

Bayesian Model Selection of Lithium-Ion Battery Models via Bayesian Quadrature

Masaki Adachi ^{*,*,*,*}

Yannick Kuhn ^{****,†} Birger Horstmann ^{****,†}

Arnulf Latz ^{****,†} Michael A. Osborne ^{*} David A. Howey ^{*,‡}
^{*} Machine Learning Research Group, University of Oxford, OX2 6ED, UK (e-mail: masaki@robots.ox.ac.uk).

^{**} Battery Intelligence Lab, University of Oxford, OX1 3PJ, UK

^{***} Toyota Motor Corporation, Shizuoka 410-1193, Japan

^{****} German Aerospace

Center (DLR), Pfaffenwaldring 38-40, 70569 Stuttgart, Germany
Helmholtz Institute Ulm, Helmholtzstraße 11, 89081 Ulm, Germany

[†] Universität Ulm, Albert-Einstein-Allee 47, 89081 Ulm, Germany

[‡] The Faraday Institution, Harwell Campus, Didcot OX11 0RA, UK

Abstract: A wide variety of battery models are available, and it is not always obvious which model ‘best’ describes a dataset. This paper presents a Bayesian model selection approach using Bayesian quadrature. The model evidence is adopted as the selection metric, choosing the simplest model that describes the data, in the spirit of Occam’s razor. However, estimating this requires integral computations over parameter space, which is usually prohibitively expensive. Bayesian quadrature offers sample-efficient integration via model-based inference that minimises the number of battery model evaluations. The posterior distribution of model parameters can also be inferred as a byproduct without further computation. Here, the simplest lithium-ion battery models, equivalent circuit models, were used to analyse the sensitivity of the selection criterion to given different datasets and model configurations. We show that popular model selection criteria, such as root-mean-square error and Bayesian information criterion, can fail to select a parsimonious model in the case of a multimodal posterior. The model evidence can spot the optimal model in such cases, simultaneously providing the variance of the evidence inference itself as an indication of confidence. We also show that Bayesian quadrature can compute the evidence faster than popular Monte Carlo based solvers.

Copyright © 2023 The Authors. This is an open access article under the CC BY-NC-ND license (<https://creativecommons.org/licenses/by-nc-nd/4.0/>)

Keywords: Bayesian, identifiability, system identification, estimation, battery, lithium-ion

1. INTRODUCTION

The lithium-ion battery is key to decarbonising power grids and electrifying vehicles. However, its behaviour can be challenging to model, control, and diagnose, and this is a practical hindrance to obtaining the optimal performance. This is compounded by the available data from operational batteries being typically limited to just three measurements: voltage, current, and temperature. Estimating internal states from these time-varying three variables is challenging or even mathematically impossible due to parameter identifiability issues (Bizeray et al., 2018). Degradation further complicates matters since the number of parameters to be identified becomes larger when considering long-term ageing.

There are dozens of plausible models for Li-ion batteries, owing to differing assumptions and levels of approximation. While electrochemists might prefer continuum

models, such as the Doyle-Fuller-Newman model (Doyle et al., 1993), that give understanding of internal chemical reactions and transport, control engineers prefer simpler approaches such as equivalent circuit models (ECMs) (He et al., 2011), for fast control and fewer parameters. Other models exist in a spectrum between these (e.g. from simple to more complex: ECM → EHM (Milocco et al., 2014) → SPM (Santhanagopalan et al., 2006) → SPM_e (Kemper et al., 2013) → DFN). System identification is the foundation of an estimation and control system, determining predictive accuracy, quick response, and reliability.

However, the ‘best’ model should be ascertained based on quantifiable performance metrics. Importantly, the optimal model strongly depends on the dataset **D** and user requirements. A widely accepted approach for defining ‘good’ models is Occam’s razor, where the simplest model to reasonably reproduce a given dataset is considered the best. Simplest here relates to the number of parameters to be identified. Rasmussen et al. (2000) showed that such a metric could be evaluated via Bayesian model evidence, obtained for a model *M* by integrating out (i.e. averaging over) the parameters θ from the likelihood,

^{*} The authors acknowledge support by Toyota Motor Corporation, Oxford Clarendon Fund, Oxford Kobe scholarship, the German Aerospace Center (DLR), and the Helmholtz Association through grant no KW-BASF-6, and contributes to the research performed at CELEST

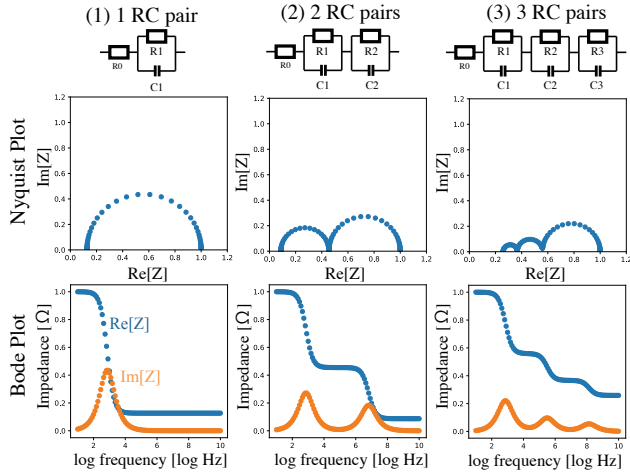


Fig. 1. Model selection from three RC pair models.

$$p(\mathbf{D}|M) = \int p(\mathbf{D}|\theta, M) d\theta, \quad (1)$$

where $p(\theta)$ is the prior distribution and $p(\mathbf{D}|\theta, M)$ is the likelihood. The mean evidence $\mathbb{E}[p(\mathbf{D}|M)]$ gives the probability of reproducing a given dataset \mathbf{D} with a given model M , the degree of model fit penalised by model complexity. The variance quantifies its uncertainty.

Surprisingly, Bayesian model selection of battery models has barely been reported, except for Miyazaki et al. (2020). Although Bayesian parameter estimation (Aitio et al., 2020; Escalante et al., 2021), and probabilistic modelling works (Huang et al., 2021; Liu et al., 2020) exist, most Bayesian approaches in the battery community use Markov chain Monte Carlo (MCMC) (Metropolis et al., 1953; Hastings, 1970), a user-friendly but sample-inefficient approach for inference. Recent work (Kuhn et al., 2022) on parameterisation applied a sample-efficient solver with Bayesian optimisation, none of the above solvers offer evidence computation. This is because estimating the evidence requires prohibitive integral computation, and this is particularly challenging when the likelihood is non-closed-form and/or expensive. A typical practice in such cases is to adopt the Bayesian information criterion (BIC), which is a coarse approximation of the evidence that assumes the posterior is a unimodal Gaussian. Unfortunately battery parameter estimation can produce multimodal or non-Gaussian posterior distributions (Aitio et al., 2020; Escalante et al., 2021), and ignoring this may cause overconfidence—previous work (Miyazaki et al., 2020) demonstrates that the identification of the best model using a variant of BIC gradually worsens as the posterior multimodality increases. This paper introduces Bayesian quadrature (BQ) as a novel technique for sample-efficient model evidence and parameter posterior estimation, and applies this to battery equivalent circuit models using synthetic data for demonstration purposes.

2. BATTERY MODEL FORMULATION

We selected ECMs for proof-of-concept here since they are relatively simple battery models that nonetheless offer identification challenges. Parameter identifiability for ECMs is often examined in the frequency domain, for example via electrochemical impedance spectroscopy (EIS)

data, although time domain data may also be used. Several plausible ECMs are usually compared when fitting EIS data, but the process is subjective and based on the user's electrochemical understanding of the target battery. For simplicity, we chose a simple resistance-capacitance (RC) pair model—this may represent various physical processes, for example kinetics and double layer capacitance, or an approximation of diffusion. Fig. 1 illustrates the circuit configurations and typical Nyquist plots of three variations of RC circuit models. The number of RC parallel connection components corresponds to the number of the semi-circles in a Nyquist plot. This correspondence is key for identifying the model from spectra. As the semi-circle shape implies, the real and imaginary parts of spectra have a mathematical relationship (Kramers-Kronig relations in Debye relaxation), where one part of spectra can be derived from the other via an equation.

We extend this formulation to make the model better suited for statistical inference using the hyperbolic formulation (Calderwood, 2003); we improve this, permitting non-dimensionalised parameterisation without positivity constraint, as follows. For a general circuit with N total RC pairs plus an additional series resistance R_0 , where R_i is the resistance of i -th RC pair $[\Omega]$, $\ln(\omega\tau_i)$ is the rescaled frequency scale to make the scale independent of the given frequency range of the dataset, rescaled with the breakpoint frequency $\omega_i := 1/\tau_i$ [rad/s], $\tau_i := R_i C_i$ is the time constant of the i -th RC pair [s], C_i is the capacitance of the i -th RC-pair [F], f is the frequency [Hz] and $\omega := 2\pi f$ is the angular frequency [rad/s], one can define the total resistance R_{total} , the log of this r_{total} (which is positive), the dimensionless resistance of i -th RC-pair r_i (constrained between zero and one), the unconstrained dimensionless resistance r'_i , the scaling factor R_{im} , and the weight of i -th hyperbolic secant distribution w_i , as follows:

$$R_{\text{re}} := R_0 + \sum_{i=1}^N R_i := R_{\text{total}} := \exp(r_{\text{total}}), \quad (2)$$

$$r_i := \frac{R_i}{R_{\text{total}}} := \exp[-\exp(r'_i)], \quad (3)$$

$$R_{\text{im}} := \frac{\pi}{2} \sum_{i=1}^N R_i \quad (4)$$

$$\lambda_i := \frac{R_i}{\sum_{i=1}^N R_i}. \quad (5)$$

From this, the real and imaginary parts of the impedance ($\text{Re}[Z]$, $\text{Im}[Z]$), are given by (see Appendix A)

$$\text{Re}[Z] = R_{\text{re}} \left[r_0 + \sum_{i=1}^N \frac{r_i}{2} [1 - \tanh(\ln \omega \tau_i)] \right], \quad (6)$$

$$\text{Im}[Z] = \underbrace{R_{\text{im}}}_{\text{scaling factor}} \underbrace{\left[\sum_{i=1}^N \frac{\lambda_i}{\pi} \text{sech}(\ln \omega \tau_i) \right]}_{\text{mixture of hyperbolic secant distributions}}. \quad (7)$$

The frequency range is also standardised according to the frequencies in the available dataset,

$$\mu_{\omega}, \sigma_{\omega} := \mathbb{E}[\ln \omega], \sqrt{\text{Var}[\ln \omega]}, \quad (8)$$

$$\omega^{\text{std}}, \tau_i^{\text{std}} := \frac{\ln \omega - \mu_{\omega}}{\sigma_{\omega}}, -\frac{\ln \tau_i + \mu_{\omega}}{\sigma_{\omega}}, \quad (9)$$

where the mean μ_ω and standard deviation σ_ω of logarithmic angular frequency¹, $\ln \omega$ can be calculated from the given frequency range of the dataset, and from this we define a standardised frequency scale ω^{std} and standardised time constants τ_i^{std} . These may be related to the actual time constants and capacitances (noting that τ_i is the unstandardised form of the time constant) via

$$\ln \omega \tau_i := \ln \omega - \sigma_\omega \tau_i^{\text{std}} - \mu_\omega, \quad (10)$$

$$C_i = \frac{\tau_i}{r_i R_{\text{total}}}. \quad (11)$$

The parameters to be fitted are unconstrained standardised ones $\theta := \{r_{\text{total}}, r'_i, \tau_i^{\text{std}}\}$. This formulation is similar to the distribution of relaxation times modelling. This canonical form provides three benefits: separation of scaling factor, unconstrained prior distribution selection for all parameters, and integral-friendly formulation. Separating the scaling factors can decompose parameter estimation problems into problems of estimating magnitudes (R_{re}) and ratios (r_i), permitting fair comparison over varied magnitudes of resistance. Logarithmically transformed parameters enable non-negativity constraints over resistance, allowing arbitrary prior distributions to be used for Bayesian inference (for instance, r_i is constrained between zero and one, but r'_i is unconstrained). The mixture of hyperbolic secant distributions offers several integral identities to analytically calculate the expectation and variance (see Appendix A). Moreover, this formulation interprets the imaginary part as a probability distribution function, allowing statistical analysis (see section 5).

3. BAYESIAN INFERENCE FORMULATION

We wish to select the likeliest model from the above-mentioned three RC pair options. In Bayesian inference, we need to assume a prior distribution $p(\Theta) := \pi(\Theta)$ and a likelihood function $p(\mathbf{D}|\Theta, M) := \ell_{\text{true}}(\Theta)$. The prior distribution is a probability distribution reflecting one's prior assumptions about possible parameters. For instance, we adopt here a multivariate normal distribution $\pi(\Theta) := \mathcal{N}(\Theta; \mu_\pi, \Sigma_\pi)$. The mean vector μ_π represents our guess of plausible parameter values and the covariance matrix Σ_π reflects our assumption on the uncertainty of each parameter, and correlations between parameters. The likelihood function $\ell_{\text{true}}(\Theta)$ is a probability distribution to evaluate how the selected parameter set Θ can reproduce the given dataset \mathbf{D} . Here we assume a univariate Gaussian with zero mean $\mathbf{0}$ and homoskedastic noise, meaning the noise variance σ_{noise} does not vary over frequency. The squared error evaluates how similar the observed data y_{obs} and ECM predicted data y_{ecm} are. Now, with the assumed prior $p(\Theta)$ and likelihood function $p(\mathbf{D}|\Theta, M)$, Bayes' rule defines the parameter posterior as $p(\Theta, M|\mathbf{D})$ and the model evidence $p(\mathbf{D}|M)$, all as follows:

$$p(\mathbf{D}|\Theta, M) := \ell_{\text{true}}(\Theta) := \prod_j^m \mathcal{N}(\text{err}_j(\theta); \mathbf{0}, \sigma_{\text{noise}}^2), \quad (12)$$

$$p(\mathbf{D}|M) := \mathcal{N}(\mathbb{E}_\pi[\ell_{\text{true}}(\Theta)], \text{Var}_\pi[\ell_{\text{true}}(\Theta)]), \quad (13)$$

$$p(\Theta|\mathbf{D}, M) = \frac{p(\mathbf{D}|\Theta, M)p(\Theta)}{p(\mathbf{D}|M)} = \frac{\ell_{\text{true}}(\Theta)\pi(\Theta)}{\mathbb{E}_\pi[\ell_{\text{true}}(\Theta)]}, \quad (14)$$

¹ Where necessary we assume arguments of logarithms are divided by appropriate units, e.g. 1 [rad/s], to ensure they are dimensionless.

where

$$\mathbf{D} := \{\mathbf{y}_{\text{obs}}, \omega^{\text{std}}\} \in \mathbb{R}^{m \times 2}, \quad (15)$$

$$\theta := \{r_{\text{total}}, r'_i, \tau_i^{\text{std}}\} \in \mathbb{R}^{d-1}, \quad (16)$$

$$\Theta := \{\theta, \sigma_{\text{noise}}^2\} \in \mathbb{R}^d, \quad (17)$$

$$y_{\text{ecm},j}(\theta) := \{y_{\text{re},j}, y_{\text{im},j}\} = M(\theta, \omega_j^{\text{std}}), \quad (18)$$

$$\text{err}_j(\theta) := [y_{\text{obs},j} - y_{\text{ecm},j}(\theta)]^2, \quad (19)$$

where subscript 'obs' refers to measured data, subscript 'ecm' to modelled data, and M is the model (equations (6)-(7)). The posterior $p(\Theta|\mathbf{D}, M)$ is a conditional probability distribution that reflects our updated estimate of the parameter space based on the observed data \mathbf{D} . We use dimensionless and unconstrained r'_i and τ_i^{std} as inputs of the model for arbitrary prior selection and fair comparison of models. The number of parameters to be estimated is $d = 2 + 2N$, as the scaling factor r_{total} and experimental noise variance σ_{noise}^2 are shared over all models.

4. BAYESIAN QUADRATURE MODELLING

We wish to estimate both the parameter posterior distribution $p(\Theta|\mathbf{D}, M)$ and the evidence $p(\mathbf{D}|M)$. We also wish to minimise the number of times that the likelihood $\ell_{\text{true}}(\Theta)$ must be queried, as this could be a computationally demanding operation in a more complex model. This problem requires a sample-efficient Bayesian inference solver. Bayesian quadrature (BQ) offers sample efficiency and solves for the posterior and the evidence in one go. This is a surrogate-model-based numerical integration approach, solving the integral as an inference problem by modelling the likelihood function $\ell_{\text{true}}(\Theta)$ with a Gaussian process (GP). Define $\ell(\Theta)$ as the surrogate likelihood function modelled by a GP. The key result is that BQ can recast the problem of Bayesian inference into one of function approximation. The more accurately $\ell(\Theta)$ can predict $\ell_{\text{true}}(\Theta)$, the more accurately the posterior and evidence can be estimated via replacing $\ell_{\text{true}}(\Theta)$ with $\ell(\Theta)$ in Eqs. (13) - (14). To achieve this, Adachi et al. (2022b) proposed *BASQ*, a discrete approximation of the kernel integral using a kernel recombination method (Hayakawa et al., 2022), yielding the following evidence computations:

$$\text{LEM} := \ln \mathbb{E}_\pi[\ell(\Theta)] \approx \ln \sum_k^L W_k \mu_f(X_k) + \beta, \quad (20)$$

$$\text{LEV} := \ln \text{Var}_\pi[\ell(\Theta)] \approx \ln \sum_{k,l}^L W_k W_l \sigma_f(X_k, X_l) + 2\beta, \quad (21)$$

where LEM and LEV refer to log evidence mean and log evidence variance, μ_f and σ_f are the predictive mean and covariance of the likelihood surrogate model $\ell(\Theta)$, β is the scaling constant, W_k, W_l and X_k, X_l are the positive weights and point configurations discretised by the kernel recombination. Recall that LEM gives the degree of model fit and the LEV quantifies the uncertainty of the fit. However, the prior work on this (Adachi et al., 2022b) assumed a narrower dynamic range of likelihood, whereas the battery model typically produces 10^{700} likelihood values. This is way beyond a typical numerical overflow limit. Thus, we improve here the prior work by adopting a four-layered warped GP method to accommodate the wide dynamic range of likelihood. (See Appendix A)

5. IDENTIFIABILITY

To evaluate the model evidence as a model selection criterion, we compare results against three classical metrics related to identifiability: number of data points m , signal-to-noise ratio (SNR), and Jensen-Shannon divergence (JS). Owing to the integral-friendly model formulation, most parts of these can be calculated analytically. The number of data points is controllable here because data are synthetically generated and equispaced over log angular frequency space. Both SNR and JS are calculated using the imaginary part of the impedance. As the canonical form can be regarded as a mixture of hyperbolic secant distributions, such statistical analysis can be applied. While SNR evaluates the identifiability along the impedance magnitude axis, JS does so along the frequency axis.

5.1 Signal-to-noise ratio

The SNR is the log fraction of the impedance variance over the noise variance, representing how much the signal is more distinct than the noise, defined as:

$$\text{SNR} := \ln \frac{\text{Var}_{P(\ln \omega)}[\text{Im}[Z]]}{\sigma_{\text{noise}}^2}. \quad (22)$$

Larger SNR means a more distinct and identifiable signal. The canonical form of the model provides an analytical form for the SNR (see derivation in Appendix A).

5.2 Jensen-Shannon divergence

The JS divergence is a distance metric quantifying how one probability distribution $P_i(x)$ is similar to a second reference probability distribution $P_j(x)$, defined as:

$$\begin{aligned} \text{JS} := & \frac{1}{2} \int \ln \left(\frac{P_i(x)}{M_{ij}(x)} \right) dP_i(x) \\ & + \frac{1}{2} \int \ln \left(\frac{P_j(x)}{M_{ij}(x)} \right) dP_j(x), \end{aligned} \quad (23)$$

where

$$M_{ij}(x) := \frac{1}{2} (P_i(x) + P_j(x)) \quad (24)$$

As the JS is defined for pairwise comparisons, the number of criteria required increases combinatorially per the number of RC pairs. For simplicity, we only consider the case of two RC pairs, which produces only one JS divergence. This represents how much the selected two peaks in the imaginary parts overlap. A smaller JS divergence means a more distinguishable and identifiable signal. While SNR is determined by the noise variance σ_{noise}^2 and scaling factor r_{total} , JS is dominated by the time constant difference $\Delta\tau_{ij}$. Again, the canonical form helps solve the integration. Note that this is formulated as noise-free. The integration calculation procedure can be seen in Appendix A. The extended JS to include noise σ_{noise}^2 is also guided, but the results shown in this paper are consistently used with noise-free formulation for simplicity.

6. NUMERICAL RESULTS

6.1 Selection criteria comparison

We now demonstrate our modified version of BASQ over several cases. We compare the model evidence metric

Table 1. Easy case

true model	1 RC pair	2 RC pairs	3 RC Pairs	4 RC pairs
		✓		
LEM	-2809233	703.6569	289.2976	225.1602
LEV	-33.52068	-27.31169	-31.91129	-31.38766
RMSE	1.147527	0.006677	0.031770	0.062151
BIC	5766999	-1405.553	-572.7390	-432.4597
ELPD	-2883641	713.7332	293.2492	213.7417

Table 2. Hard case

true model	1 RC pair	2 RC pairs	3 RC Pairs	4 RC pairs
			✓	
LEM	-150.3634	-151.8002	-147.4257	-151.9208
LEV	-15.74094	-19.07997	-19.45956	-26.57386
RMSE	0.492191	0.492643	0.492269	0.492260
BIC	302.8601	313.8921	321.0089	310.8289
ELPD	-145.7505	-148.8758	-148.4754	-146.1485

(LEM and LEV², eqs. (20) - (21)) with root-mean-square error (RMSE), Bayesian information criterion (BIC), and expected log predictive density (ELPD), based on the maximum a posteriori (MAP) parameter estimates, defined as:

$$\Theta_{\text{MAP}} := \text{argmax} \ell_{\text{true}}(\Theta), \quad (25)$$

$$\text{RMSE} := \sqrt{\frac{1}{m} \sum_j \text{err}_j(\theta_{\text{MAP}})}, \quad (26)$$

$$\text{BIC} := d \ln m - 2 \ln \ell_{\text{true}}(\Theta_{\text{MAP}}), \quad (27)$$

$$\text{ELPD} := \sum_j \ln \int \ell_{\text{true}}(\Theta) dP(\Theta | \mathbf{D}, M). \quad (28)$$

The RMSE is a noise-free formulation that does not consider parameter uncertainty. BIC is an asymptotic approximation of evidence, so it cannot evaluate multimodal likelihoods. ELPD is a similar formulation to the log mean evidence, but the probability measure is changed from prior to posterior. The motivation behind ELPD is to estimate the alternative evidence from MCMC samples, as it cannot estimate evidence when solving Bayesian inference. However, it relies on Monte Carlo (MC) integration, which requires a significant amount of posterior samples, meaning that a plethora of model evaluations $\ell_{\text{true}}(\Theta)$ will run. All these alternative criteria were calculated from the BQ estimated posteriors by post-processing. Moreover, none of these criteria quantify their own uncertainty except BQ.

We demonstrate the behaviours of the selection criteria on two different datasets—an easy case ($\Delta\tau_{ij} = 9.1$, $\ln \sigma_{\text{noise}}^2 = -9.97$) with results in Table 1, and a hard case ($\Delta\tau_{ij} = 0.36$, $\ln \sigma_{\text{noise}}^2 = -1.6$) detailed in Table 2. The easy case is clean data generated with 2 well-separated semi-circles, and the hard case is noisy data generated with an additional third semi-circle with more overlap. The "better" column shows which upward or downward direction is better for each criterion. As expected, separated peaks (large $\Delta\tau_{ij}$) and lower noise σ_{noise}^2 boost identifiability. While all criteria selected the true model in the easy case, only the evidence can select the true model in the hard case.

² LEV values in the tables are standardised via subtracting 2β from eq. (21) for a fair comparison between models.

Table 3. Linear correlation matrix

factors	m	JS	SNR	LEM	LEV
m	-	-0.0268	0.0058	0.4096	-0.1806
JS	-0.0268	-	-0.0993	-0.0297	0.2243
SNR	0.0058	-0.0993	-	0.7299	-0.4882
LEM	0.4096	-0.0297	0.7299	-	-0.2867
LEV	-0.1806	0.2243	-0.4882	-0.2867	-

The other metrics were unsuccessful in the hard case because of a multimodal posterior in the one RC pair model. As three RC pairs were used to generate the dataset, the posterior distribution of one RC pair parameter inevitably becomes multimodal, such as the peak intensity (λ_i). While the evidence correctly incorporates the multimodal distribution shape, RMSE and BIC consider only the largest peak. The BIC estimates the whole distribution from the local curvature at the maximum, which becomes erroneously overconfident in the multimodal case (Murphy, 2012). ELPD's failure could be due to its rough integral approximation. As the convergence rate of MC integration is $\mathcal{O}(1/\sqrt{n})$, the posterior samples ($n = 1,000$) is too few. This means more model evaluations $\ell_{\text{true}}(\Theta)$ are required, which would not scale to slower simulation models.

In contrast, the evidence can be estimated simultaneously during training. Moreover, the variance of the evidence successfully points out the lower confidence in the one RC pair model in the hard case, suggesting multimodality. This uncertainty over the selection criterion could avoid overconfidence toward a simpler model. Moreover, the evidence variance in the hard case is generally higher than in the easy case. This also tells us that the hard case dataset is almost unidentifiable, suggesting we should not trust these comparisons. For instance, the evidence mean and ELPD for one RC pair in the easy case are much lower than in the hard case. However, the integral variance is the opposite. Thus, only this metric quantifies its own uncertainty, suggesting the dataset or model is less informative. A similar notion can be found in Jeffreys' scale for the Bayes factor (Jeffreys, 1998), which claims the evidence is not strong when the difference between the log evidence of two models is lower than 10. This explains that the hard case is unreliable, as the difference in the log evidence shows insufficient plausibility. Contrary to Jeffreys' scale, log evidence variance is self-contained and does not require the comparison of models. Instead, it can independently spot the unreliability of the estimation.

In such an uncertain case, a typical practice is Bayesian model averaging. Rather than selecting one definite model, we sample from a *mixture of models* with probability proportional to their mean evidence. Averaging can boost predictive accuracy and reduce the uncertainty over predictions, where only evidence offers this method. As such, while the easy cases do not require advanced methods, the evidence with self-check on reliability can assist in deciphering minor differences in hardly identifiable problems.

6.2 Sensitivity analysis

A sensitivity analysis of the evidence metric was performed. We generated 1,024 datasets using two-RC-pair models while varying the following five parameters; the

Table 4. Functional ANOVA results

factors	LEM	LEV	residual
(m)	0.0012	0.0083	0.0012
(JS)	0.3258	0.3335	0.3397
(SNR)	0.2778	0.2002	0.0825
(m , JS)	0.0028	0.0088	0.0017
(m , SNR)	0.0599	0.1260	0.2427
(JS, SNR)	0.2702	0.1969	0.0879
(m , JS, SNR)	0.0623	0.1264	0.2443

number of data points m , the scaling factor r_{total} , the first resistance r'_1 , the first time constant τ_1^{std} , and noise variance σ_{noise}^2 . We calculated the SNR, JS, and the number of data points m for each dataset. In the first step of the analysis, we compared the linear correlations between the evidence estimates. Table 3 shows Pearson's correlation coefficients. This result aligns with our intuition—for instance, larger data size m and SNR can boost the evidence LEM and confidence (inverse of LEV). However, while the large correlation of evidence with SNR is instinctive, the small correlation with JS is counterintuitive.

Thus, we further investigated the variance analysis via functional ANOVA (Hutter et al., 2014), which models a partition of a functional response according to the main effects and interactions of input parameters. This method can attribute each parameter sensitivity in a non-linear manner. Table 4 illustrates that the most significant influence over the mean and variance of the evidence is the JS, contrary to the linear correlation results. This can be interpreted as meaning that a smaller JS divergence (more overlapped peaks) destabilises the evidence estimation, resulting in a more considerable variance. This viewpoint is supported by the relatively large negative correlation coefficient between JS and LEV.

Further insights can be obtained via residual analysis. The residual is defined as follows:

$$Z_{\text{pred}} := \text{slope} \times \text{BIC} + \text{intercept}, \quad (29)$$

$$\text{residual} := (Z_{\text{pred}} - \log \mathbb{E}_{\pi}[\mu_e(\Theta)])^2. \quad (30)$$

As the BIC is an approximation of the LEM, the BIC and LEV have a linear relationship. While a linear regression model with BIC can predict log evidence mean reasonably, it fails to predict in hard cases, as shown in the section 6.1. Residual refers to the squared error between the BIC and log evidence mean. Table 4 shows that the residual is mainly caused by the JS divergence and less influenced by SNR or the number of data points m . This also suggests that the BIC cannot distinguish between the models with overlapped peaks, namely, a multimodal posterior.

6.3 Computation efficiency

Lastly, we compared the computation efficiency of our modified version of BASQ with the existing MCMC solvers elliptical slice sampling (ESS) (Murray et al., 2010) and dynamic nested sampling (Speagle, 2020). Note that amongst MCMC samplers, only nested sampling can estimate the evidence. For ESS, we approximated the evidence using ELPD via posterior samples. Therefore, the estimation with ESS should converge to a larger value than the




Fig. 2. The learning curve of log evidence over the computation time and the number of samples.

actual evidence. The BASQ computation was performed using both CPU and GPU.³

Fig. 2 compares the learning curve of the above four samplers versus computation time, using the easy case dataset shown in Table 1. While BASQ in a GPU converges at 18 seconds, BASQ in a CPU converges at 131 seconds. Both ESS and nested sampling do not converge in this time. Fig. 2 contrasts the sample efficiency of the samplers. As BASQ is a parallel sampler, we generate 100 samples per iteration. The sampling efficiency of BASQ does not change over computation modes and is the best of the selected solvers. This is expected—while the convergence rate of BASQ is $\mathcal{O}(\exp(-cn^{1/d}))$ in the Gaussian case (Adachi et al., 2022b), that of MCMC is $\mathcal{O}(1/\sqrt{n})$. Furthermore, even this result does not fully represent BASQ’s potential. While ECMs return model predictions in a millisecond order, more complex models (e.g. DFN model) take seconds to query. Therefore, BASQ for such complex models will be even more beneficial. Recent work shows even faster convergence than BASQ (Adachi et al., 2023).

REFERENCES

- Adachi, M. et al. (2022a). Bayesian model selection of lithium-ion battery models via Bayesian quadrature. URL <https://arxiv.org/abs/2210.17299>.
 Adachi, M. et al. (2022b). Fast Bayesian inference with batch Bayesian quadrature via kernel recombination. *NeurIPS*, 35.
 Adachi, M. et al. (2023). SOBER: Scalable batch Bayesian optimization and quadrature using recombination constraints. *arXiv preprint arXiv:2301.11832*.
 Aitio, A. et al. (2020). Bayesian parameter estimation applied to the Li-ion battery single particle model with electrolyte dynamics. *IFAC*, 53(2), 12497.

- Bizeray, A.M. et al. (2018). Identifiability and parameter estimation of the single particle lithium-ion battery model. *IEEE Trans. Control. Syst. Technol.*, 27, 1862.
 Calderwood, J. (2003). A physical hypothesis for Cole-Davidson behavior. *IEEE Trans. Dielectr. Electr. Insul.*, 10(6), 1006.
 Doyle, M. et al. (1993). Modeling of galvanostatic charge and discharge of the lithium/polymer/insertion cell. *J. Electrochem. Soc.*, 140(6), 1526.
 Escalante, J.M. et al. (2021). On uncertainty quantification in the parametrization of Newman-type models of lithium-ion batteries. *J. Electrochem. Soc.*, 168, 110519.
 Hastings, W.K. (1970). Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, 57(1), 97.
 Hayakawa, S. et al. (2022). Positively weighted kernel quadrature via subsampling. *NeurIPS*, 35.
 He, H. et al. (2011). Evaluation of lithium-ion battery equivalent circuit models for state of charge estimation by an experimental approach. *Energies*, 4(4), 582.
 Huang, J. et al. (2021). Towards robust autonomous impedance spectroscopy analysis: A calibrated hierarchical Bayesian approach for electrochemical impedance spectroscopy (EIS) inversion. *Electrochim. Acta*, 367, 137493.
 Hutter, F. et al. (2014). An efficient approach for assessing hyperparameter importance. In *ICML*, 754. PMLR.
 Jeffreys, H. (1998). *The theory of probability*. OUP Oxford.
 Kemper, P. et al. (2013). Extended single particle model of Li-ion batteries towards high current applications. In *IEEE VPPC*, 1.
 Kuhn, Y. et al. (2022). Bayesian parameterization of continuum battery models from featurized electrochemical measurements considering noise. *Batteries & Supercaps*.
 Liu, J. et al. (2020). The Gaussian process distribution of relaxation times: A machine learning tool for the analysis and prediction of electrochemical impedance spectroscopy data. *Electrochim. Acta*, 331, 135316.
 Metropolis, N. et al. (1953). Equation of state calculations by fast computing machines. *Chem. Phys.*, 21(6), 1087.
 Milocco, R.H. et al. (2014). Generic dynamic model of rechargeable batteries. *J. Power Sources*, 246, 609.
 Miyazaki, Y. et al. (2020). Bayesian statistics-based analysis of ac impedance spectra. *AIP Adv.*, 10, 045231.
 Murphy, K.P. (2012). *Machine learning: a probabilistic perspective*. MIT press.
 Murray, I. et al. (2010). Elliptical slice sampling. *AISTATS*, 541.
 Rasmussen, C. et al. (2000). Occam’s razor. *NeurIPS*, 13.
 Santhanagopalan, S. et al. (2006). Review of models for predicting the cycling performance of lithium ion batteries. *J. Power Sources*, 156(2), 620.
 Speagle, J.S. (2020). dynesty: a dynamic nested sampling package for estimating Bayesian posteriors and evidences. *MNRAS*, 493(3), 3132.

Appendix A. MORE INFORMATION

The appendix can be found in Adachi et al. (2022a).

³ Both MCMC samplers and BASQ in CPU were computed with a MacBook Pro 2019, 2.4 GHz 8-Core Intel Core i9, 64 GB 2667 MHz DDR4. BASQ in GPU was performed on Google Colaboratory.