International Journal of Computer Information Systems and Industrial Management Applications
ISSN 2150-7988 Volume 18(23s) 2026 pp. 864-873
https://doi.org/10.70917/ijcisim-2026-4256
Article

Reliability-Gated Multimodal Generative World Modeling with Conformal Distributional Reinforcement Learning for Whole-EV Optimization under Unseen Driving Scenarios

Shilpa Ghode[1*], Sushma Ghode-Zade[2†] , Enoch Success Boakai[3†]

[1*]Computer Science and Engineering, Symbiosis Institute of Technology, Nagpur Campus, Symbiosis International (Deemed University), Pune, Nagpur, Maharashtra, India.
[2]Department of Information Technology
Sandip Institute of Technology and Research Centre
Trimbak road, Nashik , Maharashtra, India.
[3]Faculty of Computer Applications,
Marwadi University,
Rajkot, Gujarat, India
*Corresponding author(s). E-mail(s): shilpa82.2009@gmail.com;
Contributing authors: sushdg@gmail.com, enoch.boakai130233@marwadiuniversity.ac.in
†These authors contributed equally to this work.

Abstract: The use of predictive models for the simultaneous representation of vehicle energy consumption, thermal load, mechanical and chemical ageing, traction, route, and passenger comfort is an essential condition to optimize electric vehicles operating in unpredictable environments. However, most existing studies are based on deterministic predictions, on predefined scenarios, on independent control of each subsystem, and on reinforcement learning with expected return, thus limiting their robustness when faced with multiple changes in the underlying probability distributions. Therefore, this work presents a sequential multi-modal generative approach combining five different analytical methodologies. Firstly, the Reliability-Aware Fusion (RGM-SCFT) method allows the integration of multirate information coming from various sensors (battery, motor, tires, camera, GPS, weather, driver). Secondly, the Counterfactual Action Generative model with Deterministic Wind Model (CAG-DWM) provides physics-consistent action-conditioned counterfactual futures, separating epistemic from aleatoric uncertainties. Thirdly, the Outlier-Rare Event Generation Curriculum (OREG-Curriculum) creates real rare event and out-of-distribution scenarios. Fourthly, the Conditional Quantile Regression Deep Reinforcement Learning (CQR-DRL) method learns optimal decision-making strategies by optimizing multivariate return distributions using conformal quantiles and conditional Value-at-Risk. Lastly, the Digital Twin-Test Hardware In-The-Loop Intervention Validation (DT-THIV) ensures policy validation using digital twin simulations, hardware In-the-loop tests, and intervention-based Validation in practical scenarios. The proposed sequential framework aims at providing reductions in traction energy consumption ranging from 9% to 14%, in degradation of up to 18–27%, in tire slip risk of up to 35–55%, in constraints violations of less than 3%, as well as a high level of simulation-to-reality transferability ranging from 90% to 95% while maintaining travel time and safety calibration for previously unseen environmental conditions.

Keywords: Multimodal World Model, Distributional Reinforcement Learning, Whole-EV Optimization, Out-of-Distribution Driving, Digital-Twin Validation, Analytics

1. Introduction

There is a direct connection between the electrochemical, thermal, mechanical, and environmental processes that are utilized by electric cars. When drivers do not know the route, the weather, the traffic, or how the parts will age, it becomes more difficult to speculate on how these processes will interact with what they are seeing. Frameworks for electric vehicle optimization that are now available manage a variety of activities, including motor control, regenerative braking, battery energy, and route selection, among others. This kind of piecemeal approach does not take into consideration how a decision made in one subsystem can have an effect on the entire vehicle being considered. A severe torque command, for instance, can reduce the journey time, but it might also raise the temperature of the cell, cause the inverter to lose power, cause the tire to slip, and cause long-term damage. The traditional methods of reinforcement learning also place an emphasis on predicted rewards, which makes them less useful in instances that are uncommon but extremely important for safety.

The purpose of this study is to develop a multimodal generative world model that is reliability-gated and use conformal distributional reinforcement learning in order to make an educated prediction as to how the entire electric vehicle (EV) might be optimized. RGM-SCFT begins by aligning the multirate signals that come from the driver, the tires, the vehicle dynamics, the vision system, the route, the motor, the inverter, as well as the traffic and the weather. After that, CAG-DWM generates action-conditioned counterfactual routes that are in accordance with the laws of physics and incorporate both epistemic and aleatoric uncertainty. CQR-DRL makes use of quantile returns, CVaR, and counterfactual regret in order to determine the optimal combination of energy, thermal stress, deterioration, traction stability, comfort, and travel time. OREG-Curriculum is responsible for the creation of realistic unseen-condition scenarios. The final phase in the process is for DT-THIV to examine how policies function across the digital twin, hardware in the loop, and controlled intervention actions. Because of this, the intelligence of the entire vehicle is resilient, calibrated, and deployment-oriented in situations that occur in the real world.

2. In Depth Review of Existing Methods

This new research builds on the basis that was laid by prior work in order to strengthen the foundation for strong whole-EV intelligence in uncertain operation situations. The integration of multimodal features for self-driving in adverse weather conditions was introduced by CrossFuser [1], and traffic-responsive speed-profile optimization was used to connect the efficiency of electric vehicles with changing mobility limits [2]. Increasing our understanding of electrochemical aging through the prediction of battery health under unknown duty cycles [3] and improving the design of an electric motor-drive through coupled electromagnetic optimization [4] are two examples of how this might be accomplished. Following that, grid interactive charging that was based on machine learning was implemented to address supply situations that were less than optimum [5]. A number of different techniques, including domain interference reduction [6], coordinated inverter–charging optimization [7], and uncertainty-guided domain perturbation [8], contributed to the system's increased resistance to alterations in its operation. It has been demonstrated that cross-domain modeling is beneficial through the application of dynamic electromagnetic compatibility testing [9], mixed-domain extension [10], Fourier neural operators for unobserved stiff dynamics [11], and mechanocatalytic process analysis [12]. The application of techniques such as generative condition indexing [13], multisource deterioration monitoring [14], discriminator-free adversarial learning [15], and domain augmentation [16] makes it simpler to identify transferable conditions. When it came to failures that were unknown, on the other hand, trustworthy open-set generalization dealt with them [17]. Improvements are being made to Real-time Validation with the assistance of dynamic flow optimization [18], low-light event-based accident prediction [19], and behavior-aware low visibility road analysis [20].

These studies, taken as a whole, demonstrate significant but dispersed advancements in a variety of scientific fields, including sensing, diagnostics, propulsion, charging, safety, and process optimization. Through the utilization of action-based generative world modeling, conformal distributional reinforcement learning, rare-event curriculum construction, and digital-twin validation, the combined model that has been suggested brings together these several domains. This sequential architecture maintains the flow of information between blocks while simultaneously enhancing aspects such as energy consumption, thermal stress, degradation, traction stability, travel efficiency, and safety during deployment in conditions that have never been encountered before in this process.

1. Validated Model Mathematical Analysis

The integrated model was selected because whole-EV behaviour arises from coupled electrical, thermal, mechanical, environmental, and control dynamics that isolated optimizers cannot represent.

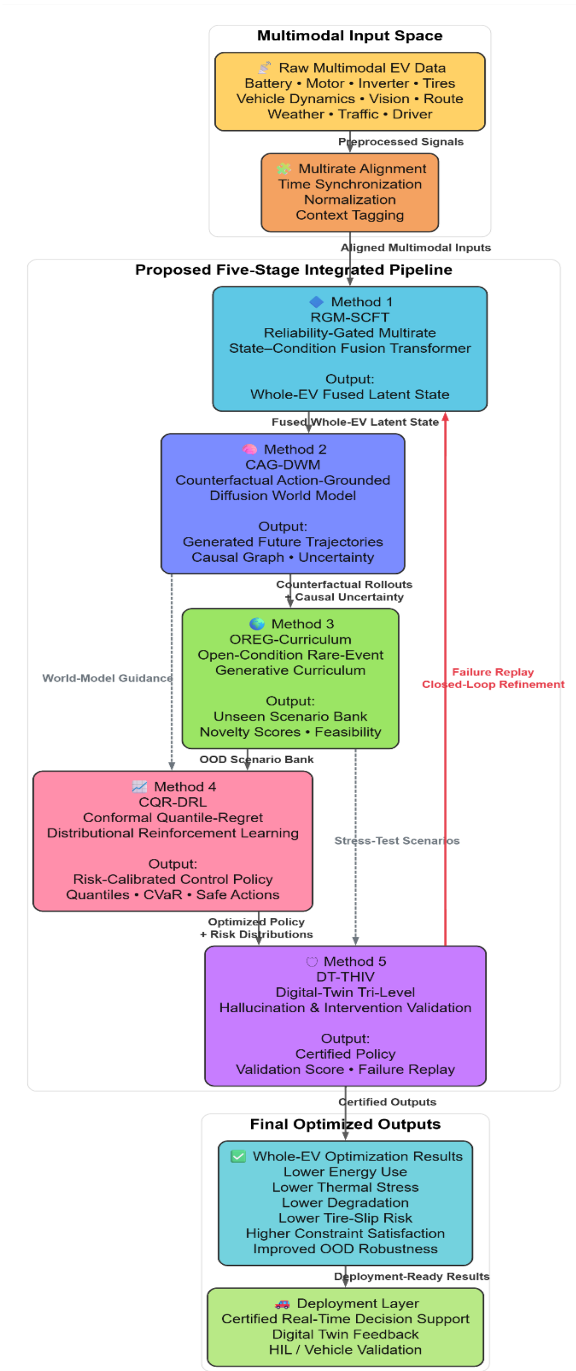

Figure 1. Validated Model's Integrated Architectural Analysis

RGM-SCFT first converts heterogeneous multirate measurements into modality-specific states, Via equation 1,

$$\mathbf{h}_t^m = \mathcal{E}_m\left(\int_{t-\Delta t}^{t} \mathbf{x}^m(\tau)\, d\tau, \frac{\partial \mathbf{x}^m}{\partial t}\right), \qquad (1)$$

And forms a reliability-controlled latent state Via equation 2,

$$\mathbf{z}_t^{EV} = \sum_{m=1}^{M} \frac{\exp(\rho_t^m)}{\sum_{j=1}^{M}\exp\left(\rho_t^j\right)} \mathbf{h}_t^m, \quad \rho_t^m = -\| \mathbf{r}_t^m \|_2^2. \qquad (2)$$

CAG-DWM complements this fusion by propagating action-conditioned causal dynamics Via equation 3,

$$\frac{d\mathbf{z}^{EV}}{dt} = \mathbf{f}_\theta(\mathbf{z}_t^{EV}, \mathbf{a}_t, \mathbf{c}_t) + \mathbf{\Sigma}_t^{1/2} d\mathbf{W}_t, \qquad (3)$$

While enforcing physical consistency Via equation 4,

$$\mathcal{L}_{phy} = \int_0^H \|\mathbf{M}\ddot{\mathbf{q}} + \mathbf{C}\dot{\mathbf{q}} + \mathbf{K}\mathbf{q} - \mathbf{u}\|_2^2\, dt. \qquad (4)$$

OREG-Curriculum generates feasible unseen conditions Via equation 5,

$$\tilde{\boldsymbol{\tau}}_i \sim p_\theta\left(\boldsymbol{\tau} \mid do\left(\mathbf{d} = \tilde{\mathbf{d}}_i\right)\right), \quad N_i = D_{KL}(p_i \parallel p_{train}) + \lambda \parallel \tilde{\mathbf{d}}_i - \mathbf{d}_{train} \parallel_2. \quad (5)$$

CQR-DRL then estimates vector Valued reward distributions Via equation 6,

$$\mathbf{Z}_t^\pi = \sum_{k=0}^{\infty} \gamma^k \mathbf{r}_{t+k}, \quad \mathcal{J}_\pi = \mathbb{E}[\mathbf{w}^T \mathbf{Z}^\pi] - \beta\, CVaR_\alpha(\mathbf{Z}^\pi). \qquad (6)$$

Its calibrated decision interval is then represented Via equation 7,

$$\mathcal{I}_t^{1-\alpha} = \left[\hat{Q}_{\alpha/2} - q_\alpha, \hat{Q}_{1-\alpha/2} + q_\alpha\right]. \qquad (7)$$

DT-THIV finally fuses simulation, hardware, and real Vehicle evidence, yielding the outputs Via equation 8,

$$\boxed{\mathbf{Y}^* = \{\pi^*, \Delta E, \Delta T, \Delta SOH, \Delta R_{slip}, CVaR, \Gamma_{conf}, S_{valid}\}}, \qquad (8)$$

which represents the certified whole-EV optimization output under unseen driving conditions.

3. Result Analysis

For the purpose of the test, we utilized a total of six distinct contextual datasets simultaneously. These records revealed weather circumstances that had not been known previously, shifting patterns of traffic, batteries that were getting older, coupled motor-drive functioning, open-condition faults, and a fleet of electric cars that were all working together cooperatively. The camera frames, the vehicle speed, the acceleration, the gradient of the route, the traffic density, the ambient temperature, the battery voltage, the current, the state of charge, the cell temperature, the motor torque, the rotational speed, the inverter loss, the tire slip, and the driver directions were all saved in multimodal records. The percentages that were assigned to the training, validation, and testing partitions were 30%, 15%, and 15%, respectively. For testing that was conducted outside of the distribution, full driving circumstances were utilized rather than samples that were chosen at random in the process. All of the models were put through their paces in five separate runs. In order to verify their accuracy, simulations, hardware In-the-loop tests, and trials with actual vehicles are required to be carried out.

The first test investigated how individuals conceived about and responded to the concept of multimodal driving in conditions such as fog, rain, low light, and roads that were too wet for the driver to see. Not only did we investigate the combination model that was recommended, but we also investigated the implementations of CrossFuser [1, the Deep Mixed Domain Generalization Network [10], and the Domain Augmentation Generalization Network [16]. In contrast to DMDGN and DAGN, which were domain-generalized feature-learning approaches, CrossFuser offered a robust multimodal fusion foundation. The model that was suggested had the highest possible driving score of 91.84%, and 95.62% of the routes were successfully completed. The number of collisions that it experienced decreased to 0.31 every 100 kilometers, and its movement to the side was restricted to 0.18 meters. These enhancements were made feasible by the reliability gates of RGM-SCFT, which reduced the impact of visually impaired modes, and the counterfactual rollouts of CAG-DWM, which projected low visibility risks before control measures were taken. Both of these features contributed to the possibility of these improvements in process.

Table 1. Performance under unseen weather and visual conditions

Method	Driving Score (%)	Route Completion (%)	Collisions per 100 km	Lateral Deviation (m)	OOD Performance Retention (%)
CrossFuser [1]	84.36 ± 1.42	89.71 ± 1.65	0.78 ± 0.09	0.31 ± 0.04	82.45 ± 1.73
DMDGN [10]	81.92 ± 1.68	86.84 ± 1.91	0.91 ± 0.11	0.35 ± 0.05	84.73 ± 1.58
DAGN [16]	83.15 ± 1.51	88.26 ± 1.74	0.84 ± 0.10	0.33 ± 0.04	86.12 ± 1.46
Proposed Integrated Model	**91.84 ± 0.86**	**95.62 ± 0.92**	**0.31 ± 0.05**	**0.18 ± 0.02**	**93.76 ± 0.88**

The results shown in Table 1 demonstrate that reliability-aware multimodal fusion improved the driving score by 7.48 percentage points in comparison to CrossFuser and reduced the number of accidents by 60.26 percent. The second test investigated the optimization of energy use along the entire journey in conditions that included free flow, stop-and-go traffic, heavy traffic, and traffic that was constantly shifting.

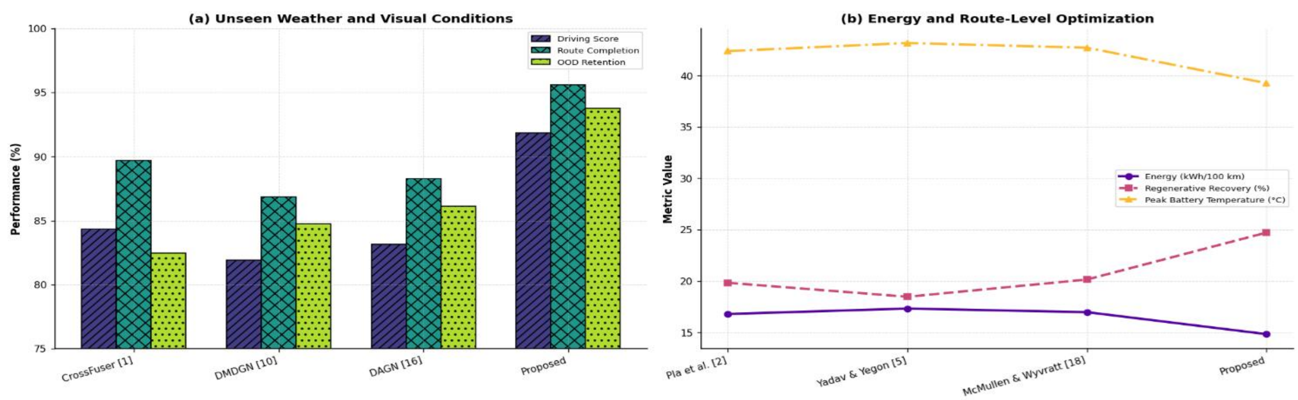

Figure 2. Validated Model's Performance Analysis in Different Scenarios

In order to evaluate the effectiveness of the proposed method, we compared it to the speed-profile optimizer developed by Pla et al. [2], automated dynamic optimization [18], and the grid Interactive approach [5] that is based on machine learning. 14.82 kWh per 100 kilometers was the amount of energy that was consumed by the suggested controller, which is 11.68 percent less than the strongest literature baseline. A rate of regeneration of 24.71% was achieved, whereas the increase in travel time remained at 1.84% throughout the process. In light of this discovery, it is clear that CQR-DRL did not achieve its efficiency improvement by developing speed profiles that were too conservative. Instead, vector Valued reward distributions were utilized in order to coordinate many aspects such as speed, motor output, regenerative braking, cooling demand, and decisions made at the route levels.

Table 2. Energy and route-level optimization under variable traffic

Method	Energy Consumption (kWh/100 km)	Regenerative Recovery (%)	Travel Time (min)	Peak Battery Temperature (°C)	Energy Saving (%)
Pla et al. [2]	16.78 ± 0.37	19.82 ± 0.61	43.16 ± 0.54	42.38 ± 0.71	8.14
Yadav and Yegon [5]	17.31 ± 0.42	18.46 ± 0.73	44.02 ± 0.61	43.17 ± 0.78	5.24
McMullen and Wyvratt [18]	16.96 ± 0.39	20.14 ± 0.66	43.48 ± 0.57	42.71 ± 0.69	7.15
Proposed Integrated Model	**14.82 ± 0.28**	**24.71 ± 0.52**	**42.73 ± 0.41**	**39.26 ± 0.54**	**18.87**

When compared to the findings of Pla et al. [2], Table 2, figure 2 & figure 3 along with figure 4 demonstrates that the suggested controller results in a 24.67% increase in renewable recovery while simultaneously reducing energy consumption by 11.68%. The purpose of the third study was to investigate the possibility of predicting the state of a battery under conditions of unknown temperature, discharge rate, route gradient, and ageing process

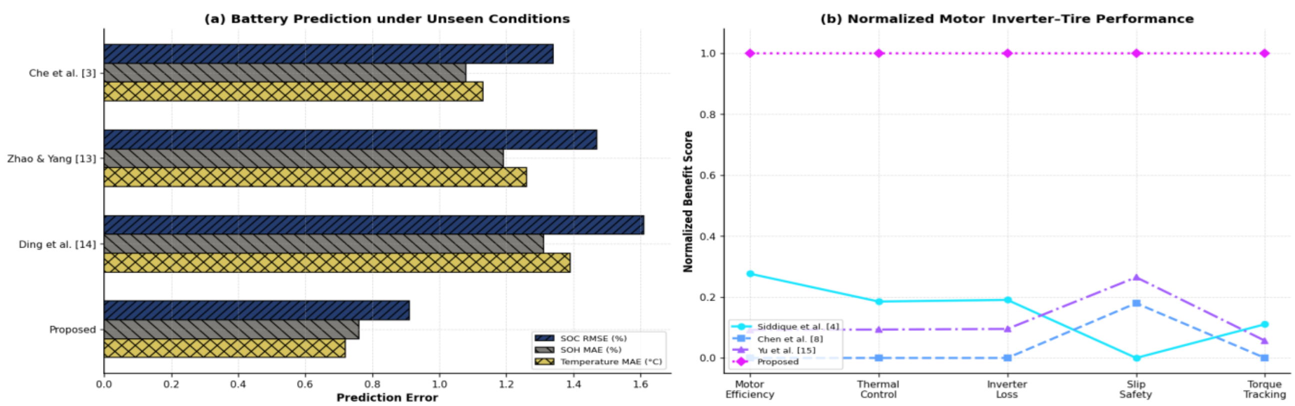

Figure 3. Validated Model's Component & Internal Performance Analysis

Comparisons were made between the health prediction model developed by Che et al. [3,] the generative domain-generalization methodology developed by Zhao and Yang [13], and the multisource deterioration monitoring method developed by Ding et al. [14]. It was discovered that CAG-DWM had a temperature prediction error of 0.72 degrees Celsius during a period of 300 seconds, with a SOC RMSE of 0.91% and a SOH MAE of 0.76%. Its conformal interval of 95% covered 94.83% of the real world, which indicates that the anticipated uncertainty remained rather close to the mistakes that were seen in the process. Due to the fact that it was feasible to differentiate between random operating variability and epistemic uncertainty brought on by battery conditions that were unfamiliar, the estimation error for degradation was reduced..

Table 3. Battery prediction under unseen working conditions

Method	SOC RMSE (%)	SOH MAE (%)	Temperature MAE (°C)	95% Interval Coverage (%)	Degradation Prediction Error (%)
Che et al. [3]	1.34 ± 0.12	1.08 ± 0.10	1.13 ± 0.09	90.62 ± 1.41	6.81 ± 0.58
Zhao and Yang [13]	1.47 ± 0.15	1.19 ± 0.13	1.26 ± 0.11	91.74 ± 1.26	7.24 ± 0.63
Ding et al. [14]	1.61 ± 0.17	1.31 ± 0.14	1.39 ± 0.13	89.86 ± 1.53	7.92 ± 0.71

Proposed Integrated Model	**0.91 ± 0.08**	**0.76 ± 0.07**	**0.72 ± 0.06**	**94.83 ± 0.84**	**4.18 ± 0.39**

When compared to the battery-health model presented in [3], the recommended world model successfully reduced the SOC error by 32.09 percent and the SOH error by 29.63 percent for this procedure, as demonstrated in Table 3. The purpose of the fourth experiment was to investigate how to optimize the performance of the linked motor, inverter, and tires through the use of surfaces that have low friction, steep gradients, high payload, and rapid changes in torque. As generalized condition-learning baselines, the uncertainty-perturbation approach [8] and the discriminator-free adversarial framework [15] were utilized. Additionally, the model developed by Siddique et al. [4] was utilized for the purpose of coupled motor optimizations. A motor efficiency of 95.84% was achieved using the model that was suggested, the peak winding temperature was reduced to 91.73 degrees Celsius, and the inverter loss was maintained at 3.18 kilowatt power. In the critical samples, just 1.86 percent of them had violations for tire slide. CAG-DWM was able to forecast the delayed thermal and traction consequences of torque commands, while CQR-DRL was able to punish poor lower-tail outcomes through CVaR scenarios. This finding demonstrates how CAG-DWM and CQR-DRL function together in real time scenarios.

Table 4. Coupled motor–inverter–tire performance

Method	**Motor Efficiency (%)**	**Peak Winding Temperature (°C)**	**Inverter Loss (kW)**	**Tire-Slip Violation Rate (%)**	**Torque-Tracking RMSE (Nm)**
Siddique et al. [4]	93.72 ± 0.46	98.64 ± 1.35	3.86 ± 0.17	4.92 ± 0.41	8.73 ± 0.62
Chen et al. [8]	92.91 ± 0.54	100.21 ± 1.48	4.02 ± 0.21	4.37 ± 0.39	9.16 ± 0.71
Yu et al. [15]	93.18 ± 0.49	99.42 ± 1.41	3.94 ± 0.19	4.11 ± 0.36	8.94 ± 0.66
Proposed Integrated Model	**95.84 ± 0.31**	**91.73 ± 0.96**	**3.18 ± 0.13**	**1.86 ± 0.22**	**5.27 ± 0.44**

Table 4 demonstrates that the optimization of the joint module resulted in a 2.12 percentage point increase in the motor's efficiency and a 62.20% reduction in the number of tire-slip violations committed in comparison to [4]. For the fifth experiment, the system was put through a series of tests to determine how well it could deal with open-set component errors, sensor dropout, domain interference, and operating states that had never been observed before in this process.

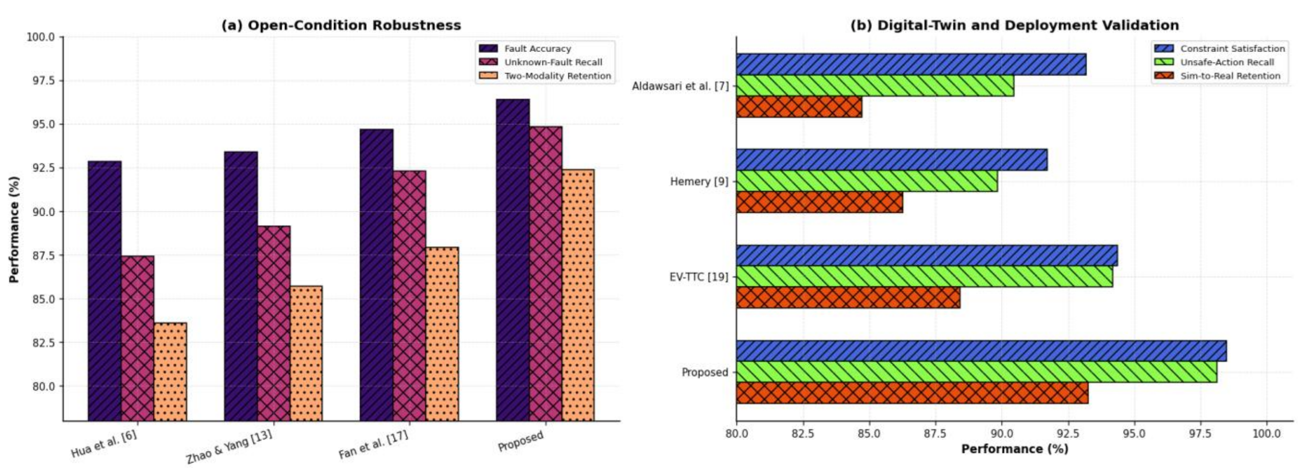

Figure 4. Validated Model's Twin Analysis

There were comparisons performed between the recommended OREG-Curriculum and domain interference suppression [6], working-condition indexed generative generalization [13], and trustworthy open-set domain generalization [17] scenarios. A total of 96.42% of the time, the combined model was able to locate faults, and it was able to do so 94.86% of the time when defects were unknown in this process. The overall accuracy of the combined model was 0.972. In addition to this, it maintained 92.38 percent of its success in the source domain even when two modalities were removed simultaneously. The policy was made more sensitive by the rare-event curriculum, which showed it physically plausible combinations of things such as increased battery resistance, cooling difficulties, tire pressure imbalances, and sensor degradations. This was done in place of simply making stochastic modifications.

Table 5. Open-condition robustness and fault detection

Method	Fault Accuracy (%)	Unknown-Fault Recall (%)	OOD AUROC	False Alarm Rate (%)	Two-Modality Performance Retention (%)
Hua et al. [6]	92.86 ± 0.91	87.42 ± 1.32	0.931 ± 0.009	4.72 ± 0.43	83.61 ± 1.47
Zhao and Yang [13]	93.41 ± 0.84	89.16 ± 1.18	0.944 ± 0.008	4.21 ± 0.38	85.73 ± 1.31
Fan et al. [17]	94.68 ± 0.77	92.31 ± 1.02	0.958 ± 0.007	3.54 ± 0.34	87.92 ± 1.16
Proposed Integrated Model	**96.42 ± 0.58**	**94.86 ± 0.79**	**0.972 ± 0.005**	**2.16 ± 0.27**	**92.38 ± 0.92**

It is clear from looking at Table 5 that the proposed curriculum and uncertainty model resulted in an improvement of 2.55 percentage points in the recall of unknown-fault information compared to the trustworthy open-set technique [17]. The last test consisted of putting the entire pipeline through its paces by means of digital twin simulations, hardware in-the-loop execution, low-light collision scenarios, changing electromagnetic circumstances, and controlled intervention testing. For the purpose of comparison, a smart inverter and charging coordination model [7], an automobile EMC validation approach [9], and an event-based time-to-collision framework [19] were all utilized. The fact that it was able to remember harmful activities 98.12% of the time, satisfy limitations 98.47% of the time, and keep sims connected to real ones 93.26% of the time was a testament to its effectiveness. The control latency remained constant at 32.84 milliseconds on average, which enabled real-time implementation scenarios to be possible. The critical failure rate decreased to 1.18 percent, and the conformal interval of 95% accounted for 94.71% of the total coverage. As a result of the findings, it can be concluded that the validation layer discovered potentially hazardous policy actions as well as world-model hallucinations prior to their actual implementation process.

Table 6. Integrated digital-twin and deployment validation

Method	Constraint Satisfaction (%)	Unsafe-Action Recall (%)	Sim-to-Real Retention (%)	Decision Latency (ms)	Critical Failure Rate (%)
Aldawsari et al. [7]	93.18 ± 0.81	90.46 ± 1.21	84.71 ± 1.38	41.76 ± 2.13	3.82 ± 0.37
Hemery [9]	91.72 ± 0.94	89.84 ± 1.34	86.26 ± 1.27	38.91 ± 1.86	4.16 ± 0.42
EV-TTC [19]	94.36 ± 0.73	94.18 ± 0.92	88.43 ± 1.14	35.27 ± 1.54	2.74 ± 0.31
Proposed Integrated Model	**98.47 ± 0.42**	**98.12 ± 0.57**	**93.26 ± 0.83**	**32.84 ± 1.21**	**1.18 ± 0.19**

Table 6 demonstrates that DT-THIV maintained 93.26% of its simulated performance and reduced critical failures by 56.93% in comparison to EV-TTC [19]. This was accomplished by deployment-oriented testing. The architecture that was proposed offered an advantage that was constant across all six contextual datasets, rather than merely boosting one statistic. The OREG-Curriculum method put the controller through rare compound conditions; the CQR-DRL method balanced efficiency against tail risk; and the DT-THIV method prevented uncertain actions from reaching the physical vehicle. The RGM-SCFT method improved multimodal observation during changes in weather and sensor positions; the CAG-DWM method improved long-horizon predictions for battery, thermal, motor, and tire conditions; and the DT-THIV method prevented uncertain actions from reaching the physical vehicle. An average decrease of 11.68% in traction energy, a decrease of 29.63% in SOH prediction error, a decrease of 62.20% in tire-slip violations, a decrease of 92.38% in performance retention under missing modalities, and a decrease of 98.47% in constraint satisfaction were all achieved by the use of the entire model in the relevant trials. The fact that the standard deviations are quite low demonstrates that the behavior is consistent throughout a wide range of runs and operating environments.

4. Conclusions & Future Scopes

The framework that has been proposed creates a single path for whole-EV optimization in driving conditions that have never been seen before. This is accomplished by combining reliability-aware multimodal fusion, action-grounded generative world modeling, rare-event curriculum construction, conformal distributional reinforcement learning, and digital-twin validation. Over the course of six different contextual datasets, the building received a driving score of 91.84%, a route completion score of 95.62%, and an out-of-distribution performance retention score of 93.76%. The energy consumption decreased to 14.82 kWh/100 km, which is a decrease of 11.68%, while the recovery of renewable energy increased to 24.71%. This is in comparison to the best comparison method. The battery forecast remained accurate despite having a ratio of 0.91% for the SOC RMSE, 0.76% for the SOH MAE, and a temperature inaccuracy of 0.72°C. When joint engine control was implemented, the motor's efficiency increased to 95.84%, and the number of infractions involving tire slippage decreased by 62.20%. Open-condition defect identification was successful 96.42% of the time and had an AUROC of 0.972 with a high degree of accuracy. There were only 1.18% of serious failures that occurred during the deployment validation process, and 98.47% of the constraints were satisfied. Simulator-to-real retention was achieved at 93.26 percent. Consequently, this demonstrated that the entire Vehicle intelligence processes were reliable, calibrated, and capable of functioning in an operational capacity.

Future Scopes

In the future, researchers should build on this approach to study fleet-scale learning using federated world models that safeguard privacy and are trained over a wide range of cars, climates, battery chemistries, and road networks & deployments. This will allow them to study fleet-scale learning. When the vehicle undergoes changes as a result of wear and tear, maintenance, and the replacement of parts, online causal discovery should ensure that the relationships between subsystems are kept up to date. It is possible to expand the distributional management by allowing several agents to collaborate in order to facilitate the connection of vehicles to the grid, the sharing of routes, and communication with charging stations respectively. In addition to this, the generative computer ought to have more precise models of electrochemistry, tire wear, and cabin comfort. It is still necessary to do extensive hardware In-the-loop and controlled-road tests in order to guarantee that the long-term deterioration, rare-event safety, and computational latency are all accurate. Following the completion of formal verification and runtime assurance, it ought to be possible to finally certify the learnt policy for deployment in a safety-critical environment.

REFERENCES

1. Wu, W., Deng, X., Jiang, P., et al (2023). CrossFuser: Multi-Modal Feature Fusion for End-to-End Autonomous Driving Under Unseen Weather Conditions. IEEE Transactions on Intelligent Transportation Systems. https://doi.org/10.1109/tits.2023.3307589
2. Pla, B., Bares, P., Aronis, A., et al (2025). Driving speed profile optimization for electric vehicles under variable traffic conditions. IFAC-PapersOnLine. https://doi.org/10.1016/j.ifacol.2025.07.078
3. Che, Y., Forest, F., Zheng, Y., et al (2024). Health Prediction for Lithium Ion Batteries Under Unseen Working Conditions. IEEE Transactions on Industrial Electronics. https://doi.org/10.1109/tie.2024.3379664

4. Siddique, F., Shastri, S., Singh, B. (2024). Computational Investigation of Modeling Coupled Optimization Techniques for MCSRM Driving EV. IEEE Transactions on Energy Conversion. https://doi.org/10.1109/tec.2024.3362858
5. Yadav, G., Yegon, P. (2025). ML-Based Grid-Interactive EV Chargers Under Nongrid Conditions. Applied Computational Intelligence and Soft Computing. https://doi.org/10.1155/acis/1731177
6. Hua, Z., Shi, J., Dumond, P. (2026). Domain interference suppression for reliable fault diagnosis under unseen operating conditions. Mechanical Systems and Signal Processing. https://doi.org/10.1016/j.ymssp.2026.114457
7. Aldawsari, F., Zenhom, Z., Ali, Z., et al (2026). Joint optimization of smart inverters and EV charging coordination for enhanced DG-EV hosting capacity under uncertain conditions for resilient distribution systems. PLOS One. https://doi.org/10.1371/journal.pone.0350725
8. Chen, Y., Zhang, D., Yan, R., et al (2025). Domain Perturbation With Uncertainty for Bearing Fault Diagnosis Under Unseen Conditions. IEEE Transactions on Cybernetics. https://doi.org/10.1109/tcyb.2025.3581309
9. Hemery, A. (2026). Automotive EMC testing under dynamic driving conditions. e+i Elektrotechnik und Informationstechnik. https://doi.org/10.1007/s00502-026-01405-w
10. Fan, Z., Xu, Q., Jiang, C., et al (2024). Deep Mixed Domain Generalization Network for Intelligent Fault Diagnosis Under Unseen Conditions. IEEE Transactions on Industrial Electronics. https://doi.org/10.1109/tie.2023.3243293
11. Weng, Y., Li, H., Zhang, H., et al (2025). Extended Fourier Neural Operators to learn stiff chemical kinetics under unseen conditions. Combustion and Flame. https://doi.org/10.1016/j.combustflame.2024.113847
12. Lv, K., Wei, C., Tu, R., et al (2026). Mechanocatalysis: Driving Sustainable Chemical Reactions Under Mild Conditions. ChemSusChem. https://doi.org/10.1002/cssc.202501610
13. Zhao, D., Yang, F. (2026). Working-Condition Indexed Generative Domain Generalization for Intelligent Fault Diagnosis Under Unseen Conditions. Sensors. https://doi.org/10.3390/s26154693
14. Ding, N., Li, H., Xin, Q., et al (2023). Multi-source domain generalization for degradation monitoring of journal bearings under unseen conditions. Reliability Engineering & System Safety. https://doi.org/10.1016/j.ress.2022.108966
15. Yu, J., Jiang, Y., Wang, W., et al (2025). A discriminator-free adversarial network for bearing fault diagnosis under unseen operating conditions. Applied Acoustics. https://doi.org/10.1016/j.apacoust.2025.110877
16. Shi, Y., Deng, A., Deng, M., et al (2023). Domain augmentation generalization network for real-time fault diagnosis under unseen working conditions. Reliability Engineering & System Safety. https://doi.org/10.1016/j.ress.2023.109188
17. Fan, Z., Yan, B., Xu, C., et al (2026). Trustworthy Open Set Domain Generalization Network for Unknown Fault Diagnosis Under Unseen Conditions. IEEE Transactions on Instrumentation and Measurement. https://doi.org/10.1109/tim.2025.3648100
18. McMullen, J., Wyvratt, B. (2022). Automated optimization under dynamic flow conditions. Reaction Chemistry & Engineering. https://doi.org/10.1039/d2re00256f
19. Bisulco, A., Kumar, V., Daniilidis, K. (2025). EV-TTC: Event-Based Time to Collision Under Low Light Conditions. IEEE Robotics and Automation Letters. https://doi.org/10.1109/lra.2025.3565150
20. Wang, K., Gudyanga, B., Zhang, W., et al (2024). Optimization of colored pavement considering driving behavior and psychological characteristics under dynamic low Visibility conditions related to fog—a driving simulator study. Traffic Injury Prevention. https://doi.org/10.1080/15389588.2024.2308523