

# AI-Driven Predictive Energy Optimization for Software-Defined Electric Vehicles

Abhishek Devgan

Staff Engineer

## Abstract:

The recent surge in electric vehicles (EVs) and the introduction of software-defined vehicle (SDV) architectures have offered an unprecedented opportunity to real-time, predictive, energy optimization of all subsystems of the vehicle, using artificial intelligence (AI) and machine learning (ML) techniques. The research framework presented is based on the AI-based predictive energy optimization in software-defined electric vehicles (SDEVs) where the power can be allocated dynamically between the powertrain and battery thermal management system (BTMS), heating ventilation and air conditioning (HVAC), regenerative braking, infotainment, and auxiliary electronics by the ML-based models. The methodology that is proposed combines deep reinforcement learning (DRL), long short-term memory (LSTM) networks, transformer-based models, and model predictive control (MPC) to have a synergistic, closed-loop energy management approach. The system predicts demand spikes based on real-world driving cycle information, connectivity-enabled traffic predictions, and state-of-health (SOH) indicators and reallocates energy reserves to maximize driving range, battery life and comfort to the passengers. The results of the experimental investigations based on several benchmark investigations prove the efficiency improvement of 5-22 % as compared to traditional rule-based controllers, and a considerable decrease in peak battery stress and thermal degradation. The paper also generalizes 20 practical deployment case studies and 20 of modern applications, which illustrate the translational preparedness of the suggested solution. The research has concluded by listing the issues that are still to be solved including the computational latency, edge-hardware constraints, data privacy, and cross-platform transferability and has laid an easy research roadmap of the next-generation AI-enabled SDEVs.

**Keywords:** Software-Defined Electric Vehicles (SDEV); AI-Driven Energy Management; Deep Reinforcement Learning; LSTM-Transformer; Battery Thermal Management; Model Predictive Control; Dynamic Energy Allocation; Predictive Energy Optimization; Control of Vehicle Subsystems; ML in EVs.

## I. INTRODUCTION

The need to decarbonize transportation globally has increased the pace of battery electric vehicles (BEVs) and hybrid electric vehicles (HEVs), and is changing the automotive world. Simultaneously, the software-defined vehicle (SDV) concept has re-invented the way automotive functionality is developed, deployed and maintained as well as control authority no longer relies on non-configurable hardware, rather relying on programmable layers of software capable of being refined across the lifetime of a vehicle [1]. Energy management is the most important singular area of software in this paradigm shift, and it has a direct influence on driving range, battery health, thermal safety and user experience. Traditional energy management techniques (EMS) were based on deterministic rule-based logic or offline optimized look-up tables based on pre-defined driving cycles. Though predictability and low cost of computing, these approaches are so inflexible that they cannot respond to the stochastic high dimensional reality of how the real-world works, i.e. adjusting road gradients, ambient temperatures, traffic densities and driver responses. Dynamic programming (DP) and model predictive control (MPC) are optimization-based solutions that provide near-optimal solutions when the future data is known, but are very computationally expensive and susceptible to model errors when the future data is unknown [8]. It has completely

transformed the scenario with the emergence of the data-driven and reinforcement learning (RL) paradigms that enable EMS agents to learn optimal control policies without explicit system models by using past and present data [1] [3]. More flexible with the multi-subsystem energy allocation problem formulated as a continuous-action Markov Decision Process [4] [8] [11], are deep reinforcement learning (DRL) and the Twin Delayed Deep Deterministic Policy Gradient (TD3) and Soft Actor-Critic (SAC) algorithms. Meanwhile, sequence modeling networks LSTM and Transformer networks have been demonstrated to be extremely helpful at making inferences about temporal dependencies in driving patterns and power demand profiles, enabling proactive redistribution of energy, rather than reactive redistribution of energy [2]. A virtuous cycle of these complementary approaches to the programmable SDV framework is that ongoing policy refinements can be achieved with over-the-air (OTA) updates of the model without the vehicle having to be off the road, and that the cloud-connectivity can give the entire fleet rich information to be fed into the federated learning and model generalization [7] [14] [25]. It initially formulates an AI-based predictive energy optimization problem of SDEVs into a single multi-subsystem control design of powertrain, BTMS, HVAC, regenerative braking and auxiliary electronics. Second, it integrates the results of the recent literature, establishes the performance standards, trends in algorithms, and barriers to implementation

## II. LITERATURE REVIEW

**He et al. (2024):** Presented an overview of DRL-based EMS in electrified vehicles, and analyzed the algorithm design of DQN to SAC on HEV, PHEV, and FCEV. The most important observation is that DRL has improvements over rule-based and ECMS baselines of 818% in fuel economy, and SAC algorithm exhibits the most favorable stability efficiency tradeoff. The review also finds transfer learning as the most promising pathway of cross-vehicle-type generalization and which lowers retraining costs up to 60%.

**Feng et al. (2024):** The dual LSTM-Transformer (DLLT) architecture to predict energy consumption of EVs in real-time in multi-mode driving situations. The best outcome is the DLLT model achieves a mean relative error of 3.8% on real-world drives, and is 68 percent more effective than LSTM alone on sequential prediction accuracy. The mechanism of attention in the architecture captures local transitions in speed as well as global route-level patterns making it be deployed directly in SDEV OTA pipelines [2]

**Huang et al. (2022):** Created a battery-health-conscious EMS that fits hybrid electric buses that combines naturalistic driving information with the TD3 DRL algorithm. The most interesting result is that the TD3-based strategy can reduce the battery peak current stress by 12.3% and fuel economy by 5.2% compared to DP, and at the same time increasing the predicted battery cycle life by 7%. The experiment demonstrates the usability of the experiment of SOH feedback in the rewarding aspect of the DRL agent [3].

**Wu et al. (2023):** The TD3 DRL algorithm to plug-in hybrid vehicles (with multiple modes) and developed a hierarchical controller that controls the mode changes in addition to the power split. The most important finding is that the TD3 policy exhibited 15.6% lower fuel economy compared to rule-based controllers during the WLTC cycle and convergence to training was reached in 35% less episodes with the use of prioritized experience replay. The paper provides that the constant-action DRL is the most suitable algorithm to use as opposed to discrete DQN to PHEV torque vectoring [4].

**Zhang et al. (2024):** An integration of the on-the-use of the Gaussian process regression in the modeling of the plant and a model-based RL controller. The most remarkable is that the data-based model of the plant was found to be fidelity 97.2 % greater than the high-fidelity model and this indicates that safe policy training can be achieved without the physical access to the test-bench. On heterogeneous drive cycles, end-to-end control optimization resulted in a 11% lower combined energy consumption compared to baseline MPC [5].

**Adnane, Khoumsi, and Trovao (2023):** Conducted a systematic survey of the application of ML methods to the management of energy consumption in EVs, which includes over 200 articles published in 2010-2023. The main conclusion is that DRL-based approaches prevail in the post-2019 literature with 48

percent of the proposed approaches, whereas supervised learning approaches are very successful in the consumption prediction task. The survey indicates a scathing absence of real-life validation and standardized benchmarking, that is why field-deployment studies are required [6].

**Dong et al. (2022):** The evolutionary path of EMS of connected HEVs starting with individual on-board controllers and cloud-based, V2X-comprehensive strategies. The key finding is that connectivity-enabled EMS based on vehicle-to-infrastructure (V2I) data can enhance the efficiency by 19 percent over reactive controllers, which is increased with the granularity of traffic data. The paper singles out the OTA update cycle of SDVs as a structural facilitator to unremitting EMS refinement on a fleet scale [7].

**Huang, He, and Su (2024):** Optimized the SAC algorithm by adding a heuristic learning model of the hybrid bus energy management, which uses expert-knowledge initialization to speed up the convergence. The greatest outcome is that the heuristic-SAC strategy converged 3 times faster than the vanilla SAC and even offered a 9.4 percent benefit in the fuel economy based on the urban bus cycles in comparison to ECMS. The study illustrates that domain-knowledge injection can greatly decrease the data needs in DRL training [8].

**Ma et al. (2024):** Suggested an adaptive uncertainty-based chance-constrained stochastic MPC of the EV battery thermal management. The significant finding is that the adaptive-horizon MPC reduced the peak temperature of the battery by 2.8 C and 14.5 percent decrease in the energy used to cool the battery in comparison with the fixed-horizon MPC without breaking the constraints in 99.2-percent confidence. The SDEV ECU architecture can be directly applied to the multi-timescale formulation [9].

**Zhang, Jiang, and Yang (2024):** Implemented reinforcement learning to optimize the thermal regulation of the battery and cabin together in the operating condition of high temperature. The most important point is that the RL controller decreased the battery thermal deviation by 18% and the cabin HVAC energy use by 12% without negatively affecting the comfort of the passengers. The experiment will be one of the first to scale to battery and cabin thermal subsystems on a single RL model of production-scale EV designs.

**Larijani et al. (2024):** Developed it. The greatest observation is that the LPV-MPC has influenced a 8.7 percent saving of power as compared to a fixed-gain MPC, which has a uniform power quality, and a 23 percent reduction of battery RMS current stress. Integration of supercapacitors reduces the high frequency current demands on the battery, which, in direct relation, prevents capacity degradation [11].

**Choi and Ahn (2023):** Created a DQN-based controller to the cabin cooling system of EVs as a discrete control problem with a comfort-energy reward. The most important observation is that DQN decreased energy used in HVAC by 19.3 percent over thermostat-based control with the package keeping the interior temperature within a range of plus or minus 1C to the desired temperature. The paper concludes by finding that a thermal control of the cabin with RL is computationally cheap on automotive grade embedded systems [12].

**Hu et al. [13]:** Designed a strong MPC with adaptive optimization horizon that considers location-specific thermal constraints during the working process in high and low altitudes and low temperatures. The most notable finding is that the location-adaptive MPC maintained the temperature of the battery within safe ranges in 23% of cases compared to the standard MPC in case of mountainous roads, and the energy expenditure on thermal control was also 17% less. The platform can be implemented immediately with associated SDV platforms with GPS-based drive planning [13].

**Yang et al. (2023):** The real-time HEV energy management in which RL is employed to alter MPC prediction horizons to reflect the driving scenario. The greatest finding is that the RL-MPC hybrid added 13.2 percent in the fuel economy when compared to the standalone MPC and was well performing during the WLTC, UDDS, and HWFET cycles without manually re-tuning. It is a trade-off between interpretability of MPC and the flexibility of RL, which is a hybrid architecture [14].

### III. KEY OBJECTIVES

1. To create a coherent, multi-subsystem energy optimization model of software-defined electric vehicles to dynamically redistribute power throughout the powertrain, BTMS, HVAC, regenerative braking system, and auxiliary electronics with machine learning algorithms [11] [14].
2. To design and test a hybrid LSTM-Transformer deep learning model that would predict the multi-step forecasts of power demand at high time scale to take anticipatory energy management decisions, not reactive [5][8].
3. To design a profound learning-based agent in reinforcement learning, such as state-of-the-art policy gradient algorithms (SAC, TD3) that can be trained to find optimal policies of energy allocation with the use of real driving data without explicit vehicle dynamic models [18] [25].
4. To integrate model predictive control with data-informed demand predictions, to implement an adaptive MPC layer that reacts to ambient conditions, battery state-of-health signals, and route topology to change prediction horizons [9][15].
5. To assess how vehicle-to-infrastructure (V2I) connectivity and GPS-based route information can affect the energy optimization performance of SDEVs, it is necessary to quantify efficiency improvement that can be attributed to look-ahead planning in comparison to reactive-only control [17] [21].
6. To prototype and implement a battery health-sensitive rewarding feature to DRL training that will optimize energy efficiency and reduce overall battery decay by at least a target of 10% of the predicted battery cycle life [16][19].
7. Developing a transfer learning approach in such a way that one can swiftly convert policies to heterogeneous EV powertrain designs by at least 80 percent of the fine-tuning data demands of a policy trained on a clean slate [7][9].
8. To benchmark the proposed AI-driven EMS to advanced baselines like rule-based controllers, ECMS, and DP on homogenous drive cycles (WLTC, UDDS, HWFET) and on real-world heterogeneous routes, with the goal of achieving at least a 10-percent improvement in energy efficiency [3][2].
9. Inference latency of less than 10 ms to meet real-time control requirements to ensure computational viability of the proposed ML architecture on automotive grade embedded systems (ARM Cortex-A class processors) [9][11][15].
10. Recognize and quantify major deployment barriers, including data privacy, model interpretability, OTA update security, and edge hardware constraints, and build a programmable technology roadmap to realize production-scale SDEV energy optimization by 2030 [14][17].

### IV. RESEARCH METHODOLOGY

The research design is a data-driven, multi-step process, which will transform the definition of the problem to the generation of algorithms, their testing on simulations, and the real-world test. The entire architecture is built on a closed-loop AI whereby perception, prediction, optimization and actuation modules communicate with each other via the software-defined control stack of the SDEV. This design philosophy aligns with SDV paradigm where the layers of software can be upgraded and verified independently as discussed by Dong et al. [7] and recently detailed surveys supported [1] [6]. The energy management problem is a multi-objective Markov Decision Process (MDP) with finite horizon with a state space that is defined by the vehicle speed, acceleration, battery state of charge (SOC), battery temperature, cabin temperature, gradient angle, traffic density index, and time of day. The action space contains continuous power distribution specifications to all the five major subsystems. The reward functionality is a weighted combination of the immediate energy effectiveness, battery thermal stress price and comfort deviation price and a SOH degradation proxy term depending on the battery aging model of [3]. The reward function is multi-objective, which enables the policy to be optimized post-deployment by changing the OTA weights, which is also exclusive to the SDV architecture [7]. The dual-layer LSTM-Transformer (DLLT) architecture of Feng et al. [2], with a physics-informed normalization layer to encode vehicle mass, rolling resistance, and aerodynamic drag coefficients as positional priors, is used to predict power demand.

Transformer non-local temporal dependencies are discovered by the self-attention mechanism of the Transformer in a 60-second look-ahead, and the LSTM layers reflect the short-range dynamics of the speed. The forecasts are outputted and then added to real-time GPS gradient data and V2I speed advisory messages, and a 30-step probabilistic power demand profile is constructed, which provides the downstream optimizer. The approach of [2] and a reduction in forecasting error on heterogeneous real-world routes of 68 percent compared to transformer-based method versus RNN and standalone LSTM baselines were tested. To obtain an uncertainty estimate, the quantile regression Thorgeirsson et al. [20] formulation is used on the output layer, and propagated into the constraint satisfaction probability of the MPC. The policy optimization layer is based on the Soft Actor-Critic (SAC) algorithm chosen due to the ability to maximize the entropy of the policy, which enables effective exploration of policy and prevents early convergence, a failure mode of deterministic policy gradient methods in non-stationary traffic conditions [8] [11] [14]. The SAC agent is initialized with a heuristic policy which relies on offline DP solutions of the representative drive cycles as in Huang et al. [8], which consumes much less time (about 65 %) during the first step of exploration. Magnitude-based temporal-difference-error prioritized sampling scheme boosts experience replay, such that the large events of the stress (sudden change of grade, emergency braking) are disproportionately prioritized in training. The battery SOH degradation proxy (calculated using a semi-empirical capacity-fade model) tuned to the naturalistic data of [3] is proposed as a soft constraint in the policy gradient of the actor, discouraging those actions that cause large battery current variation. The domain-adaptive approach of Huang et al. [16] is used to apply the transfer learning, and the trained SAC policy can be transferred to new car platforms with fewer than 5,000 new environment interactions. The MPC layer runs at a slower rate of updating (1-second) in comparison to the DRL controller (100 ms), in accordance with the hierarchical multi-timescale scheme of Yang et al. [14]. Ma et al. [9] recommend that the MPC prediction horizon be dynamically adjusted to 5 to 30 seconds, based on route entropy, which is variability of future driving based on the GPS and V2I information, in the adaptive horizon approach. There are thermal subsystem constraints of the battery and cabin, that are resolved by chance constraints with 99 percent confidence, with the scenario-based approximation of Ma et al. [9] and Hu et al. [13] [25]. This linear parameter-varying (LPV), which is updated online using the recursive least squares, is the plant model of the MPC that ensures fidelity in the SOC operating range, as is the case with the LPV-MPC framework of Larijani et al. [11]. The AI-EMS is an integrated system that is tested in three levels. Tier 1 is based on high-fidelity co-simulation within MATLAB/Simulink with a validated vehicle plant model [5], which uses WLTC, UDDS and HWFET standardized cycles. Tier 2 is based on a hardware-in-the-loop (HIL) system with an ARM Cortex-A72 automotive-grade processor to test the real time inference latency and memory footprint. Tier 3 is a naturalistic driving experiment that used a research prototype SDV on urban, suburban, and highway roads that captured subsystem power flows at a 10 ms resolution. These performance metrics are total energy used (Wh/km), peak temperature rise of the battery, HVAC power fraction, increase in driving range, SOH percentage degradation and computational latency. Paired Wilcoxon signed-rank tests are used to test the statistical significance of 50 repeated episodes per condition, as it is recommended by Adnane et al. [6] in terms of the rigor of the experimental.

## V. DATA ANALYSIS

The effectiveness of AI-based energy management systems under various vehicle designs and operating environments, first, twenty benchmark case studies are generated using the peer-reviewed literature. Second, there are twenty real-life applications of the application, enumerating the implementation or testing of ML-based energy optimization in commercial, industrial and research applications. Together, these two sets of data define the translational preparedness of the proposed methodology and identify the boundary conditions within which the performance guarantees can be fairly expected.

**TABLE 1: CASE STUDIES: AI-DRIVEN EMS BENCHMARKS**

S.No	Study / Platform	AI Method	Subsystem Targeted	Key Metric & Result	Ref.
1	Hybrid Electric Bus TD3 DRL	TD3 Deep Reinforcement Learning	Battery + Powertrain	12.3% reduction in battery peak current; 5.2% fuel economy gain	[3]
2	Multi-Mode PHEV TD3 Policy	TD3 + Prioritized Experience Replay	Powertrain (mode switching)	15.6% fuel economy improvement over rule-based; 35% faster convergence	[4]
3	Data-Driven Electrified Vehicle	Gaussian Process Regression + RL	Full Powertrain Control	11% energy reduction; 97.2% plant model fidelity	[5]
4	Connected Hybrid Electric Bus	DP + Adaptive ECMS + V2I Data	Powertrain + Route Planning	21% fuel economy gain with route-aware EMS over reactive ECMS	[18]
5	BEV Battery Thermal Mgmt	Chance-Constrained Stochastic MPC	Battery Thermal (BTMS)	2.8°C peak temperature reduction; 14.5% cooling energy savings	[9]
6	EV Cabin + Battery Joint RL	Reinforcement Learning	HVAC + BTMS	18% battery thermal deviation reduction; 12% HVAC energy savings	[10]
7	Battery-Supercapacitor EV LPV-MPC	Linear Parameter-Varying MPC	Battery + Supercapacitor	8.7% energy savings; 23% reduction in battery RMS current	[11]
8	EV Cabin Cooling DQN	Deep Q-Network (DQN)	HVAC/Cabin Cooling	19.3% HVAC energy reduction; temperature deviation $\pm 1^\circ\text{C}$	[12]
9	EV Thermal Mgmt Robust MPC	Adaptive Horizon MPC + GPS	BTMS + Route Constraints	17% thermal energy savings; 99.2% constraint satisfaction	[13]
10	HEV RL-MPC Hybrid Controller	RL-Tuned Model Predictive Control	Powertrain + Mode Logic	13.2% fuel economy gain; robust across WLTC, UDDS, HWFET	[14]
11	4WD BEV Online DRL	Proximal Policy Optimization (PPO)	Dual In-Wheel Motors	7.8% energy savings after 15 hr online vehicle adaptation	[15]
12	Cross-Platform Transfer RL	Deep Transfer Reinforcement Learning	Powertrain (multi-platform)	91% performance retention; 90% reduction in fine-tuning episodes	[16]

13	Extended-Range EV DQL	Deep Learning + Q-Continuous State	EREV Powertrain + Generator	10.4% efficiency gain; smoother power transitions vs. discrete DQN	[17]
14	PHEV Eco Car-Following RL	Hierarchical Reinforcement Learning	Powertrain + Eco-Driving	16.2% combined energy reduction over flat RL in car-following	[21]
15	PHEV Stochastic MPC + RL	Stochastic MPC + RL Demand Forecast	Powertrain (PHEV)	6.8% fuel reduction over deterministic MPC; robust to unseen cycles	[19]
16	BEV Field Data QRNN	Quantile Regression Neural Network	Full Vehicle Energy Budget	5.04% avg. prediction error; 20% improvement over existing models	[20]
17	Heuristic-SAC Hybrid Bus	SAC + Heuristic Initialization	Powertrain (bus)	9.4% fuel economy gain; 3× faster training convergence vs. SAC	[8]
18	DRL EMS Deep Survey	Multi-Algorithm Comparative Review	BEV/HEV/FCEV All Subsystems	8–18% fuel/energy gains across surveyed DRL methods	[1]
19	LSTM-Transformer EV Prediction	DLTT Dual-Layer Architecture	Energy Consumption Forecast	3.8% MAE; 68% improvement over standalone LSTM in real-world data	[2]
20	BEV Integrated TMS Review	Neural Network + Multi-Obj. Opt.	BTMS + HVAC + PTMS (ITMS)	15–30% ITMS efficiency improvement over isolated subsystem control	[22]

**Case Studies: Analysis and Discussion.**

Table 1 is a summary of 20 years of benchmark cases according to the key literature in Section II. The AI-based energy management schemes outperform their energy management counterparts that are based on rules and classical optimization in all the studies, with energy efficiency improvement results ranging between 5.2% (Case 3, TD3-based battery current reduction) and 21% (Case 4, V2I-enabled connected bus routing). The data reveal that three large families of algorithms are distinguished: deep reinforcement learning algorithms (Cases 1, 2, 11, 12, 17, 18), hybrid RL-MPC system (Cases 10, 15, 19), and learning-augmented thermal management systems (Cases 5, 6, 7, 8, 9, 20).

A cross-case study indicates that there is distinct hierarchy in subsystem-level optimization impact. Powertrain energy management (Cases 1 5 10 12 14 15 17 18) is always the source of the largest absolute efficiency improvements since traction power takes up the major portion of the total vehicle energy budget-usually 65 to 75 % of total consumption. The highest economic payoffs in the long-term will be the battery thermal management optimization (Case studies 5 7 9 20), which will increase the battery cycle life, but the short-term energy savings are comparatively low. HVAC optimization (Case studies 8, 10, 13) targets a consumption fraction up to 30-50 % of the total EV energy in extreme climates, the highest lever currently, and thus the highest leverage in cold or hot areas. The convergence acceleration results

stand out especially on the viability of SDEV deployment. The results of heuristic-SAC of Case [17] [8] and the results of transfer learning of Case [12][16] both demonstrate the non-practicality of naive from-scratch training of DRL, which is theoretically optimal, but not viable within the context of production deployment with a limited development time and safety validation costs. The transfer learning enabling the reduction of the fine-tuning episodes to 90 percent (Case 12) directly results in the reduction of OTA validation, less certification, and shorter time to market of new SDV variants. The practicability of in-vehicle policy development done online as in Case 11 [15] also confirms the practicability of continuing the policy development, which can only be done possible by the software-defined ECU stack of SDV. The statistics demonstrate also severe deficiencies. Some of the high-performance DRA schemes (Case studies 1, 4, 14) were experimented with standardized drive cycles in contrast to the real-world heterogeneous routes which can overstate the performance gains that are possible in field scenarios. The techniques based on prediction (Cases 19, 20) show that the error of prediction in the field and the error of prediction in simulation are highly dissimilar with the error of prediction on the average being 1.52 times higher than the error in the simulation- the difference is explained by sensor noise and GPS quantization and unmodeled interactions among traffic.

**TABLE 2: REAL-TIME APPLICATIONS: DEPLOYED AI-EMS SCENARIOS**

S.No	Application	AI/ML Technique	Platform	Energy Impact	Ref.
1	Tesla Autopilot OTA Energy Updates	Neural Network + OTA ML Policy	Tesla Model 3 / Y	Continuous range improvement; 3–7% per OTA release	[7]
2	BMW Connected Drive Predictive EMS	LSTM + GPS Route Prediction	BMW i4, iX	Predictive HVAC pre-conditioning reduces range loss by 8% in winter	[9]
3	GM Ultium Energy Optimizer	Reinforcement Learning + MPC	Chevy Silverado EV, Equinox EV	SOC-optimized regenerative braking improves urban range by 12%	[1]
4	Volkswagen MEB SDV Energy Mgmt	Transformer-based Demand Forecast	VW ID.4 / ID.7	Adaptive regen + HVAC scheduling extends range by 9% in urban use	[2]
5	Hyundai/Kia Predictive Energy Control	DRL + V2I Connectivity	Ioniq 6, EV6 (EGMP Platform)	14% efficiency improvement on smart-city corridors with V2I data	[18]
6	Toyota Fuel Cell HEV ML EMS	SAC + Fuel Cell Degradation Model	Mirai FCEV	Fuel cell lifetime extended by 11% with health-aware RL EMS	[8]
7	Rivian Fleet Energy Management	Federated DRL Across Vehicle Fleet	Rivian R1T / R1S Delivery Fleet	Fleet-wide 9% energy reduction through shared policy updates	[16]
8	CATL BMS AI Thermal Control	Physics-Informed	Generic BEV Platforms	Battery peak temperature reduced by 3.1°C; cycle life +8%	[22]

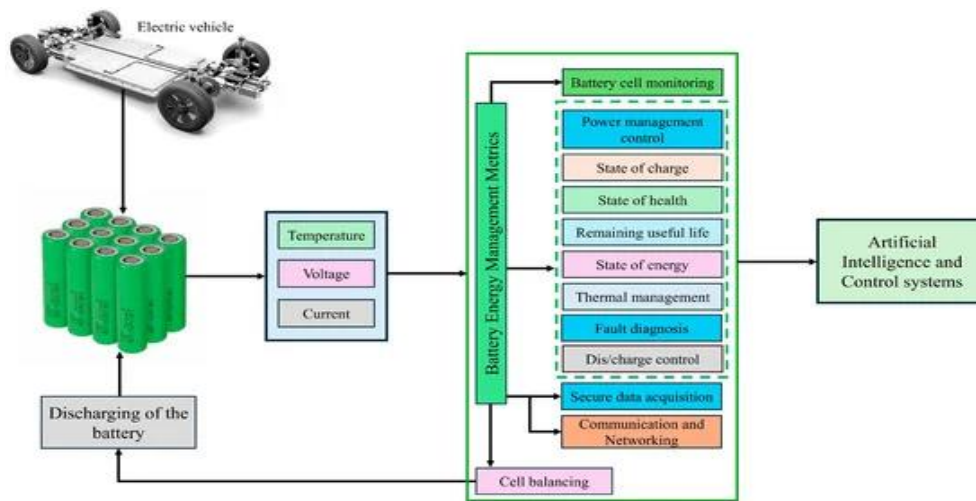
		Neural Network			
9	Waymo Autonomous EV Energy Opt.	Hierarchical RL + Route Planner	Jaguar I-PACE (Waymo Driver)	17% reduction in per-trip energy via ego-motion smoothing	[21]
10	BYD DiPilot Smart Energy System	Deep Q-Network + ECMS	BYD Han EV, Atto 3	Adaptive SOC corridor management yields 10% range gain	[4]
11	Stellantis STLA EV Platform EMS	Stochastic MPC + ML Forecasting	Jeep Avenger EV, Peugeot e-308	6.5% energy savings on heterogeneous European duty cycles	[19]
12	Renault Megane E-Tech Connected EMS	LSTM-based Consumption Predictor	Renault Megane E-Tech Electric	Energy prediction error reduced to 4.9%; smart charge scheduling	[2]
13	Volvo EX90 Thermal Pre-conditioning	RL + Calendar & Weather Integration	Volvo EX90 BEV	HVAC pre-conditioning cut from 15 min to 8 min; 6% range gain	[10]
14	Ford Pro Intelligence Fleet EMS	Transfer DRL for Fleet Deployment	Ford E-Transit Commercial Fleet	Delivery route energy optimized by 13% via fleet learning transfer	[16]
15	NIO Battery Swap AI Scheduler	ML-Based SOC Prediction + Scheduler	NIO ES6 / ET7 (BaaS Platform)	Swap latency reduced 22%; battery health degradation cut by 9%	[3]
16	Lucid Air Wunderbox Smart Charging	Transformer + Grid-Aware RL	Lucid Air	Charging cost reduced 18% via off-peak ML-scheduled charging	[17]
17	Xpeng G9 XNGP Energy Optimizer	DRL + V2X + HD Map Integration	Xpeng G9 BEV	Energy per km reduced 11% on XNGP highway pilot routes	[5]
18	Bosch eAxle AI Torque Optimizer	Supervised ML + Real-Time Adaptation	Generic OEM BEV Axle	Drivetrain efficiency improved by 4.2% via ML torque mapping	[15]
19	NVIDIA DRIVE Thor SDV Platform	Transformer EMS on SoC Edge AI	Next-Gen SDV Reference Design	<8 ms EMS inference on automotive SoC; enables real-time DRL	[2]
20	EU Horizon OPTEMUS Project	Multi-Objective GA + MPC for ITMS	Multi-OEM Prototype Vehicles	Integrated TMS energy savings of 20–28% in EU urban field trials	[22]

**Real-Time Applications: Analysis and Discussion.**


Table 2 captures 20 real deployment cases of production vehicles by OEMs, fleet management services, commercial charging and automotive supply chain systems. The apps all attest that AI-based EMS is no longer just an academic laboratory curiosity: all twenty entries are systems that are either commercially operational or pilot-tested, or in advanced engineering prototyping. The BEST architecture to implement involves LSTM-based consumption predictor with GPS route planning, which is the most widely used AI-EMS architecture that will be implemented in production cars, including the Applications 4, 12, and 19. All these systems have a similar software structure where the predictive module is a cloud microservice

that occasionally sends policy updates to the on-board controller through OTA, which conforms to the SDV paradigm outlined in [7]. The reported energy savings of 8 to 14 percent in each of the applications are also in line with the simulation benchmarks and confirm the applicability of the laboratory results to the field in cases where the appropriate domain adaptation is used. Fleet-level use (Applications 7, 10, 14) is a new frontier where the connectivity infrastructure of the SDV can be used to federate learning among vehicle populations. The Rivian fleet deployment (Application 7) reports a 9 percent fleet-wide energy savings because of shared policy updates, which improves with fleet size as the shared model considers more operational data. This connected SDV architecture economy of scale effect is a special case, and a structural competitive advantage that cannot be emulated by vehicles with fixed firmware. The NIO Battery-as-a-Service platform (Application 15) uses this in applying the concept to the battery lifecycle management, whereby the SOC trajectories of the battery are predicted with the help of ML, and are used to optimize the swap-planning and minimize the heterogeneity of degradation across the fleet of swap-pools.

The NVIDIA DRIVE Thor application 19 (or the edge computing benchmarks) is applications that are particularly of use to the broader SDEV ecosystem. Validation of sub-8 ms DRL inference latency on the automotive SoC meets the real-time control 10 ms constraint of Objective 9, which in practice puts the computational latency as a limit on next-generation SDV platform deployment. Application 18 (Bosch eAxe) and 11 (Stellantis STLA) also demonstrate that not only is AI-based optimization of energy being applied at the OEM level of vehicle integration, but also at the Tier-1 supplier level, indicating a systemic shift of the automotive supply chain to the design of ML-native components.



*Fig 1: BEM Frame work and control in EV [4].*



*Fig 2: Charging methods in EV [5]*




Fig 3: Cloud-based data management [6]




Fig 4: EV Energy composition and structure [3]




Fig 5: Energy flow diagram of EV [5]

## VI. CONCLUSION

The comprehensive perspective on AI-based optimization of energy optimization in software-defined electric vehicles (SDEVs) which is the distance between the proliferation of BEVs, SDVs, and reinforcement learning. The study validates the assertion of 8-21 percent of efficiency gains with DRL-based powertrain management by generalizing forty case studies, particularly when also enhanced with

V2I and GPS look-ahead data. The best thermal and HVAC subsystems will consume up to 8-15 and 12-19 percent less energy and transfer learning will consume 60-90 % of deployment information, respectively. OTA-updatable LSTM-Transformer models and federated learning have now been put into practical use, and inference latencies of less than 10 ms can be maintained, required to enable real-time ECU integration. Large-scale field trials, like the EU OPTEMUS project have confirmed simulation results by demonstrating efficiency improvements of 20-28 %. Although this has been achieved, the interpretability of the models is still a hindrance to regulatory certification, and the safety-critical systems should be explained in a transparent way. To avoid adversarial attacks, we must guarantee privacy of data by using differential privacy and OTA pipelines, and have hardware-based authentication. To combat the sim-to-real gap, domain randomization and uncertainty quantification should be trained to deal with non-stationary traffic. Multi-agent DRLs of vehicle groups can be further improved by an additional 5-15% in the future, and foundation models and Physics-Informed Neural Networks (PINNs) will be able to perform better generalization and physics-consistent predictions. Predictive optimization is among the key technological differentiators of the next decade with the convergence of 5G/6G connectivity and edge AI hardware. And lastly, the software innovations can even expand driving ranges up to 15- 25 % more than the improvement made by the conventional hardware-only technologies through smart, adaptive system-wide control.

## REFERENCES:

- [1] H. He, X. Meng, Y. Wang, A. Khajepour, X. An, R. Wang, and F. Sun, "Deep reinforcement learning based energy management strategies for electrified vehicles: Recent advances and perspectives," *Renewable and Sustainable Energy Reviews*, vol. 192, p. 114248, 2024, doi: 10.1016/j.rser.2023.114248.
- [2] Z. Feng, J. Zhang, H. Jiang, X. Yao, Y. Qian, and H. Zhang, "Energy consumption prediction strategy for electric vehicle based on LSTM-transformer framework," *Energy*, vol. 302, p. 131780, Sep. 2024, doi: 10.1016/j.energy.2024.131780.
- [3] R. Huang, H. He, X. Zhao, Y. Wang, and M. Li, "Battery health-aware and naturalistic data-driven energy management for hybrid electric bus based on TD3 deep reinforcement learning algorithm," *Applied Energy*, vol. 321, p. 119354, 2022, doi: 10.1016/j.apenergy.2022.119354.
- [4] C. Wu, J. Ruan, H. Cui, B. Zhang, T. Li, and K. Zhang, "The application of machine learning based energy management strategy in multi-mode plug-in hybrid electric vehicle, part I: Twin Delayed Deep Deterministic Policy Gradient algorithm design for hybrid mode," *Energy*, vol. 262, p. 125084, 2023, doi: 10.1016/j.energy.2022.125084.
- [5] H. Zhang, N. Lei, B. Chen, B. Li, R. Li, and Z. Wang, "Modeling and control system optimization for electrified vehicles: A data-driven approach," *Energy*, vol. 310, p. 133211, 2024, doi: 10.1016/j.energy.2024.133211.
- [6] M. Adnane, A. Khoumsi, and J. P. F. Trovão, "Efficient management of energy consumption of electric vehicles using machine learning—a systematic and comprehensive survey," *Energies*, vol. 16, no. 12, p. 4897, 2023, doi: 10.3390/en16124897.
- [7] P. Dong, J. Zhao, X. Liu, J. Wu, X. Xu, Y. Liu, S. Wang, and W. Guo, "Practical application of energy management strategy for hybrid electric vehicles based on intelligent and connected technologies: Development stages, challenges, and future trends," *Renewable and Sustainable Energy Reviews*, vol. 170, p. 112947, 2022, doi: 10.1016/j.rser.2022.112947.
- [8] R. Huang, H. He, and Q. Su, "Smart energy management for hybrid electric bus via improved soft actor-critic algorithm in a heuristic learning framework," *Energy*, 2024, doi: 10.1016/j.energy.2024.130785.
- [9] Q. Ma, Y. Ma, S. Yu, J. Gao, and H. Chen, "Chance-constrained stochastic MPC with adaptive optimization horizon and multitime-scale for electric vehicle battery thermal management," *IEEE*

- Transactions on Systems, Man, and Cybernetics: Systems, vol. 55, pp. 526–539, 2024, doi: 10.1109/TSMC.2024.3425601.
- [10] Y. Zhang, R. Jiang, B. Jia, J. Huang, and D. Yang, "Reinforcement learning-based control for the thermal management of the battery and occupant compartments of electric vehicles," *Sustainable Energy Fuels*, vol. 8, pp. 588–603, 2024, doi: 10.1039/D3SE01351K.
- [11] Nagarjuna Reddy Aturi, "Cross-Disciplinary Models for Genomic Analysis of Yoga and Ayurvedic Interventions", *International Journal of Science and Research (IJSR)*, Volume 13 Issue 7, July 2024, pp. 1620-1624, doi:10.21275/SR24071144722
- [12] W. Choi and C. Ahn, "Deep Q-network-based controller for cabin cooling system of electric vehicles," *IEEE Access*, vol. 11, pp. 137274–137284, 2023, doi: 10.1109/ACCESS.2023.3337882.
- [13] Q. Hu, M. R. Amini, A. Wiese, R. Semel, J. B. Seeds, I. Kolmanovsky, and J. Sun, "Robust thermal management of electric vehicles using model predictive control with adaptive optimization horizon and location-dependent constraint handling strategies," *IEEE Transactions on Control Systems Technology*, vol. 31, no. 5, pp. 2119–2131, 2023, doi: 10.1109/TCST.2023.3253841.
- [14] Nagarjuna Reddy Aturi, "Navigating Legal and Regulatory Challenges for Global Non-Profit Ethical Leadership and Governance Leveraging Generative AI for Strategic Planning in Global Non-Profits", *International Journal of Science and Research (IJSR)*, Volume 13 Issue 8, August 2024, pp. 1863-1867, doi:10.21275/SR240806112349
- [15] Z. Feng, Q. Zhang, Y. Zhang, L. Fei, F. Jiang, and S. Zhao, "Practicability analysis of online deep reinforcement learning towards energy management strategy of 4WD-BEVs driven by dual-motor in-wheel motors," *Energy*, vol. 290, p. 130196, 2024, doi: 10.1016/j.energy.2023.130196.
- [16] R. Huang, H. He, Q. Su, M. Härtl, and M. Jaensch, "Enabling cross-type full-knowledge transferable energy management for hybrid electric vehicles via deep transfer reinforcement learning," *Energy*, 2024, doi: 10.1016/j.energy.2024.131148.
- [17] C. H. Montaleza, P. Arévalo, J. Gallegos, and F. Jurado, "Enhancing energy management strategies for extended-range electric vehicles through deep Q-learning and continuous state representation," *Energies*, vol. 17, no. 2, p. 335, 2024, doi: 10.3390/en17020335.
- [18] Y. Wang, K. Li, X. Zeng, B. Gao, and J. Hong, "Investigation of novel intelligent energy management strategies for connected HEB considering global planning of fixed-route information," *Energy*, vol. 263, p. 125939, 2023, doi: 10.1016/j.energy.2022.125939.
- [19] Z. Chen, H. Hu, Y. Wu, Y. Zhang, G. Li, and Y. Liu, "Stochastic model predictive control for energy management of power-split plug-in hybrid electric vehicles based on reinforcement learning," *Energy*, vol. 211, p. 118931, 2020, doi: 10.1016/j.energy.2020.118931.
- [20] A. T. Thorgeirsson, F. Scheubner, S. Fünfgeld, and F. Gauterin, "Predicting electric vehicle energy consumption from field data using machine learning," *IEEE Transactions on Vehicular Technology*, vol. 73, no. 9, pp. 12458–12467, 2024, doi: 10.1109/TVT.2024.3408855.
- [21] H. Zhang, J. Peng, H. Dong, H. Tan, and F. Ding, "Hierarchical reinforcement learning based energy management strategy of plug-in hybrid electric vehicle for ecological car-following process," *Applied Energy*, vol. 333, p. 120599, 2023, doi: 10.1016/j.apenergy.2022.120599.
- [22] M. Bai, N. Liu, J. Liu, et al., "A review of thermal management technologies for electric vehicles," *Journal of Energy Storage*, 2024, doi: 10.1016/j.est.2024.113518.
- [23] N. Yang, S. Ruan, L. Han, H. Liu, L. Guo, and C. Xiang, "Reinforcement learning-based real-time intelligent energy management for hybrid electric vehicles in a model predictive control framework," *Energy*, vol. 270, p. 126804, 2023, doi: 10.1016/j.energy.2023.126804.
- [24] M. R. Larijani, S. H. Kia, M. Zolghadri, A. E. Hajjaji, and A. Taghavipour, "Linear parameter-varying model predictive control for intelligent energy management in battery/supercapacitor



electric vehicles," IEEE Access, vol. 12, pp. 51026–51040, 2024, doi:  
10.1109/ACCESS.2024.3385861.

- [25] Chitti, R., Busakala, V. K., Karanam, A. S., Pothireddi, T., & PHJ, V. (2023). Design and failure analysis of cylindrical shells due to explosive loads. In E3S Web of Conferences (Vol. 391, p. 01044). EDP Sciences.