

Simultaneous Material and Shape Optimization of EV Battery Plates Using Surrogate Models Trained on Limited Samples

Hiroyuki Saito ¹⁾ Takayuki Michishita ¹⁾ Kotaro Kawajiri ¹⁾

¹⁾ AIZOTH Inc.

1-5-7 Azuma, Tsukuba, Ibaraki, 305-0031, Japan (E-mail: hiro.saito@aizoth.com)

KEY WORDS: EV and HV systems, battery technology, aluminum alloy, heat and temperature management, stiffness, mass reduction, design optimization, Surrogate model, Chain analysis [A3]

Electric vehicle (EV) battery cooling plates must simultaneously satisfy thermal management and structural (mass, stiffness, strength) performance, while keeping cost low enough for practical adoption. Because obtaining a globally optimal combination of material selection and shape design requires an enormous number of evaluations, material and shape are rarely optimized jointly in practice, and a design that is optimal in both material and structure is rarely achieved under realistic development schedules.

This study proposes a chained surrogate model that connects a neural network predicting aluminum alloy properties from composition and temper (front stage) with a surrogate model predicting thermal and structural performance from channel geometry (back stage), enabling consistent prediction from material composition to component performance. Using this framework, an EV battery cooling plate is treated as a case study, addressing material, shape, and cost simultaneously. A multi-objective genetic algorithm was used to optimize channel width, pitch, and thickness together with aluminum alloy composition and temper, for five structural alloy series (2xxx–7xxx), under two cell-expansion pressures ($p = 0.1$ and 0.2 MPa). Total cost (material cost plus forming cost) was evaluated analytically from the existing optimization results from the back-stage model.




Fig. 1 Pareto-optimal solutions at $p = 0.1$ (a) and 0.2 MPa (b), and the resulting cost trade-off (c)

Fig. 1(a)(b) show the combined Pareto-optimal solutions (minimizing thermal resistance R_{th} and mass, maximizing bending stiffness EI) at $p = 0.1$ and $p = 0.2$ MPa. At $p = 0.1$ MPa, all five alloy series contribute comparably to the Pareto front, whereas at $p = 0.2$ MPa, only the high-strength, low-conductivity 2xxx and 7xxx series remain feasible, accounting for 99% of the front. Despite this shift, the range of achievable mass, R_{th} , and EI remains nearly the same at both pressures. What changes with load is not the achievable performance itself, but which alloys can achieve it. This invariance, however, only holds when cost is excluded. Fig. 1(c) shows the cheapest feasible cost at each pressure, obtained by applying the cost model to the existing Pareto results. As the load increases from $p = 0.1$ to 0.2 MPa, the cheapest alloy series (6xxx) becomes infeasible, and the minimum achievable cost rises by about 10%. In other words, a load-driven switch in alloy class raises cost while leaving the achievable performance unchanged. This is a hidden trade-off, invisible if only thermal and structural metrics are considered.

A partial derivative (PaD)-based contribution analysis further showed that, while the back-stage output (mass) is mainly affected by shape variables, the contribution of material variables is not negligible, supporting the need to couple material and shape optimization rather than treating them independently. These results demonstrate that a chained surrogate model can support simultaneous material–shape–cost optimization and reveal cost implications that a performance-only analysis would miss.