

Advanced Reinforcement Learning based Thermal Management Strategy for Battery Electric Vehicles

Tomislav Bukić*¹,
¹AVL-AST d.o.o., Croatia,
Amr Mousa², Milan Živadinović², Bernhard Peischl², Harisyam Manda², Wolfgang Koenig²,
Armin Traussnig²
²AVL List GmbH, Austria

KEYWORDS – reinforcement learning, adaptive control strategies, thermal management, artificial intelligence, battery electric vehicles

The vehicle control systems domain encounters an increasing number of parameters and calibration targets considering the emerging technologies such as connected vehicles and automated driving. Accordingly, the calibration processes for such systems have become more complex and thus error prone and tedious. Moreover, the derived control policies are not easily transferable between different vehicle configurations, hence, the calibration effort is increasing dramatically with each configuration change. Therefore, the reduction of such efforts needed to setup the control policy is inevitable to further reduce cost of drivetrain development. The proposed methodology therefore is an important means to make BEVs (Battery Electric Vehicles) more attractive for car buyers.

The fast development of the Artificial Intelligence (AI) domain is opening the door to numerous opportunities and applications in the automotive industry. We utilize Reinforcement Learning (RL) techniques to design appropriate control strategies for different vehicle systems, thus improving the conventional approaches and reducing the development effort. Combining the expertise in simulation and big data, we propose a cloud-based solution that runs a high-fidelity simulation to train, test and deploy the thermal management control strategy for a fleet of BEVs. The benefits, among others, are that the RL-based control policies can be designed more rapidly and run more efficient than the traditional rule-based approaches. After deploying the initial model trained against the simulation, we have the capability of collecting data from a fleet of vehicles operating with the latest control strategy. Using collected data, we iteratively train and customize the strategy throughout the operation time.

We have tested the abovementioned framework on the use case of cabin heating mode selection for BEVs. Our RL agents are trained and evaluated in a model-in-the-loop simulation environment. The policy evaluation is based on the agents' performance on representative vehicle test measurements (drive cycles). The metrics are selected to quantify the energy efficiency and comfort individually, as well as aggregated to enable a fair comparison. Notably the trained agents achieved better results than the original control policy on most of the individual metrics and significantly better results on the aggregated metric.

At this moment, our framework is tested against simulated vehicle fleet. The first reasonable research question is if the trained control system can be directly transferred to the real vehicle, or whether additional adjustments must be performed to achieve the needed flexibility. Another open question refers to the adequate combination of RL algorithms to achieve even better performance on telemetry data from a connected fleet. Time will tell if the idealized case with continuing on-policy training, or the more complex case using the policy-agnostic offline algorithms, will provide the stronger solution.

The main technical contribution of our work is the use case agnostic framework which iteratively improves a conventional rule-based control strategy. Following the automotive V-model, the design-, implementation-, and testing -phase is strictly separated from the in-use phase of a vehicle function. To leverage historic data from the in-use phase, our framework disrupts this classical model and embeds the DevOps and ML-Ops practices into the automotive engineering process.

In this article we suggest a framework which automates the complex and time-consuming creation process of control strategies. The trained control policies provide better results and allow for a continuous improvement after they are finally deployed on a fleet of vehicles.