

Battery lifetime prediction across diverse ageing conditions with inter-cell deep learning

Received: 6 December 2023

Accepted: 12 December 2024

Published online: 15 January 2025

 Check for updates

Han Zhang^{1,2,4}, Yuqi Li^{1,3,4}, Shun Zheng¹✉, Ziheng Lu¹, Xiaofan Gui¹, Wei Xu² & Jiang Bian¹✉

Accurately predicting battery lifetime in early cycles holds tremendous value in real-world applications. However, this task poses significant challenges due to diverse factors influencing complex battery capacity degradation, such as cycling protocols, ambient temperatures and electrode materials. Moreover, cycling under specific conditions is both resource-intensive and time-consuming. Existing predictive models, primarily developed and validated within a restricted set of ageing conditions, thus raise doubts regarding their extensive applicability. Here we introduce BatLiNet, a deep learning framework tailored to predict battery lifetime reliably across a variety of ageing conditions. The distinctive design is integrating an inter-cell learning mechanism to predict the lifetime differences between two battery cells. This mechanism, when combined with conventional single-cell learning, enhances the stability of lifetime predictions for a target cell under varied ageing conditions. Our experimental results, derived from a broad spectrum of ageing conditions, demonstrate BatLiNet's superior accuracy and robustness compared to existing models. BatLiNet also exhibits transferring capabilities across different battery chemistries, benefitting scenarios with limited resources. We expect this study could promote exploration of cross-cell insights and facilitate battery research across comprehensive ageing factors.

With high energy densities and low production costs, lithium-ion batteries have become widely adopted across modern industries, driving the growth of renewable energy solutions and electric vehicles^{1–3}. Nevertheless, the capacity of lithium-ion batteries inevitably fades with cyclic operations due to their intrinsic electrochemical mechanisms⁴. Unexpected rapid degradation not only leads to poor user experiences, such as range anxiety for electric vehicles, but can also affect the operation of essential facilities, such as the stability of power grids. To proactively mitigate these side effects, accurately predicting battery lifetime in early cycles has been identified as a critical

task^{5–8}, where the lifetime is typically measured in cycle life, which is defined as the number of charge–discharge cycles until the capacity of a battery cell drops to 80% of its nominal capacity^{9,10}. Predicting battery lifetime in early cycles is rather challenging because numerous factors, including but not limited to cycling protocols, ambient temperatures and electrode materials, collectively influence the complex battery ageing process. Additionally, achieving comprehensive coverage of various ageing conditions is prohibitively costly and time-consuming, as testing a cell under a specific condition is a lengthy process.

¹Microsoft Research, Beijing, China. ²Institute for Interdisciplinary Information Sciences, Tsinghua University, Beijing, China. ³Present address: Department of Materials Science and Engineering, Stanford University, Stanford, CA, USA. ⁴These authors contributed equally: Han Zhang, Yuqi Li. ✉e-mail: shun.zheng@microsoft.com; jiang.bian@microsoft.com




Fig. 1 | Diverse ageing conditions covered in this study and complex degradation behaviours associated. **a**, We compare the broad coverage of different ageing factors in this study against typical datasets, including MATR⁹, HUST¹², SNL¹⁹ and CALCE^{13–15}. **b**, The varying capacity degradation behaviours over long- and short-term cycles.

Existing methods for battery lifetime prediction have been developed and validated under limited ageing conditions, such as testing only lithium-iron-phosphate (LFP) cathode materials and using a certain group of cycling protocols^{9–12}. Data characteristics under these restricted conditions affect influence a thorough examination of feature designs and modelling techniques, potentially limiting the success and generalization of their conclusions. It remains questionable whether these methods perform well across broad ageing conditions. Moreover, focusing on limited ageing conditions restricts the research development of leveraging historical data collected under different conditions. This limitation separates battery datasets emphasizing different ageing factors as isolated islands, hindering the development of general modelling approaches.

In this study, we introduce BatLiNet, a deep learning framework designed for reliably predicting battery lifetime across diverse ageing conditions, such as variations in cycling protocols, ambient temperatures and even battery chemistries. At its core, the framework introduces inter-cell learning that contrasts pairs of battery cells to discern lifetime differences, a significant leap from traditional models that focus solely on individual cells. By combining this approach with conventional single-cell learning, BatLiNet not only captures the unique degradation patterns of individual batteries but also places these patterns within a broader, comparative ageing landscape.

In our experiments covering various ageing complexities, BatLiNet consistently outperforms traditional models, delivering precise and robust lifetime predictions. A noteworthy observation is BatLiNet's ability to reduce the mean absolute percentage error (MAPE) of its single-cell learning counterpart by more than 40% on average, underscoring the significance of inter-cell learning. Additionally, BatLiNet exhibits intriguing capabilities to transfer learning across different battery chemistries. These findings underscore the importance of considering a broad range of ageing conditions in battery data modelling and also reveal the untapped potential of leveraging the connections across different ageing conditions.

Results

To gain a comprehensive understanding of various ageing factors and their influences on battery lifetime, we aggregate most publicly available datasets, including MATR⁹, HUST¹², CLO¹¹, CALCE^{13–15}, HNEI¹⁶, UL-PUR¹⁷, RWTH¹⁸ and SNL¹⁹. Figure 1 compares the ageing conditions examined in this study with typical datasets, highlighting the associated complex degradation behaviours. Supplementary Note 1 provides additional detailed comparisons across all datasets.

In this study, we derive five datasets reflecting different variations in ageing conditions. First, we inherit two evaluation sets of the MATR dataset, labelled MATR-1 and MATR-2, to maintain a consistent comparison with existing models. Besides, we consider the HUST dataset,

using similar LFP cells to MATR, to investigate the adaptation of current models to different cycling protocols. Furthermore, we include other datasets by merging all datasets into a MIX set because of limited cells under other conditions, different ambient temperatures, packing structures and cathode active materials, such as lithium nickel cobalt manganese oxide (NMC), lithium cobalt oxide (LCO) and lithium nickel cobalt aluminium oxide (NCA). We develop two prediction setups based on the MIX dataset: MIX-100 and MIX-20. MIX-100 examines a typical early prediction scenario where models must predict the 80% end-of-life point using only the first 100 cycles of data. MIX-20 presents a more challenging task, wherein models must forecast the number of cycles before the capacity degrades to 90% of the initial value using just the first 20 cycles. It is worth noting that batteries that reached their end-of-life prematurely during the early cycles were excluded from the study.

The BatLiNet framework

The inherent conflict between numerous ageing conditions and limited cell cycling coverage poses a significant challenge to predictive modelling. Existing models primarily focus on intra-cell learning, capturing early variations of a single cell to predict its long-term lifetime^{9,10}. However, this approach struggles with data scarcity due to limited cell samples available. In contrast, our BatLiNet framework, presented in Fig. 2, integrates inter-cell learning, which is essentially an auxiliary task of predicting lifetime differences given a pair of cells. We designate a cell with an unknown lifetime as the 'target cell' and use a fully cycled cell from existing databases as the 'reference cell'.

Figure 2a depicts the process of constructing effective features for intra- and inter-cell learning, where we use the voltage–capacity curve during the discharge stage as a typical cycle-level feature. Traditional approaches for intra-cell learning operate on a single cell and obtain effective intra-cell features by calculating the cycle-level feature differences between a pair of early cycles^{9,10}. We refer to these operations as calculating intra-cell differences. In contrast, we calculate inter-cell differences by contrasting the cycle-level features between a target cell and a reference cell in the same cycle. The inter-cell differences aim to represent the connections between cells cycled under different conditions.

Figure 2b illuminates the potential and challenges of both intra- and inter-cell learning across various ageing conditions. The upper plot illustrates a near-linear relationship between the standard deviation of intra-cell difference curves and the corresponding lifetime for different LFP cell groups. While the linearity is consistent with previous studies⁹, this correlation highlights the complexities introduced by diverse ageing conditions, as shown by the scattered points, outliers and distinct clusters due to additional cells undergoing varied protocols and temperatures. Conversely, the lower plot delves into inter-cell




Fig. 2 | An overview of the BatLiNet framework. a, The feature construction for intra- and inter-cell learning given the curve illustrating the relationship between Q (current capacity normalized by nominal capacity) and V (voltage in volts) as

a per-cycle feature map. **b**, The correlations between constructed features and prediction targets for both intra-cell (the upper half) and inter-cell (the lower half) learning. **c**, The data flow of BatLiNet.

learning by mapping the correlation between inter-cell differences and associated lifetime differences. Here, the same LFP cells serve as target cells, and a shared LCO cell acts as a reference. Despite the challenges posed by diverse ageing factors, a clear connection is evident. This suggests that a simple feature computed on inter-cell difference curves can differentiate lifetime differences, even with a reference cell from a different battery chemistry. Supplementary Note 2 further supports the feasibility of inter-cell learning within or across different battery chemistries with more data-driven evidence.

Moreover, these nonlinear patterns have indicated the necessity of leveraging deep learning for effective representation learning. Figure 2c concludes the data flow of BatLiNet. Given a target cell, we feed its cycle-level features into two separate branches. The first branch calculates intra-cell difference curves by comparing features in all early cycles with that of a reference cycle and uses a specific neural network f for intra-cell learning. In addition, the second branch includes a reference cell to obtain inter-cell difference curves and further feeds them into another neural network g for inter-cell learning. After the encoding procedures of f and g , the resulting representations are mapped to an aligned linear space, allowing us to leverage a shared linear layer to emit outputs for these two branches. The rationale for separated encoding followed by a shared linear layer is detailed in the Methods. During training, we can simply select any pair of cells from a training set to stimulate robust inter-cell modelling across various combinations

of ageing conditions. To ensure robust prediction during inference, we randomly select a batch of reference cells from the training set and use the average of outputs from intra- and inter-cell branches as the final prediction.

Figure 3 offers a comparison of BatLiNet with state-of-the-art baselines across different datasets, using two metrics: root mean squared error (r.m.s.e.) and MAPE. We can observe that basic linear models relying on hand-crafted features, such as the 'Var.', 'Dis.' and 'Full' models⁹, demonstrated commendable performance on their initial datasets. However, their performance degraded significantly when faced with datasets encompassing complex ageing conditions. This trend is also evident in some advanced statistical learning techniques applied to raw capacity–voltage curves¹⁰, such as ridge regression²⁰, partial least-squares regression²¹, principal component regression²² and support vector machine²³. An exception to this pattern is random forest²⁴, which consistently delivered relatively robust predictions, even on the MIX-100 and MIX-20 datasets. Despite this, random forest generally yielded suboptimal performances across all datasets under evaluation.

Compared with statistical methods, deep learning models operating in the feature space developed by ref. 10 are susceptible to high variability resulting from different random initializations, including multi-layer perceptron²⁵, long-short-term memory network²⁶ and convolutional neural network (CNN)²⁷. This may be due to the limited number of cells available for learning, combined with diverse ageing




Fig. 3 | Performance comparisons of BatLiNet against linear models along with statistical models and deep learning models. Linear models (Var., Dis. and Full) developed in ref. 9, and statistical models (ridge regression, partial least-squares regression, principal component regression, support vector machine, random forest) and deep learning models (multi-layer perceptron (MLP), long-short-term memory (LSTM), convolutional neural network (CNN)) from ref. 10. These comparisons span five datasets, MATR-1, MATR-2, HUST, MIX-100 and MIX-20,

and cover two evaluation metrics, r.m.s.e. and MAPE. For deep learning models, we run various experiments with eight unique initializations ($n = 8$ trials). The bar plot displays the mean and standard deviation values. The dashed line represents a simple baseline, using the average lifetime of training cells as a fixed prediction to indicate the difficulty of learning. PCR, principal component regression; PLSR, partial least-squares regression; RF, random forest; Ridge, ridge regression; SVM, support vector machine.

patterns and significant noise. BatLiNet also belongs to the deep learning family and uses the same CNN architecture for feature encoding. We observe that it has significantly mitigated the high variability associated with deep learning baselines, consistently delivering highly competitive results across diverse ageing conditions. Specifically, BatLiNet reduces the r.m.s.e. of the best-performing baseline (which varies by dataset) by 36.5, 6.8, 20.1, 27.4 and 40.1%, respectively. When compared with its single-cell learning counterpart, CNN, BatLiNet can reduce the average MAPE by up to 40%. It is worth noting that even for BatLiNet, the MAPE score on the MIX-100 and MIX-20 datasets is distinctly higher than that on previous datasets such as MATR-1 and MATR-2. We believe this highlights the significance of considering comprehensive ageing factors in battery data modelling.

Furthermore, we delve into investigating cross-chemistry transfer of BatLiNet. More specifically, we aim to confirm if BatLiNet can enhance the transfer from a resource-rich chemistry to other chemistries that have limited cells. This arrangement can highlight its potential in expediting precise lifetime predictions for future chemistries. In our experiments, we use 275 LFP cells as the rich-resource battery chemistry and aim to investigate how well we can leverage these LFP cells to develop models for the remaining 37 LCO cells, 22 NCA cells and 69 NMC cells. For consistent evaluation, we randomly sampled 21 LCO cells, 14 NCA cells and 53 NMC cells to form respective test sets. Then, from the remaining cells, we sampled 1, 2, 4, 8 and 16 cells as the available data for model training, corresponding to different low-resource conditions for each target battery chemistry. Figure 4 presents the overall comparisons of these models. We can observe that

in most cases, BatLiNet has demonstrated distinct error reductions over other learning models. For instance, direct learning cannot leverage information from LFP cells and thus struggle when the number of target cells is extremely rare. Parameter-based transfer learning relies on pretrained parameters to enable implicit transfer, which we also find could encounter generalization challenges in certain cases, such as transferring from LFP cells to only one NCA or LCO cell. In contrast, BatLiNet enables an explicit transfer via inter-cell learning, which delivers more robust and accurate predictions in most cases.

Finally, we explain the reference selection strategy in BatLiNet. Figure 5 presents extended results on the MIX-100 dataset. Figure 5a shows that the choice of reference cells for a target cell can significantly affect prediction errors, as seen from the gap between the best and worst choices across different test set cells. While selecting the optimal reference cell is challenging, we mitigate performance variations by sampling a batch of reference cells and using their median predictions, as shown in Fig. 5b. Figure 5c reveals that this batch sampling mechanism does result in a significant reduction in inference speed as the batch size increases exponentially. However, given the extensive battery cycling time, these differences in inference speed are negligible. We find that using a batch of 64 reference cells is effective in maintaining a reasonable inference speed yet delivering robust predictions.

Discussion

We introduce BatLiNet, a deep learning framework that integrates inter- and intra-cell learning to improve battery lifetime predictions across varied ageing conditions. Inter-cell learning complements




Fig. 4 | Performance comparison of BatLiNet when transferring from 275 LFP cells to limited NCA, LCO and NMC target cells. We simulate various low-resource scenarios by varying the number of available target cells for training (1, 2, 4, 8, 16). BatLiNet's inter-cell learning is compared to two baselines: a CNN model trained directly on target cells and a CNN model pretrained on LFP cells then fine-tuned. Each case and model is tested over eight unique initializations and data splits ($n = 8$ trials), with bar charts showing average error and error bars indicating standard deviation across runs.

intra-cell learning by enabling knowledge transfer across different ageing conditions.

This approach has potential applications in other critical prediction tasks, such as estimating battery state of charge and health^{5,28}, adapting to diverse ageing conditions such as fast-charging protocols²⁹ and emerging chemistries such as solid-state and sodium-ion batteries^{30–32}, predicting calendar life under real-world conditions for more accurate lifetime estimates³³.

While this study demonstrates the data-driven feasibility of inter-cell learning across various conditions, a deeper physical understanding and modelling remain essential³⁴. Future research should prioritize a comprehensive approach to understanding and modelling diverse ageing factors to advance battery research, development and application.

Methods

Building per-cycle feature maps

We use the normalized capacity ($Q \in [0, 1]$) as a unified index to align different electrical signals in different cycles and compute their differentiation within the same stage or between charge and discharge stages. In this way, we do not need to worry about the variable lengths

of different signals across different cycles and charge–discharge stages. Besides, due to the redundant recording of electrical signals, we can obtain these Q -indexed series by performing interpolations on raw time-indexed signals and easily control the feature dimensions by adjusting interpolation granularity. After this step, we can obtain four types of processed series indexed by Q , $V_c(Q)$, $V_d(Q)$, $I_c(Q)$ and $I_d(Q)$, where Q varies from 0 to 1 with a prespecified step, V denotes voltage, I denotes current and d and c are superscripts to denote discharge and charge stages, respectively. We further derive two additional signals to characterize the connections between the charge and discharge stages: (1) $\Delta V(Q) = V_c(Q) - V_d(Q)$, meaning the gap between charge voltages and discharge voltages, and (2) $R(Q) = (V_c(Q) - V_d(Q)) / (I_c(Q) - I_d(Q))$, corresponding to the status of internal resistance as Q varies. These cycle-level features serve as the basic elements for calculating intra- and inter-cell differences.

There are two critical processing steps to ensure the effectiveness of calculating these differences from cycle-level feature maps. This first is to normalize the capacity of all batteries by their respective nominal capacities, ensuring the range of Q falls within $[0, 1]$, to adapt to various battery types with very different nominal capacities. The second is to eliminate the data noise introduced by abrupt changes in current and voltage signals due to the variations in cycling protocols. We use a rolling-median-based filter to alleviate this issue:

$$\mathbf{M} = \text{rolling_median}(\mathbf{r}, w), \quad (1)$$

$$\Delta \mathbf{M} = \text{abs}(\text{rolling_median}(\mathbf{M}, w)), \quad (2)$$

$$r_t^{\text{new}} = \begin{cases} M_t & \text{if } \Delta M_t > 3 \times \text{median}(\Delta \mathbf{M}), \\ r_t & \text{otherwise.} \end{cases} \quad (3)$$

Here $\mathbf{r}$ is the raw signal and w represents the window size of the rolling operations. These multi-faceted feature maps will both intra- and inter-cell learning as separate channels.

Unifying intra- and inter-cell learning

Given the intra-cell feature of a battery cell, denoted $\mathbf{x}$, and its cycle life y , we want to optimize the following objective to obtain a perfect cycle life predictor.

$$\min_{\theta} \mathbb{E}_{(\mathbf{x}, y) \in D} \|f_{\theta}(\mathbf{x}) - y\|_2^2, \quad (4)$$

where D denotes the data distribution, f_{θ} is an encoding function parameterized by θ and $\mathbf{w}_y$ is the weight of the last prediction layer. In practice, we need to perform the empirical risk minimization over limited training instances, that is

$$\min_{\theta} \sum_{i=1}^N \|f_{\theta}(\mathbf{x}_i) - y_i\|_2^2. \quad (5)$$

However, battery lifetime prediction is a special task that suffers from the data-scarce challenge due to the necessity of nonlinear modelling and the huge cost of obtaining data labels. In this scenario, instantiating f_{θ} as a neural network has a high risk of overfitting. To alleviate the data-scarce issue, we consider modelling the differences between two distinct cells:

$$\min_{\theta} \mathbb{E}_{(\Delta \mathbf{x}, \Delta y)} \|g_{\phi}(\Delta \mathbf{x}) - \Delta y\|_2^2, \quad (6)$$

where $\Delta \mathbf{x} = \mathbf{x} - \mathbf{x}'$, $\Delta y = y - y'$, $(\mathbf{x}, y)$ and $(\mathbf{x}', y')$ are independently sampled from D and g_{ϕ} is a function parameterized by ϕ that operates in the space of $\Delta \mathbf{x}$.




Fig. 5 | Analysis of reference cell selection in BatLiNet. a, Best, worst and median prediction errors for all test cells in the MIX-100 dataset, using a single reference cell traversing all training cells. Shaded areas represent the standard deviation of prediction errors from randomly selecting a reference cell. **b**, Effects of increasing reference cell numbers on predictions in the MIX-100 dataset, with eight random selections per case ($n = 8$ trials). Every box plot displays medians as centre lines, the 25th and 75th percentiles as lower and upper quartiles, with whiskers extending to 1.5 times the interquartile range, including outlier points. **c**, Changes in inference throughput as the number of reference cells increases, tested on an NVIDIA GTX 4090 graphical processing unit.

The inter-cell learning formulation (6) assumes that the differences in feature representations for any pair of battery cells hold a unified relationship with the differences in their lifetimes. Specifically, we can establish a clear connection between f_θ and g_ϕ in the linear setting. For example, if y is zero-centred (can be done via preprocessing), and the optimal solution for the original objective (4) is $f_{\theta^*}(x) = \mathbf{w}^* \mathbf{x}$, then it is easy to verify that $g_{\phi^*} = f_{\theta^*}$ is also the optimal solution for the

contrastive objective (6) because $\mathbb{E}_{(\Delta \mathbf{x}, \Delta y)} \|\mathbf{w}^* \Delta \mathbf{x} - \Delta y\|_2^2$ can be decomposed into

$$\mathbb{E}_{(\mathbf{x}, y) \in D} \|\mathbf{w}^* \mathbf{x} - y\|_2^2 + \mathbb{E}_{(\mathbf{x}', y') \in D} \|\mathbf{w}^* \mathbf{x}' - y'\|_2^2. \quad (7)$$

In the nonlinear setting, such as using neural networks as function approximators, it is intractable to establish the exact connection between f_θ and g_ϕ , which leads to two separate optimization procedures. Inspired by the same optimality of the objectives (4) and (6) under the linear setting, we propose to share the last linear layer of f_θ and g_ϕ when using neural networks, that is

$$f_\theta(\mathbf{x}) = \mathbf{w}^T h_\theta(\mathbf{x}), \quad g_\phi(\Delta \mathbf{x}) = \mathbf{w}^T h_\phi(\Delta \mathbf{x}), \quad (8)$$

where $\mathbf{w}$ is the shared parameter, $h_\theta(\cdot)$ and $h_\phi(\cdot)$ are two separate neural networks parameterized by θ and ϕ , respectively. In this way, we can connect the optimization of the objectives (4) and (6) via the shared parameter $\mathbf{w}$ and enjoy the complementarity between intra- and inter-cell modelling.

Moreover, we also need to perform empirical risk minimization to mimic the desired objective (6). Specifically, given that we only have N training instances, we use $N(N-1)$ pairs of instances to substitute the expectation over independently sampled instance pairs. Together with the original regression, we have the following joint empirical risk minimization problem:

$$\min_{\mathbf{w}, \theta, \phi} \sum_{i=1}^N \|\mathbf{w}^T h_\theta(\mathbf{x}_i) - y_i\|_2^2 + \lambda \sum_{i=1}^N \sum_{j=1, j \neq i}^N \|\mathbf{w}^T h_\phi(\Delta \mathbf{x}_{i,j}) - \Delta y_{i,j}\|_2^2, \quad (9)$$

where $\Delta \mathbf{x}_{i,j} = \mathbf{x}_i - \mathbf{x}_j$, $\Delta y_{i,j} = y_i - y_j$ and λ is a hyper-parameter to balance two regression objectives.

After the optimization stage, we can leverage the neural networks for encoding intra- and inter-cell differences to make predictions. Given an unseen instance $\mathbf{x}$, we have

$$\hat{y}^o = \mathbf{w}^T h_\theta(\mathbf{x}), \quad \hat{y}^c = \mathbf{w}^T h_\phi(\mathbf{x} - \mathbf{x}') + y', \quad (10)$$

where $\hat{y}^o$ and $\hat{y}^c$ are the predictions in the original and contrastive space, respectively, and $(\mathbf{x}', y')$ can be sampled from the training instances. Last, we predict the lifetime as $\hat{y} = \alpha \hat{y}^o + (1 - \alpha) \hat{y}^c$, where α is a hyper-parameter to balance the two types of prediction.

Data availability

All datasets used in this study are publicly accessible and include MATR and CLO (<https://data.matr.io/1/>), HUST (<https://data.mendeley.com/datasets/nsc7hnsq4s/2>), CALCE (<https://web.calce.umd.edu/batteries/data>) and RWTH (<https://publications.rwth-aachen.de/record/818642/files>). HNEI, SNL and UL-PUR datasets are hosted by BatteryArchive (https://www.batteryarchive.org/study_summaries.html). Moreover, we leverage an open-source tool, BatteryML³⁵, to preprocess these datasets into a unified format to support experiments reported in this article.

Code availability

Our code is available on CodeOcean³⁶, including instructions to train all models from scratch and to perform direct inference with our provided checkpoints.

References

1. Tarascon, J. M. & Armand, M. Issues and challenges facing rechargeable lithium batteries. *Nature* **414**, 359–367 (2001).
2. Armand, M. & Tarascon, J. M. Building better batteries. *Nature* **451**, 652–657 (2008).

3. Dunn, B. S., Kamath, H. & Tarascon, J. M. Electrical energy storage for the grid: a battery of choices. *Science* **334**, 928–935 (2011).
4. Arora, P., White, R. E. & Doyle, M. Capacity fade mechanisms and side reactions in lithium-ion batteries. *J. Electrochem. Soc.* **145**, 3647 (1998).
5. Ng, M.-F., Zhao, J., Yan, Q., Conduit, G. J. & Seh, Z.W. Predicting the state of charge and health of batteries using data-driven machine learning. *Nat. Mach. Intell.* **2**, 161–170 (2020).
6. Yao, Z. et al. Machine learning for a sustainable energy future. *Nat. Rev. Mater.* **8**, 202–215 (2023).
7. Biggio, L., Bendinelli, T., Kulkarni, C. & Fink, O. Ageing-aware battery discharge prediction with deep learning. *Appl. Energy* **346**, 121229 (2023).
8. Li, T., Zhou, Z., Thelen, A., Howey, D. A. & Hu, C. Predicting battery lifetime under varying usage conditions from early aging data. *Cell Rep. Phys. Sci.* **5**, 101891 (2024).
9. Severson, K. A. et al. Data-driven prediction of battery cycle life before capacity degradation. *Nat. Energy* **4**, 383–391 (2019).
10. Attia, P. M., Severson, K. A. & Witmer, J. D. Statistical learning for accurate and interpretable battery lifetime prediction. *J. Electrochem. Soc.* **168**, 090547 (2021).
11. Attia, P. M. et al. Closed-loop optimization of fast-charging protocols for batteries with machine learning. *Nature* **578**, 397–402 (2020).
12. Ma, G. et al. Real-time personalized health status prediction of lithium-ion batteries using deep transfer learning. *Ener. Environ. Sci.* **15**, 4083–4094 (2022).
13. He, W., Williard, N., Osterman, M. & Pecht, M. Prognostics of lithium-ion batteries based on Dempster-Shafer theory and the Bayesian Monte Carlo method. *J. Power Sources* **196**, 10314–10321 (2011).
14. Xing, Y., Ma, E., Tsui, K.L. & Pecht, M. An ensemble model for predicting the remaining useful performance of lithium-ion batteries. *Microelectron. Reliab.* **53**, 811–820 (2013).
15. Williard, N., He, W., Osterman, M. & Pecht, M. Comparative analysis of features for determining state of health in lithium-ion batteries. *Int. J. Progn. Health Manag.* **4**, 1–7 (2013).
16. Devie, A., Baure, G. & Dubarry, M. Intrinsic variability in the degradation of a batch of commercial 18650 lithium-ion cells. *Energies* **11**, 1031 (2018).
17. Juarez-Robles, D., Jeevarajan, J. A. & Mukherjee, P. P. Degradation-safety analytics in lithium-ion cells: part I. Aging under charge/discharge cycling. *J. Electrochem. Soc.* **167**, 160510 (2020).
18. Li, W. et al. One-shot battery degradation trajectory prediction with deep learning. *J. Power Sources* **506**, 230024 (2021).
19. Preger, Y. et al. Degradation of commercial lithium-ion cells as a function of chemistry and cycling conditions. *J. Electrochem. Soc.* **167**, 120532 (2020).
20. Hoerl, A. E. & Kennard, R. W. Ridge regression: biased estimation for nonorthogonal problems. *Technometrics* **12**, 55–67 (1970).
21. Geladi, P. & Kowalski, B. R. Partial least-squares regression: a tutorial. *Anal. Chim. Acta* **185**, 1–17 (1986).
22. Jolliffe, I. T. Principal components in regression analysis. In *Principal Component Analysis*. Springer Series in Statistics 167–198 (Springer, 2002).
23. Drucker, H. et al. Support vector regression machines. *Adv. Neural Inform. Proc. Syst.* **9**, 155–161 (1996).
24. Breiman, L. Random forests. *Mach. Learn.* **45**, 5–32 (2001).
25. Cybenko, G. Approximation by superpositions of a sigmoidal function. *Math. Control Signal. Syst.* **2**, 303–314 (1989).
26. Hochreiter, S. & Schmidhuber, J. Long short-term memory. *Neural Comput.* **9**, 1735–1780 (1997).
27. Krizhevsky, A., Sutskever, I. & Hinton, G. E. ImageNet classification with deep convolutional neural networks. *Commun. ACM* **60**, 84–90 (2012).
28. Roman, D., Saxena, S., Robu, V., Pecht, M. & Flynn, D. Machine learning pipeline for battery state-of-health estimation. *Nat. Mach. Intell.* **3**, 447–456 (2021).
29. Zeng, Y. et al. Extreme fast charging of commercial Li-ion batteries via combined thermal switching and self-heating approaches. *Nat. Commun.* **14**, 3229 (2023).
30. Randau, S. et al. Benchmarking the performance of all-solid-state lithium batteries. *Nat. Energy* **5**, 259–270 (2020).
31. Shraer, S. D. et al. Development of vanadium-based polyanion positive electrode active materials for high-voltage sodium-based batteries. *Nat. Commun.* **13**, 4097 (2022).
32. Li, Y. et al. Origin of fast charging in hard carbon anodes. *Nat. Energy* **9**, 134–142 (2024).
33. McBrayer, J. D. et al. Calendar aging of silicon-containing batteries. *Nat. Energy* **6**, 866–872 (2021).
34. Brosa Planella, F. et al. A continuum of physics-based lithium-ion battery models reviewed. *Prog. Energy* **4**, 042003 (2022).
35. Zhang, H. et al. BatteryML: an open-source platform for machine learning on battery degradation. *Zenodo* 10.5281/zenodo.14040125 (2024).
36. Zhang, H. et al. Code for ‘battery lifetime prediction across diverse aging conditions with inter-cell deep learning’. *Code Ocean* <https://doi.org/10.24433/CO.8904065.v2> (2024).

Acknowledgements

The authors received no specific funding for this work.

Author contributions

H.Z. led the development of BatLiNet and conducted most experiments. Y.L. provided valuable insights into cross-chemistry and facilitated the experiment analysis. S.Z. led the overall research, built the initial code base and introduced the concept of inter-cell learning. S.Z., H.Z. and Y.L. wrote the paper. Z.L. contributed to refining the manuscript and inspiring cross-chemistry explorations. X.G. ensured the reproduction of all experiments and refined paper presentation. Both W.X. and J.B. participated in the discussion of ideas and analysis of results. J.B. provided overarching supervision for the research and strong support for the entire team.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s42256-024-00972-x>.

Correspondence and requests for materials should be addressed to Shun Zheng or Jiang Bian.

Peer review information *Nature Machine Intelligence* thanks Kandler Smith, Qingyu Yan and the other, anonymous, reviewer(s) for their contribution to the peer review of this work.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line

to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2025