea—namely DDPG, TD3, and SAC—each has an appropriate domain. With regard to standard settings, the TD3 algorithm converges more rapidly and yields a higher rate of energy recovery, the present choice of this method is now regarded as preferable for engineering applications, with the SAC algorithm proving to be notably more robust—under extreme cases such as high SOC or low coefficient of friction; Third, the general performance of deep learning control methods vastly exceeds that of standard rule-based or PID control. From an experimental standpoint it is known that such strategies help to raise the proportion of regenerative energy by about 18% and to expand driving range beyond 15% of the previous value, at the same time as braking jerk is curtailed, thus enabling multi-objective coordination of safety, comfort, and energy economy.

Deep learning presents an entirely new optimization framework for braking energy recovery systems, one that is well suited to solving such industry problems as nonlinear behavior or complex coupling of constraints (usually hard for traditional controls to manage). Among the engineering issues still unresolved—training too slowly, difficulty in meeting real-time demands of the car, or poor validation under extreme operating conditions remain valid, and developments in lightweight algorithms, in transfer learning and in vehicle-road coordination now seem feasible, so it is reasonable to anticipate that deep-learning-oriented control of braking energy will be part of standard technology for future generations of new energy vehicles—a core technical aid both in reducing range anxiety and in raising the efficiency of energy use.

References

- [1] International Energy Agency. (2026). *Global EV outlook 2026*.
- [2] Wager, G., Whale, J., & Braunl, T. (2018). Performance evaluation of regenerative braking systems. *Proceedings of the Institution of Mechanical Engineers, Part D: Journal of Automobile Engineering*, 232(10), 1414–1427.
- [3] Wang, Y., & Pei, K. (2022). Research on optimization of braking energy recovery strategy for pure electric vehicles. *Mechanical Science and Technology for Aerospace Engineering*, 41(9), 1436–1441.
- [4] Zhu, J., Liu, H., Xing, B., Chen, J., & Fan, A. (2026). A review of control strategies for brake energy recovery systems. *Energies*, 19(10), 2275.
- [5] Prakash, Z. (2025). Integration of AI and ML in regenerative braking for electric vehicles: A review. *Frontiers in Artificial Intelligence*, 8, 1626804.
- [6] Szumska, E. M. (2025). Regenerative braking systems in electric vehicles: A comprehensive review of design, control strategies, and efficiency challenges. *Energies*, 18(10), 2422.
- [7] Ananganó-Alvarado, G., Umaña-Morel, I., & Keith-Norambuena, B. (2025). Reinforcement learning in electric vehicle energy management: A comprehensive open-access review of methods, challenges, and future innovations. *Frontiers in Future Transportation*, 6, 1555250.
- [8] Naseri, F., Nordin, B. R. a., Spiliotopoulos, K., & Schatz, E. (2026). Regenerative–frictional brake blending in electric vehicles considering energy recovery and dynamic battery charging limit: A reinforcement learning-based approach. *Machines*, 14(4), 416.
- [9] Fujimoto, S., Hoof, H., & Meger, D. (2018, July). Addressing function approximation error in actor-critic methods. In *International Conference on Machine Learning* (pp. 1587–1596). PMLR.

- [10] Jeyanthi, P. A., & Manoraja, S. (2026). Smart energy optimization in electric vehicles using reinforcement learning algorithms. *OPSEARCH*.
- [11] Ouyang, T., Jin, S., Xie, X., Gong, Y., & Zhang, Z. (2025). Adaptive energy management in dual-motor electric vehicles using deep deterministic policy gradient. *IEEE Transactions on Transportation Electrification*.
- [12] Li, N., Lin, Y., Huang, Z., Hong, Y., & Ning, X. (2026, April). Research on speed planning and energy management strategy for distributed-drive electric vehicles based on deep deterministic policy gradient algorithm. In *Actuators* (Vol. 15, No. 5, p. 248). MDPI.
- [13] Wahid, M. R., Joelianto, E., Budiman, B. A., Dewanata, M. P., & Aziz, M. (2026). Optimizing regenerative braking in light electric vehicles using deep deterministic policy gradient reinforcement learning. *Egyptian Informatics Journal*, 33, 100893.