Self-supervised battery health estimation from partial fast-charging data: evaluating JEPA

Keywords. Battery state of health, time series, self-supervised learning, Joint-Embedding Predictive Architecture, limited labels

Supervisor: Shankar A. Deka, shankar.deka@aalto.fi.

Advisor: Hassan Rouhi, hassan.rouhi@aalto.fi.

Research Problem Statement. Battery health estimation from short charging records could reduce dependence on frequent capacity measurements. This project asks whether a compact Joint-Embedding Predictive Architecture (JEPA) learns useful health representations with limited labels and changes in charging policy, and whether conditioning the predictor on the recorded current contributes anything beyond pretraining itself.

Battery pretraining and learning from incomplete measurements already exist [1]. Joint-embedding predictive architectures learn by predicting representations of hidden inputs rather than reconstructing them [2], and have been adapted to time series [3]. The contribution will be a controlled comparison with reconstruction and supervised learning, including normalization and computational-budget controls. Better performance from JEPA is a hypothesis to test.

Three features of the problem make it difficult. Only a short, partial segment of a charging event is typically observed, so the model must work from incomplete records. Reliable capacity measurements are far rarer than the charging data itself, so health labels are scarce. And the charging current shapes the measured voltage response, so operating conditions and ageing are entangled in the same signal: a model can easily learn to recognise the charging protocol rather than the state of the cell.

The project will investigate whether a compact JEPA, conditioned on the recorded current, learns representations that separate ageing from operating conditions, and whether those representations remain useful when the charging policy changes. Public fast-charging cycling data [4] will be used, with health labels withheld to emulate label scarcity. The evaluation is a controlled comparison against reconstruction-based and supervised learning under matched conditions.

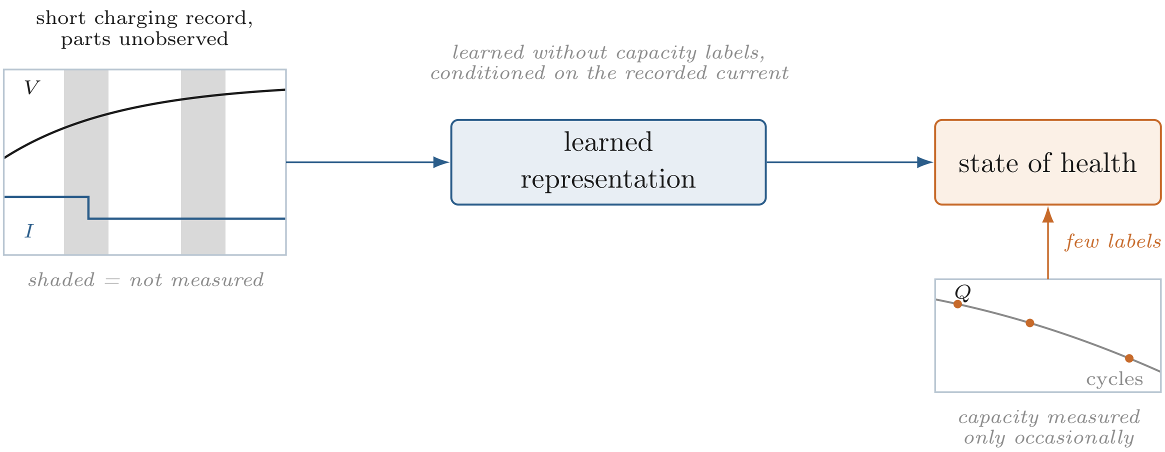

Figure 1: The setting. Only part of a charging record is observed, and capacity is measured far less often than the charging data itself. A representation is learned from the unlabelled records; a small health-estimation head is fitted from the few available capacity labels.

The outcome will be a reproducible comparison, including an analysis of situations where latent prediction does not improve performance. One dataset and one prediction task form the complete core project.

Expected workload: The work will include:

1. Review the closest battery and JEPA methods; audit raw signals, cell identities, batch differences and fast-charging windows. Define cell-disjoint sets with held-out charging-policy groups

where the data support a meaningful comparison.

2. Establish feature-based and supervised baselines, including a supervised normalization control [5] that preserves absolute current and voltage-level information.

3. Implement a compact current-conditioned JEPA, a conditioning ablation and a matched reconstruction model. Share preprocessing and encoder size; cap pretraining compute and tuning effort. Log embedding variance and effective rank from the first run.

4. Reuse frozen pretrained encoders at three simulated label budgets, initially 5%, 10% and 100%, the last serving as a fully supervised reference. Report cell-averaged health-estimation error, variation across cells and three training seeds, batch effects, and computational cost. All methods are evaluated on identical cells, and the primary comparison is the paired per-cell difference reported with a confidence interval rather than a difference of group means.

Feasibility and completion. Approximately six months of full-time work, using Aalto's Triton computing cluster. Fix the data protocol by week 4 and establish supervised/reconstruction baselines by week 6. At week 10, continue JEPA experiments only if a stable, non-collapsed encoder exists; otherwise complete the reconstruction-versus-supervised study and document the JEPA implementation limits. A stable JEPA with no performance advantage remains a valid result.

Optional extension: After completing the core results, investigate either a second segment length or a small HUST transfer experiment. Neither is required for completion.

Prerequisite: Good knowledge of machine learning, linear algebra and optimisation; Python and PyTorch experience. Time-series experience is useful. Battery-specific knowledge can be acquired during the project, supported by regular methodological supervision.

Type of work: Approximately 30% methods and analysis, 70% implementation. Writing proceeds throughout the project. Foundation-model training from scratch, RUL prediction, new experiments and controller development are outside scope.

References

[1] J. Li et al., "LiPM: Foundation Model for Lithium-Ion Battery Analysis," KDD, 2025. DOI: 10.1145/3711896.3737027.

[2] M. Assran et al., "Self-Supervised Learning from Images with a Joint-Embedding Predictive Architecture," CVPR, 2023. arXiv:2301.08243.

[3] S. Ennadir, S. Golkar and L. Sarra, "Joint Embeddings Go Temporal," 2025. arXiv:2509.25449.

[4] K. A. Severson et al., "Data-driven prediction of battery cycle life before capacity degradation," Nature Energy 4, 383–391, 2019. DOI: 10.1038/s41560-019-0356-8. Data: data.matr.io/1.

[5] T. Kim et al., "Reversible Instance Normalization for Accurate Time-Series Forecasting against Distribution Shift," ICLR, 2022. The normalization control is adapted to health regression, not copied unchanged from forecasting.