From Prompt to Protocol: Fast Charging Batteries with Large Language Models

Ge Lei[1], Ferran Brosa Planella[2], Sterling G. Baird[3], Samuel J. Cooper[1]

[1] Dyson School of Design Engineering, Imperial College London, London SW7 2AZ, UK
[2] University of Warwick, Coventry CV4 7AL, UK
[3] Acceleration Consortium, University of Toronto, Toronto, Canada

g.lei23@imperial.ac.uk

Abstract

Optimizing lithium-ion battery charging protocols is fundamentally challenging due to the vast design space of possible current trajectories and the slow, costly, and non-differentiable nature of experimental evaluations [1,2]. Conventional approaches typically constrain the protocol representation to predefined analytical forms or multi-step constant current (CC) segments to make the optimization tractable [1]. While effective, such restrictions limit expressiveness and may prevent discovery of higher-performing strategies.

In this work, we introduce two gradient-free large language model (LLM)-driven closed-loop optimization frameworks, Prompt-to-Optimizer (P2O) and Prompt-to-Protocol (P2P), that expand the accessible search space beyond human-designed parameterizations. P2O adopts a two-loop evolutionary structure in which the LLM generates lightweight neural-network architectures (<35 trainable parameters) that define adaptive charging policies, while an inner Sparse Axis-Aligned Subspace Bayesian Optimization (SAASBO) loop performs sample-efficient parameter tuning under limited evaluation budgets [3]. In contrast, P2P follows a single-loop paradigm where the LLM directly synthesizes executable charging protocols, jointly determining structural form and numerical parameters without a dedicated inner optimizer, consistent with recent evaluator-in-the-loop LLM program synthesis approaches [5].

Both methods are fully gradient-free and therefore directly applicable to real experimental settings. We validate the frameworks using PyBaMM-based electrochemical simulations [4], built on a Doyle–Fuller–Newman (DFN) model [6] augmented with thermal dynamics and degradation submodels, including reaction-limited SEI growth, stress-driven loss of active material, particle cracking, and porosity change.

Three progressively complex case studies are conducted: (1) predefined heat-profile recovery to isolate outer-loop architecture evolution; (2) constant-heat protocol search under non-differentiable feedback to assess inner-loop optimization; and (3) adaptive degradation-aware fast charging under fixed charging-time constraints.

In the third realistic fast-charging scenario, both LLM-driven approaches outperform a state-of-the-art multi-step CC baseline derived from prior closed-loop optimization literature [1]. Specifically, P2O and P2P improve final state-of-health (SOH) by approximately 4.2% relative to the baseline under matched evaluation budgets.

Notably, P2P achieves competitive performance without an inner-loop optimizer by encoding voltage-safety constraints directly in natural language, demonstrating that semantic constraint injection can guide efficient search in high-cost regimes. P2O further improves performance through evolutionary refinement when additional computational budget is available.

The results indicate that LLM-driven optimization can remain effective even when each evaluation is expensive and slow, without requiring thousands of iterations typical of algorithm-evolution studies. More broadly, the proposed framework illustrates how generative models can expand functional priors, integrate qualitative design preferences, and operate beyond rigid hypothesis classes imposed by manual parameterization. This capability is particularly valuable for battery-protocol design, where safety constraints, degradation mechanisms, and multi-physics coupling make traditional gradient-based optimization impractical. The methodology is readily transferable to experimental workflows by replacing simulation feedback with laboratory measurements, offering a scalable pathway toward AI-assisted materials and energy-system optimization.

References

[1] Attia P. M., Grover A., Jin N., Severson K. A., Markov T. M., Liao Y.-H., Chen M. H., Cheong B., Perkins N., Yang Z., et al., Nature, 578 (2020) 397–402.

[2] Keil P., Jossen A., Journal of Energy Storage, 6 (2016) 125–141.

[3] Eriksson D., Jankowiak M., Proceedings of the Thirty-Seventh Conference on Uncertainty in Artificial Intelligence, Proceedings of Machine Learning Research, 161 (2021) 493–503.

[4] Sulzer V., Marquis S. G., Timms R., Robinson M., Chapman S. J., Journal of Open Research Software, 9(1) (2021) 14.

[5] Romera-Paredes B., Barekatain M., Novikov A., Balog M., Kumar M. P., Dupont E., Ruiz F. J. R., Ellenberg J. S., Wang P., Fawzi O., Kohli P., Fawzi A., Nature, 625(7995) (2024) 468–475.

[6] Doyle M., Fuller T. F., Newman J., Journal of The Electrochemical Society, 140(6) (1993) 1526–1533.

Figures

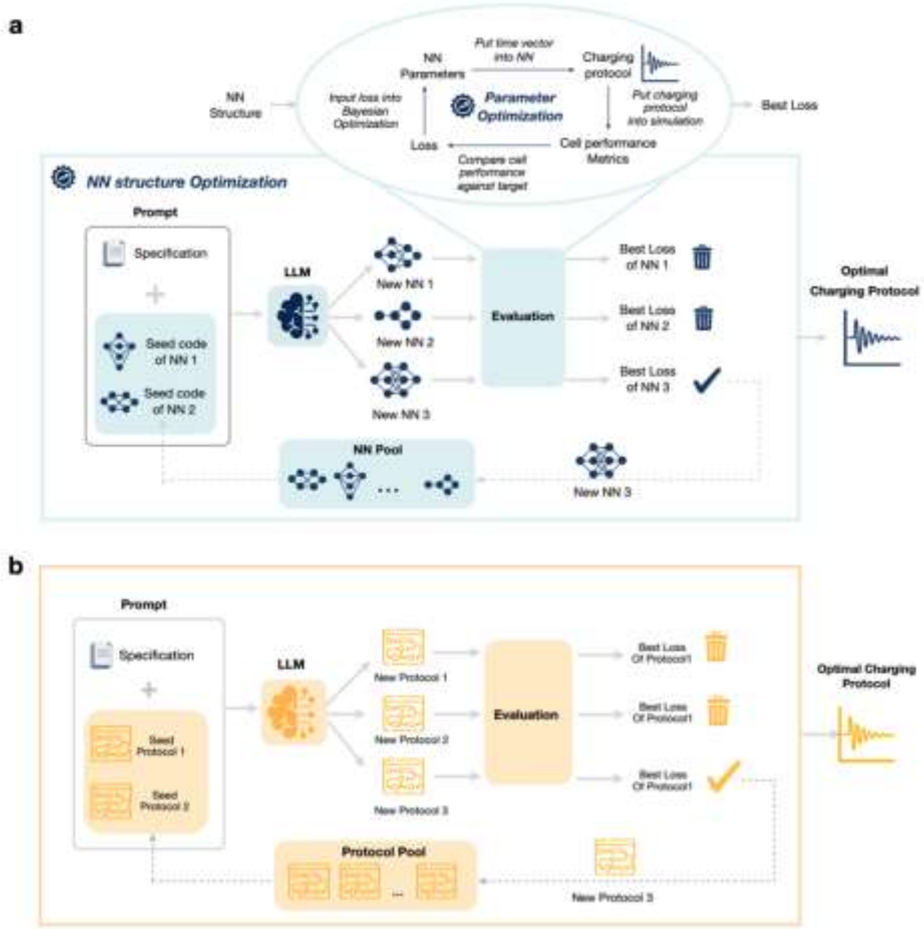

Figure 1. Overview of the proposed LLM-driven optimization frameworks. (a) P2O (two loops): LLM-guided neural network generation and refinement with a SAASBO-based inner optimization loop. (b) P2P (single loop): Direct LLM generation of explicit charging protocols without inner-loop optimization.

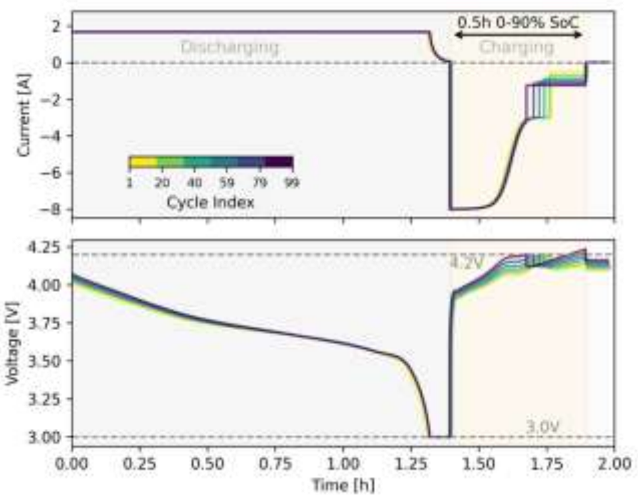

Figure 2. Best P2O protocol across cycles, showing earlier voltage-limit entry with degradation.