

Closed-loop optimization of fast-charging protocols for batteries with machine learning

<https://doi.org/10.1038/s41586-020-1994-5>

Received: 6 August 2019

Accepted: 19 December 2019

Published online: 19 February 2020

 Check for updates

Peter M. Attia^{1,7}, Aditya Grover^{2,7}, Norman Jin¹, Kristen A. Severson³, Todor M. Markov², Yang-Hung Liao¹, Michael H. Chen¹, Bryan Cheong^{1,2}, Nicholas Perkins¹, Zi Yang¹, Patrick K. Herring⁴, Muratahan Aykol⁴, Stephen J. Harris^{1,5}, Richard D. Braatz^{3,✉}, Stefano Ermon^{2,✉} & William C. Chueh^{1,6,✉}

Simultaneously optimizing many design parameters in time-consuming experiments causes bottlenecks in a broad range of scientific and engineering disciplines^{1,2}. One such example is process and control optimization for lithium-ion batteries during materials selection, cell manufacturing and operation. A typical objective is to maximize battery lifetime; however, conducting even a single experiment to evaluate lifetime can take months to years^{3–5}. Furthermore, both large parameter spaces and high sampling variability^{3,6,7} necessitate a large number of experiments. Hence, the key challenge is to reduce both the number and the duration of the experiments required. Here we develop and demonstrate a machine learning methodology to efficiently optimize a parameter space specifying the current and voltage profiles of six-step, ten-minute fast-charging protocols for maximizing battery cycle life, which can alleviate range anxiety for electric-vehicle users^{8,9}. We combine two key elements to reduce the optimization cost: an early-prediction model⁵, which reduces the time per experiment by predicting the final cycle life using data from the first few cycles, and a Bayesian optimization algorithm^{10,11}, which reduces the number of experiments by balancing exploration and exploitation to efficiently probe the parameter space of charging protocols. Using this methodology, we rapidly identify high-cycle-life charging protocols among 224 candidates in 16 days (compared with over 500 days using exhaustive search without early prediction), and subsequently validate the accuracy and efficiency of our optimization approach. Our closed-loop methodology automatically incorporates feedback from past experiments to inform future decisions and can be generalized to other applications in battery design and, more broadly, other scientific domains that involve time-intensive experiments and multi-dimensional design spaces.

Optimal experimental design (OED) approaches are widely used to reduce the cost of experimental optimization. These approaches often involve a closed-loop pipeline where feedback from completed experiments informs subsequent experimental decisions, balancing the competing demands of exploration—that is, testing regions of the experimental parameter space with high uncertainty—and exploitation—that is, testing promising regions based on the results of the completed experiments. Adaptive OED algorithms have been successfully applied to physical science domains, such as materials science^{1,2,12–14}, chemistry^{15,16}, biology¹⁷ and drug discovery¹⁸, as well as to computer science domains, such as hyperparameter optimization for machine learning^{19,20}. However, while a closed-loop approach is designed to

minimize the number of experiments required for optimizing across a multi-dimensional parameter space, the time (and cost) per experiment may remain high, as is the case for lithium-ion batteries. Therefore, an OED approach should account for both the number of experiments and the cost per experiment. Multi-fidelity optimization approaches have been developed to learn from both inexpensive, noisy signals and expensive, accurate signals. For example, in hyperparameter optimization for machine learning algorithms, several low-fidelity signals for predicting the final performance of an algorithmic configuration (for example, extrapolated learning curves^{19,20}, rapid testing on a subset of the full training dataset²¹) are used in tandem with more complete configuration evaluations^{22,23}. For lithium-ion batteries, classical

¹Department of Materials Science and Engineering, Stanford University, Stanford, CA, USA. ²Department of Computer Science, Stanford University, Stanford, CA, USA. ³Department of Chemical Engineering, Massachusetts Institute of Technology, Cambridge, MA, USA. ⁴Toyota Research Institute, Los Altos, CA, USA. ⁵Materials Science Division, Lawrence Berkeley National Laboratory, Berkeley, CA, USA. ⁶Applied Energy Division, SLAC National Accelerator Laboratory, Menlo Park, CA, USA. ⁷These authors contributed equally: Peter M. Attia, Aditya Grover. ✉e-mail: braatz@mit.edu; ermon@cs.stanford.edu; wchueh@stanford.edu




Fig. 1 | Schematic of our CLO system. First, batteries are tested. The cycling data from the first 100 cycles (specifically, electrochemical measurements such as voltage and capacity) are used as input for an early outcome prediction of cycle life. These cycle life predictions from a machine learning (ML) model are subsequently sent to a BO algorithm, which recommends the next protocols to test by balancing the competing demands of exploration (testing protocols with high uncertainty in estimated cycle life) and exploitation

(testing protocols with high estimated cycle life). This process iterates until the testing budget is exhausted. In this approach, early prediction reduces the number of cycles required per tested battery, while optimal experimental design reduces the number of experiments required. A small training dataset of batteries cycled to failure is used both to train the early outcome predictor and to set BO hyperparameters. In future work, design of battery materials and processes could also be integrated into this closed-loop system.

methods such as factorial design that use predetermined heuristics to select experiments have been applied^{24–26}, but the design and use of low-fidelity signals is challenging and unexplored. These previously considered approaches do not discover and exploit the patterns present in the parameter space for efficient optimization, nor do they address the issue of time per experiment.

In this work, we develop a closed-loop optimization (CLO) system with early outcome prediction for efficient optimization over large parameter spaces with expensive experiments and high sampling variability. We employ this system to experimentally optimize fast-charging protocols for lithium-ion batteries; reducing charging times to approach gasoline refuelling time is critical to reduce range anxiety for electric vehicles^{8,9} but often comes at the expense of battery lifetime. Specifically, we optimize over a parameter space consisting of 224 unique six-step, ten-minute fast-charging protocols (that is, how current and voltage are controlled during charging) to find charging protocols with high cycle life (defined as the battery capacity falling to 80% of its nominal value). Our system uses two key elements to reduce the optimization cost (Extended Data Fig. 1). First, we reduce the time per experiment by using machine learning to predict the outcome of the experiment based on data from early cycles, well before the batteries reach the end of life⁵. Second, we reduce the number of experiments by using a Bayesian optimization (BO) algorithm to balance the exploration–exploitation tradeoff in choosing the next round of experiments^{10,11}. Testing a single battery to failure under our fast-charging conditions requires approximately 40 days, meaning that when 48 experiments are performed in parallel, assessing all 224 charging protocols with triplicate measurements takes approximately 560 days. Here, using CLO with early outcome prediction, only 16 days were required to confidently identify protocols with high cycle lives (48 parallel experiments). In a subsequent validation study, we find that CLO ranks these protocols by lifetime accurately (Kendall rank correlation coefficient, 0.83) and efficiently (15 times less time than a baseline ‘brute-force’ approach that uses random search without early prediction). Furthermore, we find that the charging protocols identified as optimal by CLO with early prediction outperform existing fast-charging protocols designed to avoid lithium plating (a common fast-charging degradation mode), the approach suggested by conventional battery wisdom^{4,8,9,26}. This work highlights the utility of combining CLO with inexpensive early outcome predictors to accelerate scientific discovery.

CLO with early outcome prediction is depicted schematically in Fig. 1. The system consists of three components: parallel battery cycling, an early predictor for cycle life and a BO algorithm. At each sequential round, we iterate over these three components. The first component is a multi-channel battery cyler; the cyler used in this work tests 48 batteries simultaneously. Before starting CLO, the charging protocols for the first round of 48 batteries are chosen at random (without replacement) from the complete set of 224 unique multi-step protocols (Methods). Each battery undergoes repeated charging and discharging for 100 cycles (about 4 days; average predicted cycle life 905 cycles), beyond which the experiments are terminated.

These cycling data are then fed as input to the early outcome predictor, which estimates the final cycle lives of the batteries given data from the first 100 cycles. The early predictor is a linear model trained via elastic net regression²⁷ on features extracted from the charging data of the first 100 cycles (Supplementary Table 1), similar to that presented in Severson et al.⁵. Predictive features include transformations of both differences between voltage curves and discharge capacity fade trends. To train the early predictor, we require a training dataset of batteries cycled to failure. Here, we used a pre-existing dataset of 41 batteries cycled to failure (cross-validation root-mean-square error, 80.4 cycles; see Methods and Supplementary Discussion 1). Whereas obtaining this dataset itself requires running full cycling experiments for a small training set of batteries (the cost we are trying to offset), this one-time cost could be avoided if pretrained predictors or previously collected datasets are available. If unavailable, we pay an upfront cost in collecting this dataset; this dataset could also be used for warm-starting the BO algorithm. The size of the dataset collected should best tradeoff the upfront cost in acquiring the dataset to train an accurate model with the anticipated reduction in experimentation requirements for CLO.

Finally, these predicted cycle lives from early-cycle data are fed into the BO algorithm (Methods and Supplementary Discussion 2), which recommends the next round of 48 charging protocols that best balance the exploration–exploitation tradeoff. This algorithm (Methods and Supplementary Discussion 2) builds on the prior work of Hoffman et al.¹⁰ and Grover et al.¹¹. The algorithm maintains an estimate of both the average cycle life and the uncertainty bounds for each protocol; these estimates are initially equal for all protocols and are refined as additional data are collected. Crucially, to reduce the total optimization cost, our algorithm performs these updates using estimates from





Fig. 2 | Structure of our six-step, ten-minute fast-charging protocols.

Currents are defined as dimensionless C rates; here, 1C is 1.1 A, or the current required to fully (dis)charge the nominal capacity (1.1 A h) in 1 h. **a**, Current versus SOC for an example charging protocol, 7.0C–4.8C–5.2C–3.45C (bold lines). Each charging protocol is defined by five constant current (CC) steps followed by one constant voltage (CV) step. The last two steps (CC5 and CV1) are identical for all charging protocols. We optimize over the first four constant-current steps, denoted CC1, CC2, CC3 and CC4. Each of these steps comprises a 20% SOC window, such that CC1 ranges from 0% to 20% SOC, CC2

ranges from 20% to 40% SOC, and so on. CC4 is constrained by specifying that all protocols charge in the same total time (10 min) from 0% to 80% SOC. Thus, our parameter space consists of unique combinations of the three free parameters CC1, CC2 and CC3. For each step, we specify a range of acceptable values; the upper limit is monotonically decreasing with increasing SOC to avoid the upper cutoff potential (3.6 V for all steps). **b**, CC4 (colour scale) as a function of CC1, CC2 and CC3 (on the x, y and z axes, respectively). Each point represents a unique charging protocol.

the early outcome predictor instead of using the actual cycle lives. The mean and uncertainty estimates for the cycle lives are obtained via a Gaussian process (Methods), which has a smoothing effect and allows for updating the cycle life estimates of untested protocols with the predictions from related protocols. The closed-loop process repeats until the optimization budget, in our case 192 batteries tested (100 cycles each), is exhausted.

Our objective is to find the charging protocol which maximizes the expected battery cycle life for a fixed charging time (ten minutes) and state-of-charge (SOC) range (0 to 80%). The design space of our 224 six-step extreme fast-charging protocols is presented in Fig. 2a. Multi-step charging protocols, in which a series of different constant-current steps are applied within a single charge, are considered advantageous over single-step charging for maximizing cycle life during fast charging^{4,8}, though the optimal combination remains unclear. As shown in Fig. 2b, each protocol is specified by three independent parameters (CC1, CC2 and CC3); each parameter is a current applied over a fixed SOC range (0–20%, 20–40% and 40–60%, respectively). A fourth parameter, CC4, is dependent on CC1, CC2, CC3 and the charging time. Given constraints on the current values (Methods), a total of 224 charging protocols are permitted. We test commercial lithium iron phosphate (LFP)/graphite cylindrical batteries (A123 Systems) in a convective environmental chamber (30 °C ambient temperature). A maximum voltage of 3.6 V is imposed. These batteries are designed to fast-charge in 17 min (rate testing data are presented in Extended Data Fig. 2). The cycle life decreases dramatically with faster charging time^{4,5}, motivating this optimization. Since the LFP positive electrode is generally considered to be stable^{4,5}, we select this battery chemistry to isolate the effects of extreme fast charging on graphite, which is universally employed in lithium-ion batteries.

In all, we ran four CLO rounds sequentially, consisting of 185 batteries in total (excluding seven batteries; see Methods). Using early prediction, each CLO round requires four days to complete 100 cycles, resulting in a total testing time of sixteen days—a major reduction from the 560 days required to test each charging protocol to failure three times. Figure 3 presents the predictions and selected protocols (Fig. 3a), as well as the evolution of cycle life estimates over the parameter space

as the optimization progresses (Fig. 3a). Initially, the estimated cycle lives for all protocols are equal. After two rounds, the overall structure of the parameter space (that is, the dependence of cycle life on charging protocol parameters CC1, CC2 and CC3) emerges, and a prominent region with high cycle life protocols has been identified. The confidence of CLO in this high-performing region is further improved from round 2 to round 4, but overall the cycle life estimates do not change substantially (Extended Data Fig. 3). By learning and exploiting the structure of the parameter space, we avoid evaluating charging protocols with low estimated cycle life and concentrate more resources on the high-performing region (Extended Data Figs. 3–5). Specifically, 117 of 224 protocols are never tested (Fig. 3c); we spend 67% of the batteries testing 21% of the protocols (0.83 batteries per protocol on average). CLO repeatedly tests several protocols with high estimated cycle life to decrease uncertainties due to manufacturing variability and the error introduced by early outcome prediction. The uncertainty is expressed as the prediction intervals of the posterior predictive distribution over cycle life (Extended Data Figs. 3g, 4, 5).

To the best of our knowledge, this work presents the largest known map of cycle life as a function of charging conditions (Extended Data Fig. 5). This dataset can be used to validate physics-based models of battery degradation. Most fast-charging protocols proposed in the battery literature suggest that current steps decreasing monotonically as a function of SOC are optimal to avoid lithium plating on graphite, a well-accepted degradation mode during fast charging^{4,8,9,26}. In contrast, the protocols identified as optimal by CLO (for example, Fig. 3d) are generally similar to single-step constant-current charging (that is, $CC_1 \approx CC_2 \approx CC_3 \approx CC_4$). Specifically, of the 75 protocols with the highest estimated cycle lives, only ten are monotonically decreasing (that is, $CC_i \geq CC_{i+1}$ for all i) and two are strictly decreasing (that is, $CC_i > CC_{i+1}$). We speculate that minimizing parasitic reactions caused by heat generation may be the operative optimization strategy for these cells, as opposed to minimizing the propensity for lithium plating (Supplementary Discussion 3). While the optimal protocol for a new scenario would depend on the selected charge time, SOC window, temperature control conditions and battery chemistry, this unexpected result highlights the need for data-driven approaches for optimizing fast charging.




Fig. 3 | Results of closed-loop experiments. **a**, Early cycle life predictions per round. The tested charging protocols and the resulting predictions are plotted for rounds 1–4. Each point represents a charging protocol, defined by CC1, CC2 and CC3 (the *x*, *y* and *z* axes, respectively). The colour scale represents cycle life predictions from the early outcome prediction model. The charging protocols in the first round of testing are randomly selected. As the BO algorithm shifts from exploration to exploitation, the charging protocols selected for testing by the closed loop in subsequent rounds fall primarily into the high-performing region. **b**, Evolution of the parameter space per round. The colour scale represents cycle life, as estimated by the BO algorithm. The initial cycle life

estimates are equivalent for all protocols; as more predictions are generated, the BO algorithm updates its cycle life estimates. The CLO-estimated mean cycle lives after four rounds for all fast-charging protocols in the parameter space are also presented in Extended Data Fig. 5 and Supplementary Table 3. **c**, Distribution of the number of repetitions for each charging protocol (excluding failed batteries). Only 46 of 224 protocols (21%) are tested multiple times. **d**, Current versus SOC for the top three fast-charging protocols, as estimated by CLO. CC1–CC4 are displayed in the legend. All three protocols have relatively uniform charging (that is, CC1 ≈ CC2 ≈ CC3 ≈ CC4).

We validate the performance of CLO with early prediction on a subset of nine extreme fast-charging protocols. For each of these protocols, we cycle five batteries each to failure and use the sample average of the final cycle lives as an estimate of the true lifetimes. We use this validation study to (1) confirm that CLO is able to correctly rank protocols based on cycle life, (2) compare the cycle lives of protocols recommended by CLO to protocols inspired by the battery literature and (3) compare the performance of CLO to baseline ablation approaches for experimental design. The charging protocols used in validation,

some of which are inspired by existing battery fast-charging literature (see Methods), span the range of estimated cycle lives (Extended Data Fig. 6 and Extended Data Table 1). We adjust the voltage limits and charging times of these literature protocols to match our protocols, while maintaining similar current ratios as a function of SOC. Whereas the literature protocols used in these validation experiments are generally designed for batteries with high-voltage positive electrode chemistries, fast-charging optimization strategies generally focus on the graphitic negative electrode^{4,8}. For these nine protocols, we validate

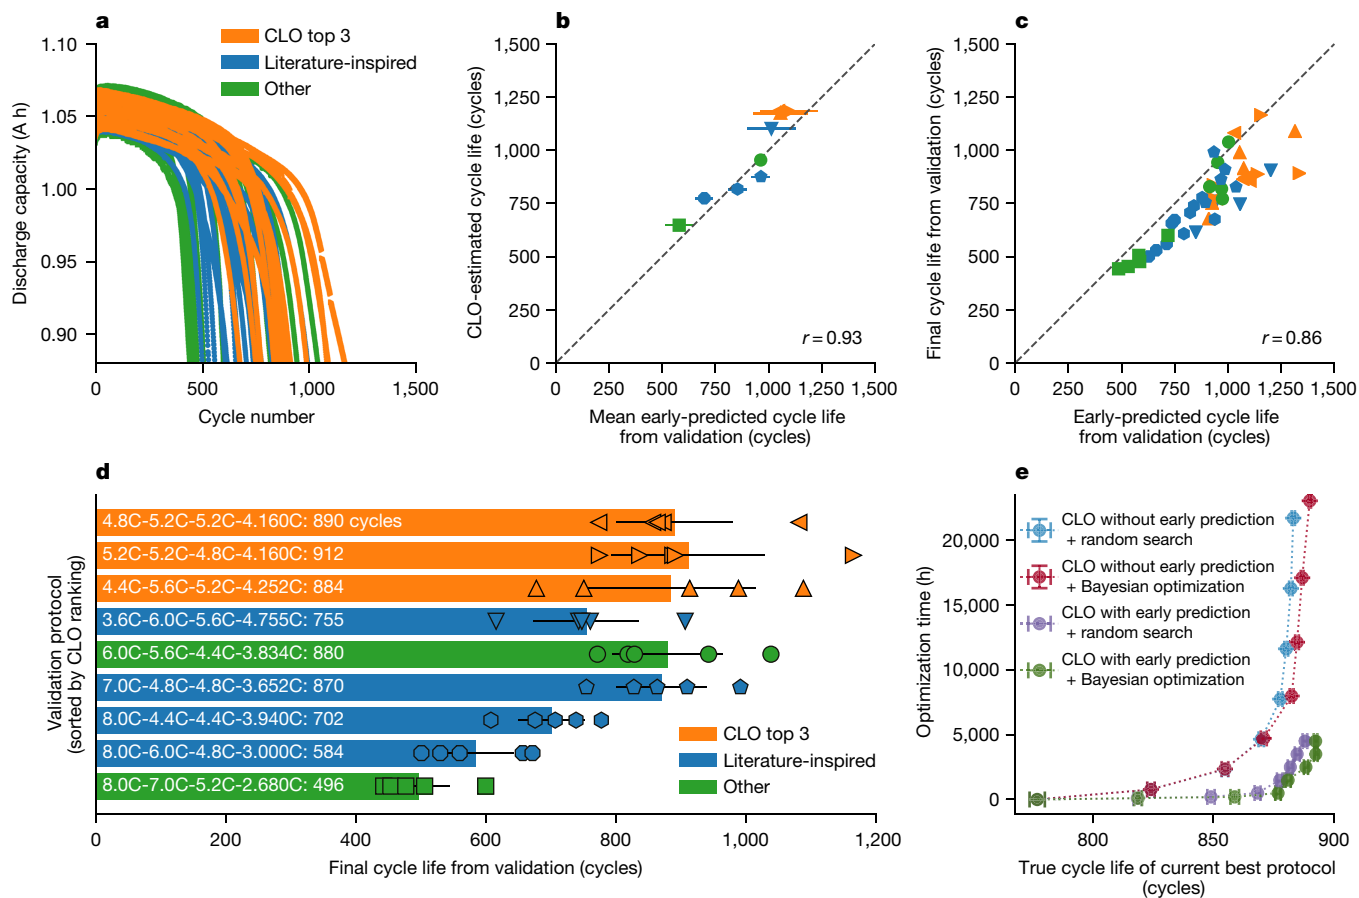


Fig. 4 | Results of validation experiment. **a**, Discharge capacity versus cycle number for all batteries in the validation experiment. The nine validation protocols include the top three protocols as estimated by CLO ('CLO top 3'), four protocols inspired by the battery literature^{39–44} ('Literature-inspired') and two protocols selected to obtain a representative sampling from the distribution of CLO-estimated cycle lives among the validation protocols ('Other'). **b**, Comparison of early-predicted cycle lives from validation to closed-loop estimates, averaged on a protocol basis. Each ten-minute charging protocol is tested with five batteries. Error bars represent the 95% confidence intervals. **c**, Observed versus early-predicted cycle life for the validation experiment. Although our early predictor tends to overestimate cycle life,

the 'CLO-estimated' cycle lives against the sample average of the five final cycle lives.

The validation results are presented in Fig. 4. The discharge capacity fade curves (Fig. 4a) exhibit the nonlinear decay typical of fast charging^{5,7}. If we apply our early-prediction model to the batteries in the validation experiment, these early predictions (averaged over each protocol) match the CLO-estimated mean cycle lives well (Pearson correlation coefficient $r = 0.93$; Fig. 4b). This result validates the performance of the BO component of CLO in particular, since the CLO-estimated cycle lives were inferred from early predictions. However, our early-prediction model exhibits some bias (Fig. 4c), probably owing to calendar ageing effects from different battery storage times²⁸ (Supplementary Table 2 and Supplementary Discussion 4). Despite this bias in our predictive model, we generally capture the ranking well (Kendall rank correlation coefficient, 0.83; Fig. 4d and Extended Data Fig. 7). At the same time, we note that the final cycle lives for the top-ranked protocols are similar. Furthermore, the optimal protocols identified by CLO outperform protocols inspired by previously published fast-charging protocols (895 versus 728 cycles on average; Extended Data Fig. 6 and Extended Data Table 1). This result suggests that the efficiency of our approach does not come at the expense of accuracy.

probably owing to calendar ageing effects (Supplementary Discussion 4), the trend is correctly captured (Pearson correlation coefficient $r = 0.86$). **d**, Final cycle lives from validation, sorted by CLO ranking. The length of each bar and the annotations represents the mean final cycle life from validation per protocol. Error bars represent the 95% confidence intervals. **e**, Ablation study of various optimization approaches using the protocols and data in the validation set (Methods). Error bars represent the 95% confidence intervals ($n = 2,000$). With contributions from both early prediction and Bayesian optimization, CLO can rapidly identify high-performing charging protocols. The gains from Bayesian optimization are larger when resources are constrained (Extended Data Fig. 8).

Our method greatly reduces the optimization time required compared to baseline optimization approaches (Fig. 4e). For instance, a procedure that does not use early outcome prediction and simply selects protocols randomly to test begins to saturate at a competitive performance level after about 7,700 battery-hours of testing. To achieve a similar level of performance, CLO with both early outcome prediction and the BO algorithm requires only 500 battery-hours of testing. For this small-scale validation experiment, we observe that the early-prediction component of CLO greatly reduces the time per experiment. Here, random selection is equivalent to a pure exploration strategy and can achieve a performance similar to the BO-based approaches for smaller experimental budgets. In later stages, random selection is eventually outperformed by BO-based approaches, which exploit the structure across the protocols and focus on reducing the uncertainty in the promising regions of the parameter space. Although these results are specific to this validation study, we observe similar or larger gains in simulations when fewer batteries or fewer parallel experiments (relative to the size of the parameter space) are available (Extended Data Fig. 8). The relative gains from BO over random selection are largest with minimal resources.

Finally, we compare our early predictor with other low-fidelity predictors proposed in state-of-the-art multi-fidelity optimization algorithms

in the literature^{19,20}, and find that our approach outperforms these algorithms (Supplementary Discussion 2 and Supplementary Table 4). The generic early-prediction models in these previous works fit composites of parametric functions to the capacity fade curves, while our model uses additional features recorded at every cycle (for example, voltage). This result highlights the value of designing predictive models for the target application in multi-fidelity optimization.

In summary, we have successfully accelerated the optimization of extreme fast charging for lithium-ion batteries using CLO with early outcome prediction. This method could extend to other fast-charging design spaces, such as pulsed^{26,28} and constant-power³ charging, as well as to other objectives, such as slower charging and calendar ageing. Additionally, this work opens up new applications for battery optimization, such as formation²⁹, adaptive cycling³⁰ and parameter estimation for battery management system models³¹. Furthermore, provided that a suitable early outcome predictor exists, this method could also be applied to optimize other aspects of battery development, such as electrode materials and electrolyte chemistries. Beyond batteries, our CLO approach combining black-box optimization with early outcome prediction can be extended to efficiently optimize other physical^{1,2,18} and computational^{22,32} multi-dimensional parameter spaces that involve time-intensive experimentation, illustrating the power of data-driven methods to accelerate the pace of scientific discovery.

Online content

Any methods, additional references, Nature Research reporting summaries, source data, extended data, supplementary information, acknowledgements, peer review information; details of author contributions and competing interests; and statements of data and code availability are available at <https://doi.org/10.1038/s41586-020-1994-5>.

- Tabor, D. P. et al. Accelerating the discovery of materials for clean energy in the era of smart automation. *Nat. Rev. Mater.* **3**, 5–20 (2018).
- Butler, K. T., Davies, D. W., Cartwright, H., Isayev, O. & Walsh, A. Machine learning for molecular and materials science. *Nature* **559**, 547–555 (2018).
- Baumhöfer, T., Brühl, M., Rothgang, S. & Sauer, D. U. Production caused variation in capacity aging trend and correlation to initial cell performance. *J. Power Sources* **247**, 332–338 (2014).
- Keil, P. & Jossen, A. Charging protocols for lithium-ion batteries and their impact on cycle life—an experimental study with different 18650 high-power cells. *J. Energy Storage* **6**, 125–141 (2016).
- Severson, K. A. et al. Data-driven prediction of battery cycle life before capacity degradation. *Nat. Energy* **4**, 383–391 (2019).
- Schuster, S. F., Brand, M. J., Berg, P., Gleissenberger, M. & Jossen, A. Lithium-ion cell-to-cell variation during battery electric vehicle operation. *J. Power Sources* **297**, 242–251 (2015).
- Harris, S. J., Harris, D. J. & Li, C. Failure statistics for commercial lithium ion batteries: a study of 24 pouch cells. *J. Power Sources* **342**, 589–597 (2017).
- Ahmed, S. et al. Enabling fast charging—a battery technology gap assessment. *J. Power Sources* **367**, 250–262 (2017).
- Liu, Y., Zhu, Y. & Cui, Y. Challenges and opportunities towards fast-charging battery materials. *Nat. Energy* **4**, 540–550 (2019).
- Hoffman, M. W., Shahriari, B. & de Freitas, N. On correlation and budget constraints in model-based bandit optimization with application to automatic machine learning. In

- Proc. 17th Int. Conf. on Artificial Intelligence and Statistics (AISTATS)* Vol. 33, 365–374 (Proceedings of Machine Learning Research, 2014); <http://proceedings.mlr.press/v33/hoffman14.html>.
- Grover, A. et al. Best arm identification in multi-armed bandits with delayed feedback. In *Proc. 21st Int. Conf. on Artificial Intelligence and Statistics (AISTATS)* Vol. 84, 833–842 (Proceedings of Machine Learning Research, 2018); <http://proceedings.mlr.press/v84/grover18b.html>.
 - Nikolaev, P. et al. Autonomy in materials research: a case study in carbon nanotube growth. *npj Comput. Mater.* **2**, 16031 (2016).
 - Ling, J., Hutchinson, M., Antono, E., Paradiso, S. & Meredig, B. High-dimensional materials and process optimization using data-driven experimental design with well-calibrated uncertainty estimates. *Integr. Mater. Manuf. Innov.* **6**, 207–217 (2017).
 - Balachandran, P. V., Kowalski, B., Sehrioglu, A. & Lookman, T. Experimental search for high-temperature ferroelectric perovskites guided by two-step machine learning. *Nat. Commun.* **9**, 1668 (2018).
 - Bédard, A.-C. et al. Reconfigurable system for automated optimization of diverse chemical reactions. *Science* **361**, 1220–1225 (2018).
 - Granda, J. M., Donina, L., Dragone, V., Long, D.-L. & Cronin, L. Controlling an organic synthesis robot with machine learning to search for new reactivity. *Nature* **559**, 377–381 (2018).
 - King, R. D. et al. The automation of science. *Science* **324**, 85–89 (2009).
 - Schneider, G. Automating drug discovery. *Nat. Rev. Drug Discov.* **17**, 97–113 (2018).
 - Domhan, T., Springenberg, J. T. & Hutter, F. Speeding up automatic hyperparameter optimization of deep neural networks by extrapolation of learning curves. In *Proc. 24th Int. Conf. on Artificial Intelligence* 3460–3468 (AAAI Press, 2015).
 - Klein, A., Falkner, S., Springenberg, J. T. & Hutter, F. Learning curve prediction with Bayesian neural networks. In *Proc. 2017 Int. Conf. on Learning Representations* 1–16 (2017); <https://openreview.net/forum?id=S11KBVcX>.
 - Petrak, J. *Fast Subsampling Performance Estimates for Classification Algorithm Selection*. Technical Report TR-2000-07, 3–14 (Austrian Research Institute for Artificial Intelligence, 2000); <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.28.3305&rep=rep1&type=pdf>.
 - Li, L., Jamieson, K., DeSalvo, G., Rostamizadeh, A. & Talwalkar, A. Hyperband: a novel bandit-based approach to hyperparameter optimization. *J. Mach. Learn. Res.* **18**, 1–52 (2018).
 - Hutter, F., Hoos, H. H. & Leyton-Brown, K. Sequential model-based optimization for general algorithm configuration. In *Proc. 5th Int. Conf. on Learning and Intelligent Optimization* 507–523 (Springer, 2011).
 - Luo, Y., Liu, Y. & Wang, S. Search for an optimal multistage charging pattern for lithium-ion batteries using the Taguchi approach. In *Region 10 Conf. (TENCON 2009)* 1–5, <https://doi.org/10.1109/TENCON.2009.5395823> (IEEE, 2009).
 - Liu, Y., Hsieh, C. & Luo, Y. Search for an optimal five-step charging pattern for Li-ion batteries using consecutive orthogonal arrays. *IEEE Trans. Energy Convers.* **26**, 654–661 (2011).
 - Schindler, S., Bauer, M., Cheetamun, H. & Danzer, M. A. Fast charging of lithium-ion cells: identification of aging-minimal current profiles using a design of experiment approach and a mechanistic degradation analysis. *J. Energy Storage* **19**, 364–378 (2018).
 - Zou, H. & Hastie, T. Regularization and variable selection via the elastic net. *J. R. Stat. Soc. Ser. B* **67**, 301–320 (2005).
 - Keil, P. et al. Calendar aging of lithium-ion batteries. I. Impact of the graphite anode on capacity fade. *J. Electrochem. Soc.* **163**, A1872–A1880 (2016).
 - Wood, D. L., Li, J. & Daniel, C. Prospects for reducing the processing cost of lithium ion batteries. *J. Power Sources* **275**, 234–242 (2015).
 - Zimmerman, A. H., Quinzio, M. V. & Monica, S. Adaptive charging method for lithium-ion battery cells. US Patent US6204634B1 (2001).
 - Park, S., Kato, D., Gima, Z., Klein, R. & Moura, S. Optimal experimental design for parameterization of an electrochemical lithium-ion battery model. *J. Electrochem. Soc.* **165**, A1309–A1323 (2018).
 - Smith, J. S., Nebgen, B., Lubbers, N., Isayev, O. & Roitberg, A. E. Less is more: sampling chemical space with active learning. *J. Chem. Phys.* **148**, 241733 (2018).

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

© The Author(s), under exclusive licence to Springer Nature Limited 2020

Methods

Experimental

Commercial high-power lithium iron phosphate (LFP)/graphite A123 APR18650M1A cylindrical cells were used in this work (packing date 2015-09-26, lot number EL1508007-R). These cells have a nominal capacity of 1.1 A h and a nominal voltage of 3.3 V. All currents are defined in units of C rate; here, 1C is 1.1 A, or the current required to fully (dis)charge the nominal capacity (1.1 A h) in 1 h. The manufacturer's recommended fast-charging protocol is 3.6C (3.96 A) CC-CV. The rate capability of these cells is shown in Extended Data Fig. 2. The graphite and LFP electrodes are 40 μm thick and 80 μm thick, respectively, as quantified via X-ray tomography (Zeiss Xradia 520 Versa).

The cells were cycled with various charging protocols but identically discharged. Cells were charged with one of 224 candidate six-step, ten-minute charging protocols from 0% to 80% SOC, as detailed below. After a five-second rest, all cells then charged from 80% to 100% SOC with a 1C CC-CV charging step to 3.6 V and a current cutoff of $C/20$. After another five-second rest, all cells subsequently discharged with a CC-CV discharge at 4C to 2.0 V and a current cutoff of $C/20$. The cells rested for another five seconds before the subsequent charging step started. The lower and upper cutoff voltages were 2.0 V and 3.6 V, respectively, as recommended by the manufacturer. In this work, cycle life is defined as the number of cycles until the discharge capacity falls below 80% of the nominal capacity.

All cells were tested in cylindrical fixtures with 4-point contacts on a 48-channel Arbin Laboratory Battery Testing battery cyler placed in an environmental chamber (Amerex Instruments) at 30 $^{\circ}\text{C}$. The cyler calibration was validated before the state of the experiment.

In the closed-loop experiment, four experiments did not reach 100 cycles owing to contact issues either at the start or partially through the experiment. These experiments were run on channels 17 and 27 in round 1 (oed_0) and channels 4 and 5 in round 2 (oed_1). Additionally, in each round, one protocol per round that should have been selected (that is, with a top-48 upper bound) was not selected and replaced with the protocol with the 49th-highest upper bound owing to a processing error (Extended Data Fig. 4), but this error is not expected to have a large effect. Additional experimental issues are documented in the notes of the data repository.

Charging protocol and parameter space design

Cells were charged with one of 224 different four-step charging protocols. Each of the first four steps is a single constant-current step applied over a 20% SOC range; thus, the 224 charging protocols represent different combinations of current steps within the 0% to 80% SOC range. We can define the charging time from 0% to 80% SOC by:

$$t_{0-80\%} = \frac{0.2}{CC1} + \frac{0.2}{CC2} + \frac{0.2}{CC3} + \frac{0.2}{CC4}$$

In all protocols considered here, we constrain $t_{0-80\%}$ to be 10 min. We now write CC4 as a function of the first three charging steps, as:

$$CC4 = \frac{0.2}{\frac{10}{60} - \left(\frac{0.2}{CC1} + \frac{0.2}{CC2} + \frac{0.2}{CC3} \right)}$$

Thus, each protocol can be uniquely defined by CC1, CC2 and CC3.

Each independent parameter can take on one of the following discrete values: 3.6C, 4.0C, 4.4C, 4.8C, 5.2C and 5.6C. Furthermore, CC1 can take on values of 6.0C, 7.0C and 8.0C, and CC2 can take on values of 6.0C and 7.0C. CC4 is not allowed to exceed 4.81C. The maximum allowable current for each parameter decreases with increasing SOC to avoid reaching the upper cutoff voltage of 3.6 V. With these constraints, a total of 224 charging protocols are permitted.

For a consistent protocol nomenclature, we define each fast-charging protocol as CC1-CC2-CC3-CC4. For example, the charging protocol with the highest CLO-estimated mean cycle life is written 4.8C-5.2C-5.2C-4.160C.

Early outcome predictor

The early outcome predictor for cycle life is similar to that presented in Severson et al.⁵. This linear model predicts the final \log_{10} cycle life (number of cycles to reach 80% of nominal capacity, or 0.88 A h) using features from the first 100 cycles. The training set is identical to the one used in Severson et al.⁵ and consists of 41 batteries. The linear model takes the form:

$$\hat{y}_i = \hat{\mathbf{w}}^T \mathbf{x}_i$$

Here \hat{y}_i is the predicted cycle life for battery i , \mathbf{x}_i is a p -dimensional feature vector for battery i and $\hat{\mathbf{w}}$ is a p -dimensional model coefficient vector. Features are z-scored (mean-subtracted and normalized by the standard deviation) to the training set before model evaluation.

Regularization, or simultaneous feature selection and model fitting, was performed using the elastic net²⁷. Regularization penalizes overly complex fits to improve both generalizability and interpretability. Specifically, the coefficient vector $\hat{\mathbf{w}}$ is found via the following expression:

$$\hat{\mathbf{w}} = \operatorname{argmin}_{\mathbf{w}} \left[\|\mathbf{y} - \mathbf{X}\mathbf{w}\|_2^2 + \lambda \left(\frac{1-\alpha}{2} \|\mathbf{w}\|_2^2 + \alpha \|\mathbf{w}\|_1 \right) \right]$$

Here λ and α are hyperparameters; λ is a non-negative scalar and α is a scalar between 0 and 1. The first term minimizes the squared loss, and the second term performs both continuous shrinkage and automatic feature selection. During model development, we apply fourfold cross-validation and Monte Carlo sampling with the training set to optimize the values of the hyperparameters λ and α .

As in Severson et al.⁵, the available features were based on the difference between discharge voltage curves of cycles 100 and 10, or trends in the discharge capacity. The five selected features, their corresponding weights and the z-scored values are presented in Supplementary Table 1. The training (cross-validated) error was 80.4 cycles (10.2%); the test error on a test set from Severson et al.⁵ was 122 cycles (12.6%).

The early predictor automatically flags predictions as anomalous if the 95% prediction interval exceeds 2,000 cycles. The two-tailed 95% prediction interval is computed by:

$$95\% \text{PI} = 2t_{(\alpha/2, n-p)} \times \text{RMSE} \sqrt{1 + \mathbf{x}_i^T (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{x}_i}$$

where t is the Student's t value, α is the significance level (0.05 for 95% confidence), n is the number of samples, p is the number of features, RMSE is the root-mean-square error of the training set (in units of cycles), \mathbf{x}_i is the vector of selected features for battery i and \mathbf{X} is the matrix of selected features for all observations in the training set.

In the closed-loop experiment, three tests returned predictions with a prediction interval outside of the threshold; these anomalous predictions were excluded. These tests were run on channel 27 in round 1 (oed_0), channel 12 in round 3 (oed_2) and channel 6 in round 4 (oed_3). Furthermore, in the validation experiment, one test returned a prediction with a prediction interval outside of the threshold (channel 12; 3.6C-6.0C-5.6C-4.755C), although the final cycle life was reasonable.

We note that the predictions from this model exhibited systematic bias for the cells in the validation experiments, which we attribute to the increased calendar ageing of these cells relative to the training set (Supplementary Table 2 and Supplementary Discussion 4).

Bayesian optimization algorithm

To perform optimal experimental design, we consider the setting of best-arm identification using multi-armed bandits. Here each arm is a charging protocol and the goal is to identify the best arm, or equivalently the charging protocol with the highest expected cycle life. Many variants of the problem have been studied in prior work^{33–35}; our algorithm builds on the approaches of Hoffman et al.¹⁰ and Grover et al.¹¹. We consider further modifications in Supplementary Discussion 2.

In particular, we assume a Bayesian regression setting, where there exists an unknown set of parameters ($\theta \in \mathcal{R}^d$) that relate a charging protocol x to its cycle life (a scalar) via a Gaussian likelihood function. Here, x denotes the CC1, CC2, CC3 configurations of a charging protocol, which is projected onto a d -dimensional feature vector $\varphi(x)$. We set $d = 224$, and the feature representations $\varphi(x)$ are obtained by approximating a radial-basis function kernel, $K(x_i, x_j) = \exp(\gamma \|x_i - x_j\|_2^2)$, using Nystroem's method. Here, x_i and x_j are the CC1, CC2 and CC3 configurations for two arbitrary charging protocols and the inverse of the kernel bandwidth, $\gamma > 0$ is treated as a hyperparameter.

The Gaussian likelihood function relates a charging protocol to its cycle life distribution. For a protocol x , the mean of this likelihood function is given as $\theta^T \varphi(x)$. The variance of this likelihood function is the sum of two uncertainty terms, both of which we assume to be homoskedastic (that is, uniform across all protocols). The first term is the empirical variance averaged across the repeated runs of individual protocols present in the training dataset (same as that used for training the early predictor). This accounts for variability due to exogenous factors such as manufacturing. Second, since we do not wait for an experiment to complete, the likelihood variance additionally needs to accommodate an additional uncertainty term due to the early outcome prediction component of the pipeline. We do so by computing the residual variance of the early predictions on the held-out portion of the dataset and set the aforementioned uncertainty term to be the maximum of the residual variances. We assume that the two sources of uncertainty are independent, and hence the overall variance of the likelihood distribution is given by the sum of the squares of both variance terms described above.

To perform inference over the unknown parameters θ and subsequent predictions of cycle lives, we employ a Gaussian process. In a Gaussian process, the prior over θ is assumed to be isotropic Gaussian; such a prior is conjugate to the Gaussian likelihood, and as a consequence the Gaussian posterior can be obtained in closed-form via the Bayes rule. This posterior is used to define a Gaussian predictive distribution over the cycle life for any given charging protocol with mean μ and variance σ^2 .

Finally, to select a charging protocol, we optimize an acquisition function based on upper confidence bounds. The acquisition function selects protocols where the noisy predictive distribution over cycle life has high mean μ (to encourage exploitation) and high variance σ^2 (to encourage exploration). The mean and upper and lower confidence bounds for any arm i is given by $\mu_{k,i} \pm \beta_k \sigma_{k,i}$ at round k , such that the relative weighting of the two terms is controlled by the exploration tradeoff hyperparameter, $\beta > 0$. The exploration tradeoff hyperparameter at round k , β_k , is decayed multiplicatively at every round of the closed loop by another hyperparameter, $\varepsilon \in (0,1]$, as given by $\beta_k = \beta_0 \varepsilon^k$.

BO hyperparameter optimization

The BO algorithm relies on eight hyperparameters, each of which can be categorized as either a resource hyperparameter, a parameter space hyperparameter or an algorithm hyperparameter. We note that the BO algorithm runs in the fixed-budget setting; here, the budget refers to the number of iterations of the closed loop we run, excluding validation experiments. We describe each category of hyperparameters below; the values of each hyperparameter are tabulated in Supplementary Table 5.

Resource hyperparameters are specified by the available testing resources. The 'batch size' represents the number of parallel tests. We set a batch size of 48 given our 48-channel battery cycler. The 'budget' represents the number of batches tested during CLO. The budget excludes batches used to develop the early predictor and validation batches. The budget is typically constrained by either the available testing time or the number of cells. In this case, we set a budget of 4, yielding a cell budget of 192 cells and a time budget of 16 days (4 days per batch of 48 cells tested for 100 cycles).

Parameter space hyperparameters are specified by the optimization problem. Here, we use the same data available from the training set of the early predictor to estimate these parameters, despite a different charging protocol structure. The 'standardization mean' represents the estimated mean cycle life across all protocols. The 'standardization standard deviation' represents the estimated standard deviation of cycle life across all protocols; in other words, this parameter represents the range of cycle lives in the parameter space. The 'likelihood standard deviation' represents the estimated standard deviation of a single protocol tested multiple times, which is a measure of the sampling error; this sampling error includes both the intrinsic variability and the prediction error.

Algorithm hyperparameters control the performance of the Bayesian optimization algorithm. γ is the kernel bandwidth, which controls the interaction strength between neighbouring protocols in the parameter space. High γ favours under-smoothing of the parameter space, that is, the protocols have weak relationships with their neighbours. β_0 represents the initial value of β , the exploration tradeoff hyperparameter; β controls the balance of exploration versus exploitation. High β_0 favours exploration over exploitation. ε represents the decay constant of beta per round; as the experiment progresses, ε shifts towards stronger exploitation (given by $\beta_k = \beta_0 \varepsilon^k$, where β_k represents the exploration constant at round k , 0-indexed). High ε favours a rapid transition from exploration to exploitation.

The algorithm hyperparameters were estimated by creating a physics-based simulator based on the range of cycle lives obtained in the preliminary batch, testing all hyperparameter combinations on the simulator, and selecting the hyperparameter combination with the best performance (that is, that which most consistently obtains the true cycle life). These results are visualized in Extended Data Fig. 9; we note that the performance of BO is relatively insensitive to the selected combination of algorithm hyperparameters, meaning sufficiently high performance can be achieved even with suboptimal algorithm hyperparameters. Other approaches, such as using the early-predictor training dataset, are also possible for optimization of the algorithm hyperparameters (Supplementary Discussion 1).

Physics-based simulator

We used a physics-based simulator for hyperparameter optimization; this simulator allows us to estimate the shape and range of cycle lives in the parameter space, although the simulator is not designed to be an accurate representation of battery degradation during fast charging. This finite element simulator was originally designed to estimate the heat generation during charging in an 18650 cylindrical battery by approximating the battery as a long cylinder, which simplifies to a one-dimensional radial heat transfer problem. The equations and thermal properties were sourced from Drake et al.³⁶ and Çengel and Boles³⁷. The output from these simulations is a matrix of temperature as a function of both radial position and time. We use total solid-electrolyte interphase (SEI) growth as a proxy for degradation. First, we estimate the temperature dependence of SEI growth from the C/10 series of figure 7 from Smith et al.³⁸ (Supplementary Table 6). Simultaneously, we compute the expected temperature profiles in the battery as a function of charging protocol with respect to time and position. We then approximate the kinetics of SEI growth with an Arrhenius equation, such that SEI growth increases with increasing temperature. SEI growth

(in arbitrary units) is calculated for each temperature element in the position-time array via:

$$D = \sum_r \sum_t \exp\left(-\frac{E_a}{k_B T}\right)$$

where D is the degradation parameter, E_a is the effective activation energy for SEI growth (Supplementary Table 6) and k_B is Boltzmann's constant. The cycle life is then calculated from the degradation parameter using the range of expected cycle lives (as estimated from the early-predictor training dataset):

$$\text{Cycle life} = 500 + C/D$$

where C is a constant (5×10^{-11}) that scales D to reasonable values of cycle life.

Validation experiments

After the closed-loop experiment completed, we selected nine protocols to test to failure (five batteries per charging protocol). This experiment allowed us to (1) evaluate the performance of the closed loop by comparing the CLO-estimated mean cycle lives to the mean cycle life of multiple batteries tested to failure for multiple protocols, (2) compare the protocols with the highest CLO-estimated mean cycle lives to conventional fast-charging protocol design principles from the battery literature, and (3) generate a small dataset with which we can evaluate the performance of the closed loop relative to baseline optimization approaches.

The selected protocols are displayed in Extended Data Fig. 6 and Extended Data Table 1. Of our nine fast-charging protocols, three were the top three CLO-estimated protocols; four were based on approximations of multi-step fast-charging protocols in the battery literature (see Extended Data Table 1); and two were selected to obtain a representative sampling from the distribution of CLO-estimated cycle lives. The four protocols based on approximations of multi-step fast-charging protocols in the battery literature were obtained by determining the current ratios between various steps and translating those ratios to our ten-minute fast-charging space. The voltage limits were consistent with our charging protocols, that is, 2.0 V and 3.6 V.

Five batteries per charging protocol were tested to obtain a reasonable estimate of the true cycle lives. In this experiment, one test returned a prediction with a prediction interval outside of the threshold (channel 12; 3.6C-6.0C-5.6C-4.755C) and was excluded. A comparison of the three different methods for cycle life results (CLO, early predictions from validation, and final measurements from validation) are presented in Extended Data Fig. 7.

Validation ablation study

For the ablation study using the charging protocols and data from the validation experiments, we systematically compared the full closed-loop system against three other ablation baselines which use (1) only early prediction (no BO exploration-exploitation, purely random exploration), (2) only BO exploration-exploitation (no early prediction), (3) purely random exploration without any early prediction. As highlighted earlier, since the final cycle lives for the protocols in the validation study have a noticeable bias that can be explained by calendar ageing (Supplementary Discussion 4), we perform a simple additive bias correction for each of the final cycle lives beforehand to suppress any undesirable influence of this bias in interpreting the results.

We run the four ablation baselines for a varying number of sequential rounds. Since our validation space is relatively small (nine charging protocols, five batteries tested per protocol in our validation dataset), we run only one battery per round (that is, we assume a one-channel battery cyler). The baselines that use BO exploration-exploitation additionally require hyperparameters to be specified before beginning

the experiment, as described in the Methods section 'BO hyperparameter optimization'. The best hyperparameters are chosen separately for each round based on the performance obtained on the physics-based simulator, averaged over 100 random seeds.

When an ablation baseline queries for the cycle life of a given charging protocol, the returned value corresponds to one of the five runs in our validation dataset, chosen via random sampling with replacement (that is, bootstrapped). The experimental time cost of this query is equal to 100 cycles for ablation baselines that use early prediction and equals the full cycle life otherwise. Finally, to account for the randomness at the beginning of the experiment (that is, round 0 when every ablation baseline randomly selects a protocol), we report the performance of each ablation baseline averaged over a sequence of 2,000 randomly initialized experiments. To specify the y-axis of Fig. 4e, we assume that each full cycle (charging, discharging, resting) requires one hour of experimental testing.

Overpotential analysis

To determine the dependence of overpotential on current and SOC during charging (Extended Data Fig. 2e-f), we perform a pseudo-galvanostatic intermittent titration technique experiment on two minimally cycled batteries and two degraded batteries (80% of nominal capacity remaining). We probe currents ranging from 3.6C to 8C at 20%, 40%, 60% and 80% SOC, mirroring the current and SOC values used in charging protocol design. In this experiment, we start at an initial SOC 20% lower than the target, for example, we start at 0% SOC to probe 20% SOC. We then charge at a given current rate, for example, 3.6C, until we reach 20% SOC. The cell rests for 1 h, and then the cell discharges at 1C back to 0% SOC. We repeat this sequence for all current values, after which we charge the cell at 1C to the next initial SOC, for example, 20% SOC to probe 40% SOC, and repeat for each SOC of interest.

To compute the overpotential, we compare the voltage at the start and end of the 1-h rest periods. Nearly all of the potential drop occurs immediately (<100 ms) after the start of the rest period. Given the linear trends observed (implying ohmic-limited rate capability), we then perform a linear fit on each overpotential-current series. In these fits, the slope represents the ohmic resistance.

Data availability

The datasets used in this study are available at <https://data.matr.io/1>.

Code availability

The CLO code, data and figures associated with this manuscript are available at <https://github.com/chueh-ermon/battery-fast-charging-optimization>. The data processing and early-prediction code are available at <https://github.com/chueh-ermon/BMS-autoanalysis>. The charging protocol generation code (automated creation of battery cyler tests) is available at <https://github.com/chueh-ermon/automate-Arbin-schedule-file-creation>.

- Shahriari, B., Swersky, K., Wang, Z., Adams, R. P. & de Freitas, N. Taking the human out of the loop: a review of Bayesian optimization. *Proc. IEEE* **104**, 148–175 (2016).
- Audibert, J.-Y., Bubeck, S. & Munos, R. Best arm identification in multi-armed bandits. In *Proc. 23rd Conf. on Learning Theory (COLT)* 41–53 (2010); <http://certis.enpc.fr/~audibert/Mes%20articles/COLT10.pdf>.
- Srinivas, N., Krause, A., Kakade, S. M. & Seeger, M. W. Information-theoretic regret bounds for Gaussian process optimization in the bandit setting. *IEEE Trans. Inf. Theory* **58**, 3250–3265 (2012).
- Drake, S. J. et al. Measurement of anisotropic thermophysical properties of cylindrical Li-ion cells. *J. Power Sources* **252**, 298–304 (2014).
- Çengel, Y. A. & Boles, M. A. *Thermodynamics: An Engineering Approach* (McGraw-Hill Education, 2015).
- Smith, A. J., Burns, J. C., Zhao, X., Xiong, D. & Dahn, J. R. A high precision coulometry study of the SEI growth in Li/graphite cells. *J. Electrochem. Soc.* **158**, A447–A452 (2011).
- Zhang, S. S. The effect of the charging protocol on the cycle life of a Li-ion battery. *J. Power Sources* **161**, 1385–1391 (2006).

Article

40. Kim, J. M. et al. Battery charging method and battery pack using the same. US Patent Application US20160226270A1 (2016).
41. Lee, M.-S., Song, S.-B., Jung, J.-S. & Golovanov, D. Battery charging method and battery pack using the same. US Patent US9917458B2 (2018).
42. Notten, P. H. L., Op het Veld, J. H. G. & van Beek, J. R. G. Boostcharging Li-ion batteries: a challenging new charging concept. *J. Power Sources* **145**, 89–94 (2005).
43. Paryani, A. Low temperature charging of Li-ion cells. US Patent US8552693B2 (2013).
44. Mehta, V. H. & Straubel, J. B. Fast charging with negative ramped current profile. US Patent US8643342B2 (2014).

Acknowledgements This work was supported by the Toyota Research Institute through the Accelerated Materials Design and Discovery programme. P.M.A. was supported by the Thomas V. Jones Stanford Graduate Fellowship and the National Science Foundation Graduate Research Fellowship under grant number DGE-114747. A.G. was supported by a Microsoft Research PhD Fellowship and a Stanford Data Science Scholarship. N.P. was supported by the SAIC Innovation Center through the Stanford Energy 3.0 industry affiliates programme. S.J.H. was supported by the Assistant Secretary for Energy Efficiency, Vehicle Technologies Office of the US Department of Energy under the Advanced Battery Materials Research Program. X-ray tomography was performed at the Stanford Nano Shared Facilities, supported by the National Science Foundation under award ECCS-1542152. We thank A. Anapolsky, L. Attia, C. Cundy, J.

Hirshman, S. Jorgensen, G. McConohy, J. Song, R. Smith, B. Storey and H. Thaman for discussions.

Author contributions P.M.A., N.J., Y.-H.L., M.H.C., N.P. and W.C.C. conceived and conducted the experiments. A.G., T.M.M., B.C. and S.E. developed the Bayesian optimization algorithm and incorporated early outcome predictions into the closed loop. P.M.A. and K.A.S. performed the early-prediction modelling. P.M.A., Z.Y., P.K.H. and M.A. performed data management. P.M.A., A.G., N.J., S.J.H., S.E. and W.C.C. interpreted the results. All authors edited and reviewed the manuscript. R.D.B., S.E. and W.C.C. supervised the work.

Competing interests S.E., W.C.C., A.G., T.M.M., N.P. and P.M.A. have filed a patent application related to this work: US Patent Application No. 16/161,790 (16 October 2018).


Additional information

Supplementary information is available for this paper at <https://doi.org/10.1038/s41586-020-1994-5>.

Correspondence and requests for materials should be addressed to R.D.B., S.E. or W.C.C.

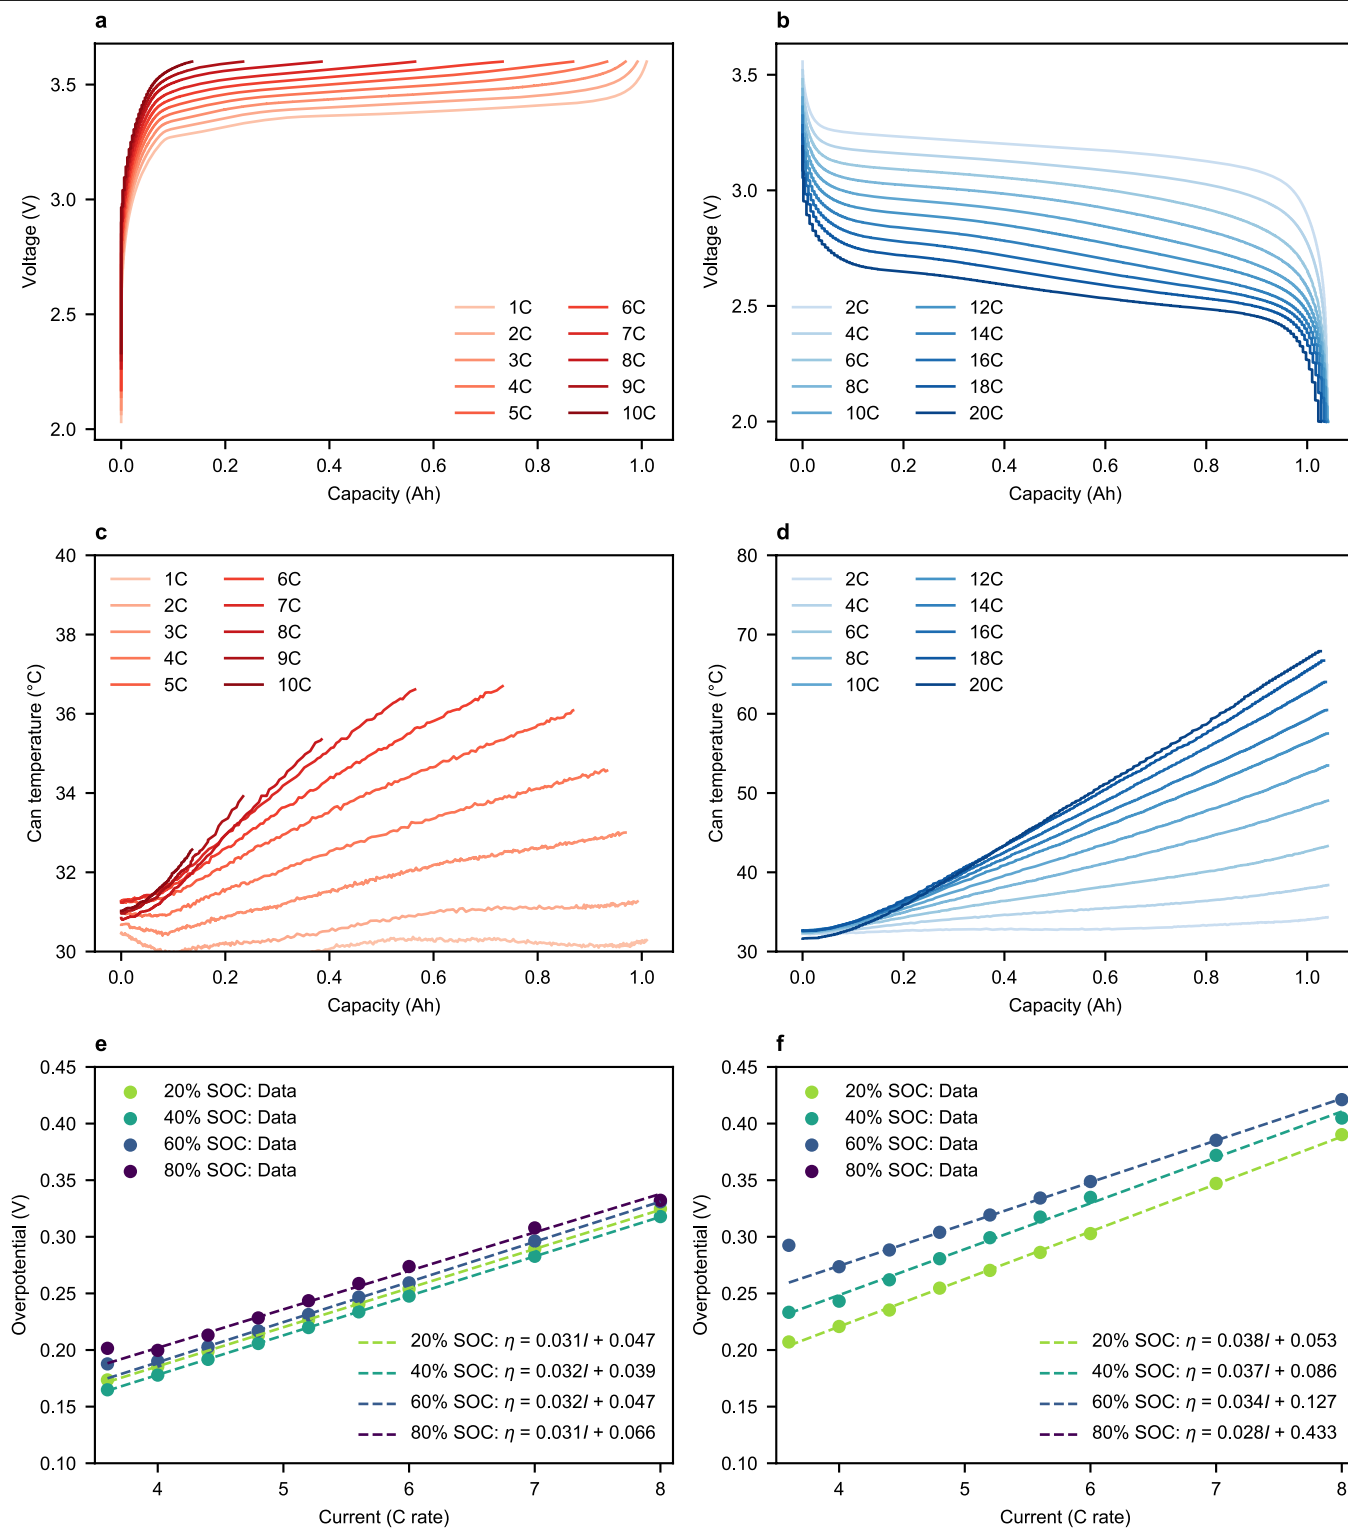
Peer review information *Nature* thanks Marius Bauer, Matthias Seeger and the other, anonymous, reviewer(s) for their contribution to the peer review of this work.

Reprints and permissions information is available at <http://www.nature.com/reprints>.




Extended Data Fig. 1 | Illustrations of early outcome predictor and BO components of CLO. a, Illustration of early outcome prediction for two cells (A and B) using data from only the first 100 cycles. Two discharge capacity features are generated: the second-cycle discharge capacity, $Q_{d,2}$, and the difference between the maximum and second-cycle discharge capacities, $\max(Q_d) - Q_{d,2}$. Three voltage features are generated: the logarithm of the minimum, variance and the skewness of the difference in voltage curves between cycles 100 and 10. These five features are combined in a linear model to predict the final cycle life, or the number of cycles until the capacity falls below 0.88 A h. The weights and scalings of each feature are determined by training the model on a training set using the elastic net; the weights and scaling values are presented in Supplementary Table 1. See Severson et al.⁷ and

Methods for additional details. **b**, Illustration of the BO principle. The desired output, cycle life, has a true functional dependence on charging protocol parameters (such as CC1). Here, we show a one-dimensional model (that is, just dependent on one parameter, CC1) for simplicity. By performing Gaussian process regression on the available data, we develop a probabilistic estimate of the true function; our goal is for the estimate to match the true function. The next data point selected is that which maximizes the upper confidence bound (UCB), which is selected by either high uncertainty (exploration) or high predicted value (exploitation). Once this point is selected (right panel), the next point selected is, again, that which maximizes the upper confidence bound.




Extended Data Fig. 2 | Cell characterization. **a, b**, Voltage versus capacity during rate testing of A12318650M1A cylindrical cells under charge (**a**) and discharge (**b**). The (dis)charge step not under investigation is cycled at 1C to isolate the rate of each step; for example, the charge rate test is performed with 1-C discharge steps. We note that the discharge rate capability is much higher than that of charge. **c, d**, Battery surface temperature ('can temperature') versus capacity during rate testing under charge (**c**) and discharge (**d**). The can temperature is measured via a type T thermocouple secured with thermal

epoxy. **e, f**, Overpotential as a function of SOC and C rate (see Methods section 'Overpotential analysis' for details of the measurement) for a minimally cycled cell (**e**) and an aged cell at 80% of nominal capacity (**f**). The trend lines are linear fits of the overpotential as a function of current at fixed SOC (excluding outliers). We note that both of the relationships are linear (indicating that the rate capability is ohmically limited) and that the SOC dependence is weak, particularly for the minimally cycled cell. The initial internal resistance, averaged over two cells and all four SOCs, is 33 mΩ.



Extended Data Fig. 3 | Additional optimization results. **a, b**, Mean of the absolute difference in CLO-estimated cycle lives with increasing rounds, expressed as both percentage change (**a**) and absolute change (**b**). These changes are relatively small beyond round 2, suggesting that the closed loop can perform well with even smaller time or battery budgets. **c**, Change in Kendall rank correlation coefficient with increasing rounds. From round 3 to round 4, the ranking of the top protocols shifts, but the cycle lives of these top protocols are similar. **d**, Distribution of CLO-estimated mean cycle lives after round 4. The mean and standard deviation are 943 cycles and 126 cycles, respectively. **e**, Correlation between CLO-estimated mean cycle lives and the sum of squared currents, a simplified measure of heat generation ($P = I^2 R$). This relationship suggests that minimizing heat generation, as opposed to avoiding


lithium plating, may be the operative optimization strategy for these cells under these conditions. **f**, Standard deviation ($\sigma_{4,i}$) versus mean ($\mu_{4,i}$) of the BO predictive distribution over cycle life after round 4. The standard deviation quantifies the uncertainty in the cycle life estimates and is generally low for protocols estimated to have high mean cycle life, since these protocols are probed more frequently. We start with a relatively wide, flat prior (standard deviation 164) and therefore the uncertainty intervals after four rounds are also wide. **g**, Mean \pm standard deviation of the predictive distribution over cycle life after round 4 ($\mu_{4,i} \pm \sigma_{4,i}$) for all charging protocols, sorted by their rank after round 4. The legend indicates the number of repetitions for each protocol (excluding failed batteries).



Extended Data Fig. 4 | See next page for caption.


Extended Data Fig. 4 | Means and upper/lower confidence bounds ($\mu_{k,i} \pm \beta_k \sigma_{k,i}$) on cycle life per round k . Protocol indices on the x -axis are sorted by rank after round 4. The weighted interval around the estimated mean, $\beta_k \sigma_{k,i} = (\beta_0 \varepsilon^k) \sigma_{k,i}$, weights the protocol-specific standard deviation at round k , $\sigma_{k,i}$ (estimated by the Gaussian process model) with the exploration tradeoff hyperparameter at round k , β_k . The upper and lower confidence bounds are plotted for all charging protocols before round 1 (**a**) and after rounds 1 (**b**), 2 (**c**), 3 (**d**) and 4 (**e**). The predictive distributions for all charging protocols have identical means and standard deviations before the first round of testing. Because the standard deviations are weighted by $\beta_k = \beta_0 \varepsilon^k$ and $\varepsilon = 0.5$, the

weighted confidence bounds rapidly decrease with increasing round number, favouring exploitation (examination of protocols with high means). The BO algorithm recommends the 48 protocols with the highest upper bounds (red points); the upper bounds are high either due to high uncertainty (exploration) or high means (exploitation). The algorithm rapidly shifts from exploration to exploitation as ε_k rapidly shrinks the upper bounds with increasing round index. We note that one protocol per round that should have been selected (that is, with a top-48 upper bound) was not selected owing to a processing error; instead, the protocol with the 49th-highest upper bound was selected.




Extended Data Fig. 5 | Mean and standard deviation of the CLO-estimated predicted distribution over cycle lives after round 4. In this two-dimensional representation, mean estimated cycle life (colour scale) and standard deviation of cycle life (marker size) after round 4 are presented as a function of CC1, CC2 and CC3 (the x axis, y axis and panels a–f, respectively). Panels a–f represent

CC3 = 3.6C, 4.0C, 4.4C, 4.8C, 5.2C, 5.6C and 6.0C, respectively. CC4 is represented by the contour lines. Note that the protocols with the highest cycle lives generally have the smallest standard deviations, since these protocols have been tested repeatedly.




Extended Data Fig. 6 | Selected protocols for validation. The three protocols with the highest CLO-estimated mean cycle lives are shown in panels **b**, **c** and **d**. The protocols shown in panels **a**, **f**, **g** and **h** are approximations of previously proposed battery fast-charging protocols (Extended Data Table 1). The remaining two protocols, shown in panels **e** and **i**, were selected to obtain a representative sampling from the entire distribution of CLO-estimated cycle lives. The annotations on each panel represent the cycle lives of each protocol

as estimated by CLO ('CLO'), early outcome prediction from validation ('Early prediction'), and the final cycle lives from validation ('Final'). In the annotations, the errors represent the CLO-estimated standard deviation after round 4 ($\sigma_{k,4}$) for the CLO-estimated cycle lives and the 95% confidence intervals for the early-predicted and final cycle lives from validation ($n = 5$; $n = 4$ for the early predictions of 3.6C-6.0C-5.6C-4.755C) (**a**).



Extended Data Fig. 7 | Validation ablation analysis. We perform pairwise comparisons of the cycle lives of the nine validation protocols, as estimated from three sources: closed-loop estimates after four rounds, early predictions from the validation experiment and final cycle lives from the validation experiment. Panels **a–c** compare closed-loop estimates to early predictions from validation, panels **d–f** compare final cycle lives from validation to early predictions from validation, and panels **g–i** compare final cycle lives from validation to closed-loop estimates. The first column (**a, d** and **g**) compares cycle lives averaged on a protocol basis; the second column (**b, e** and **h**)


compares cycle lives on a battery (cell) basis; and the third column (**c, f** and **i**) compares the predicted ranking by cycle life via each method. Orange points represent the top three CLO-estimated protocols, blue points represent protocols inspired by the battery literature (Methods), and green points represent protocols selected to sample the distribution of estimated cycle lives. The error bars represent the 95% confidence intervals ($n = 5$; $n = 4$ for the early predictions of 3.6C-6.0C-5.6C-4.755C). The Pearson correlation coefficient and Kendall rank correlation coefficients are displayed for all relevant cycle life and ranking plots, respectively.



Extended Data Fig. 8 | Closed-loop performance under resource constraints.

Comparison of the closed loop with and without the Bayesian optimization algorithm (that is, with and without the explore/exploit component) as a function of number of channels and number of rounds in the 224-protocol space, using the first-principles simulator as the ground-truth source for cycle lives. Early prediction is not included. Each point represents the mean of 100

simulations; error bars represent the 95% confidence intervals ($n=100$). Early prediction is not incorporated into these simulations. The complete closed loop (that is, with Bayesian optimization) consistently outperforms the closed loop without Bayesian optimization. Bayesian optimization offers the largest advantage when the number of channels is low relative to the number of protocols.



Extended Data Fig. 9 | Hyperparameter sensitivity analysis on a cycle life simulator. The true cycle life of the best charging protocol as estimated by CLO, averaged over ten random seeds, is plotted as a function of the initial exploration constant (β_0), the exploration decay factor (ϵ) and the kernel bandwidth (γ). The values of all other hyperparameters are consistent with the values indicated in the ‘BO hyperparameter optimization’ Methods section and

in Supplementary Table 5. Overall, CLO achieves acceptable performance over a range of hyperparameter combinations; the highest-cycle-life protocols as estimated by the best and worst hyperparameter combinations differ by only 60 cycles. In our real-world CLO experiment, the selected hyperparameters are $\beta_0 = 5.0$, $\epsilon = 0.5$ and $\gamma = 1$; this combination performed well on a variety of simulated parameter spaces and budgets.

Extended Data Table 1 | Selected charging protocols for validation

Charging protocol	CLO-estimated cycle life	Early-predicted cycle life (from validation)	Final cycle life (from validation)	Source
3.6C-6.0C-5.6C-4.755C	1103 ± 131	1013 ± 115	755 ± 81	Zhang ³⁹
4.4C-5.6C-5.2C-4.252C	1174 ± 76	1056 ± 127	884 ± 132	Protocol with third-highest CLO-estimated mean cycle life
4.8C-5.2C-5.2C-4.160C	1185 ± 78	1047 ± 49	890 ± 90	Protocol with highest CLO-estimated mean cycle life
5.2C-5.2C-4.8C-4.160C	1183 ± 86	1098 ± 134	912 ± 118	Protocol with second-highest CLO-estimated mean cycle life
6.0C-5.6C-4.4C-3.834C	954 ± 164	963 ± 26	880 ± 85	
7.0C-4.8C-4.8C-3.652C	876 ± 183	964 ± 43	870 ± 70	Samsung patents ^{40,41}
8.0C-4.4C-4.4C-3.940C	818 ± 212	854 ± 44	702 ± 51	Notten <i>et al.</i> ⁴²
8.0C-6.0C-4.8C-3.000C	775 ± 273	698 ± 40	584 ± 60	Tesla patents ^{43,44}
8.0C-7.0C-5.2C-2.680C	648 ± 174	580 ± 68	496 ± 49	

The columns represent the CLO-estimated mean cycle lives of each protocol, early predictions in the validation experiment and the final tested cycle lives. For the CLO-estimated cycle lives, the errors represent the CLO-estimated standard deviation after round 4 ($\sigma_{4,d}$). For the early-predicted and final cycle lives from validation, the errors represent 95% confidence intervals ($n = 5$; but $n = 4$ for the early predictions of 3.6C-6.0C-5.6C-4.755C). The two protocols without a source were selected to obtain a representative sampling from the distribution of CLO-estimated cycle lives. Literature fast-charging protocols are from refs. ³⁹⁻⁴⁴.