Article

An active learning framework for the rapid assessment of galvanic corrosion

Check for updates

Aditya Venkatraman [1], Ryan Michael Katona [1], Demitri Maestas[1], Matthew Roop[1,2], Philip Noell [1] & David Montes de Oca Zapiain [1] ✉

The current present in a galvanic couple can define its resistance or susceptibility to corrosion. However, as the current is dependent upon environmental, material, and geometrical parameters it is experimentally costly to measure. To reduce these costs, Finite Element (FE) simulations can be used to assess the cathodic current but also require experimental inputs to define boundary conditions. Due to these challenges, it is crucial to accelerate predictions and accurately predict the current output for different environments and geometries representative of in-service conditions. Machine learned surrogate models provides a means to accelerate corrosion predictions. However, a one-time cost is incurred in procuring the simulation and experimental dataset necessary to calibrate the surrogate model. Therefore, an active learning protocol is developed through calibration of a low-cost surrogate model for the cathodic current of an exemplar galvanic couple (AA7075-SS304) as a function of environmental and geometric parameters. The surrogate model is calibrated on a dataset of FE simulations, and calculates an acquisition function that identifies specific additional inputs with the maximum potential to improve the current predictions. This is accomplished through a staggered workflow that not only improves and refines prediction, but identifies the points at which the most information is gained, thus enabling expansion to a larger parameter space. The protocols developed and demonstrated in this work provide a powerful tool for screening various forms of corrosion under in-service conditions.

Materials degradation is a pervasive challenge in engineering that manifests in diverse forms across a wide variety of applications and are often highly dimensional containing variations in part dimensions, service environments, or material properties. Digitial twins or other models, including Finite Element Methods (FEM), are frequently utilized to inform upon materials degradation, helping to determine the rate and extent of degradation. Despite the utility of models, boundary conditions (BC) are often necessary, in combination with the geometry, to describe a material's performance in a given environment. This becomes especially difficult if there is a wide range of geometric and environmental considerations for a material's performance. Further, if BC's are non-trivial to measure, substantial time could be spent on boundary condition acquisition. Therefore, there exists a need to speed up materials degradation in complex, high dimensional cases.

Active Learning (AL) is a viable solution to maximize the efficacy of materials degradation models that utilize FEM simulations and experimentally measured BCs. This approach provides a standardized, probabilistically informed pathway to identify optimal input configurations in an efficient manner[1-3]. AL establishes an accurate surrogate model using an iterative process that starts with a sparse dataset and sequentially selects additional datapoints that exhibit the highest potential to improve the performance of the model. The additional datapoints are selected by maximizing an acquisition function[3,4]. This acquisition function is informed using a surrogate model[5,6] that is trained on the sparse initial dataset[7-9]. Prior studies have used various acquisition functions including predictive uncertainty, expected improvement[10,11], and expected information gain[8,12,13]. While the utility of AL has been shown, limited prior work has incorporated AL in a stratified manner to address problems with environment and geometric variations.

One high dimensional problem that FEM can be utilized for is corrosion. Specifically, galvanic corrosion is a phenomenon in which dissimilar metals, such as those in fasteners, come into contact under a corrosive environment[14,15]. These structures are exposed to a wide range of corrosive environments, including atmospheric conditions in which the relative humdity in combination with salt load will dictate the concentration and thickness of the brine layer. When two dissimilar metals are in contact under

¹Sandia National Laboratories, Albuquerque, New Mexico 87123, USA. ²University of New Mexico, Albuquerque, New Mexico 87131, USA.
✉e-mail: dmonte@sandia.gov

1

a corrosive brine, the corrosion rates experienced by the metals are governed by various environmental, chemical, and geometric factors[16–19]. In addition, the composition of the alloys that constitute the galvanic couple also plays an important role in controlling the reaction rates at the surface[14,20]. Such a wide range of environments as well as reaction mechanisms present in corrosion provides a complex, stratified dataset which can be difficult to understand without advanced analytical techniques such as machine learning (ML) and AL[21].

A metric that provides useful insight into the susceptibility to galvanic corrosion is the current (either anodic or cathodic) for a given environment, material, and geometric configuration as the current can often inform upon materials loss. It is important to note that under the consideration of charge conservation, the anodic and cathodic currents will be equal and opposite. Computational modeling of galvanic couples has served as a convenient means of estimating the available cathodic current across multiple conditions[16,17,22–25]. For corrosion FEM modeling, electrochemical boundary conditions and solution properties are necessary at every relative humidity (i.e., chloride concentration) and often requires performing experiments through the measurement of polarization curves. Performing experimental investigations across multiple configurations of environmental parameters and electrode materials contains a specific cost, including: (i) materials, (ii) specialty equipment (i.e., temperature control), and (iii) time of material preparation (i.e., polishing) and experiments. Therefore, it is not feasible to measure and model all potential combinations of environments and alloy configurations.

Previous studies have modeled metallic corrosion under environmental conditions with varying levels of fidelity using the Finite Element (FE) Method[20]. These FE models consider differ in terms of time-dependency (i.e., including transient charge transport[24,26]), the dimensionality of the specimen (considering either one[27,28] or two-dimensional specimens[22,24,26]), and the coupling between electrodynamics and the charge transport[24,26]. Recently, Liu and Kelly[17] have demonstrated that the Laplace equation can be used to model galvanic corrosion under steady-state, electroneutral conditions. This model form has been utilized in prior works, along with experimentally derived cathodic boundary conditions to predict the current output of a SS-AA7050 galvanic couple[16,17]. Despite successful application of FE-simulations in literature, significant challenges remain for the rapid exploration of the vast space of inputs that influence corrosion. The main limiting factor is the need to reconcile the boundary conditions (BC) in FE simulations with experimental results. Some previous studies[16,17,23] utilized cathodic polarizations obtained from a rotating disk electrode to capture the coupled, nonlinear BC for the solution of the Laplace equation in atmospheric conditions. These BCs are developed based on the Butler-Volmer equation[15], and are dependent on the temperature and salt concentration within the electrolyte. Thus, a minimum of two experimental polarization scans (i.e., anodic and cathodic) is necessary for every environmental configuration, with oftentimes replicate scans conducted to ensure fidelity of the data. Even after experimental data is obtained, several FE simulations are required to identify interactions between the boundary conditions and geometric configuration[16]. This leads to significant run time for the FE method on distributed computing clusters depending upon the geometric complexity of the specimen and the fidelity of the physics-based model.

As experimental cost to measure BCs and solve the FEM models can be excessive, a protocol to accelerate the predictions of current in a corrosion models is desirable. These accelerated predictions are typically obtained with the help of a machine learned surrogate model[4–6,25] trained on a dataset of FEM simulations. The surrogate model emulates the response of performing an experiment/simulation at a specific input configuration for which the data is currently unavailable in a fast and accurate manner. Depending upon the run time for the FE simulation and the form of the surrogate model, it is possible to accelerate the predictions by multiple orders of magnitude[1,5,7,12,29]. However, training the surrogate model with a set of previously performed simulations (and experiments) is still necessary. It is critical to perform the calibration experiments at optimally selected

input configurations, so that the accuracy of the surrogate model is maximized with the fewest number of data acquisition steps.

For instance, Bansal et al.[30] used the cumulative distribution function of the classification confidence value to iteratively improve the GP surrogate model predictions for the material loss experienced by metallic joints in spot welds, and obtained a drastic (75%) reduction in prediction errors after 30 AL iterations. These AL approaches can be utilized to incrementally increase the explored ranges of the input parameters. In other words, the acquisition function predictions can be estimated beyond the input domain that was originally used to calibrate the surrogate model, which in turn allows for an enlarged domain in a cost-efficient manner. However, a direct application of AL approaches to accelerate galvanic corrosion predictions faces several hurdles.

First, the existing protocols for AL search the entire input domain using the acquisition function to sequentially identify additional inputs for data acquisition purposes. However, the available cathodic current of a galvanic couple is dependent on environmental (temperature, WL thickness), chemical (salt concentration), and geometric parameters (cathode length). Consequently, the experimental polarization curves need to be discerned for each new combination of temperature and salt concentration, since they inform the BCs used in the simulations[16,17,23]. With these BCs, the FE simulations are used to calculate the cathodic current. In other words, the simulations can probe the dependency of the current on the cathode length and WL thickness, for a given combination of temperature and salt concentration. For that reason, an AL framework must follow for a "staggered" prediction of the acquisition function. The framework needs to identify subsets of inputs (temperature and salt concentration) to experimentally measure polarization curves and, subsequently, identify cathode lengths and WL thicknesses that have the highest potential to improve the predictions. Existing protocols are inadequate to function in this manner as searches for additional inputs occur over the entire input domain. This necessitates the adoption of a versatile surrogate modeling approach capable of estimating the acquisition function with limited information about the input configuration. Second, the optimization of the acquisition function typically only results in the identification of a single input configuration[7]. However, acquiring the data for a single input configuration may not always represent the best trade-off between exploration and exploitation[31]. Therefore, a new framework capable of leveraging the acquisition function to identify a set or batch of input configurations that collectively refine the predictions to the highest possible extent is necessary. Third, while the surrogate model can be utilized to extrapolate the output predictions beyond the input domain, the accuracy of these predictions are likely to be poor[5]. Therefore, the optimal choice for the acquisition function needs to be identified to aid in the enlargement of the input domain. This is crucial to ensure that we obtain the fastest improvements to the model performance with additional data. Lastly, it is unlikely that a surrogate model-form will be identified that performs optimally across the entire input-domain due to the complex, multi-physics nature of corrosion phenomena. Therefore, responses from multiple surrogate model forms will be optimally combined to improve the robustness of the FE predictions.

This work seeks to address the aforementioned challenges to optimally establish a highly accurate surrogate model for the current in a corrosion system. Specifically, this goal will be achieved by developing an AL framework that enables the identification of input configurations for a galvanic couple with the highest potential to increase the performance of the surrogate model. The framework developed is initiated through training a Gaussian Process (GP)[5,30,32] and Neural Network (NN)[6] surrogate model for the current output using a simulation dataset that was recently made available by Katona et al.[16]. Information fusion approaches[33,34] are used to combine the NN and GP predictions in an optimal manner. A 'staggered' GP model is then constructed by analytically marginalizing the influence of the WL layer thickness and cathode length. This 'staggered' surrogate model, which only depends on the temperature and salt concentration, is then used to maximize the acquisition function. In addition, two acquisition functions are considered for active learning - the predictive standard
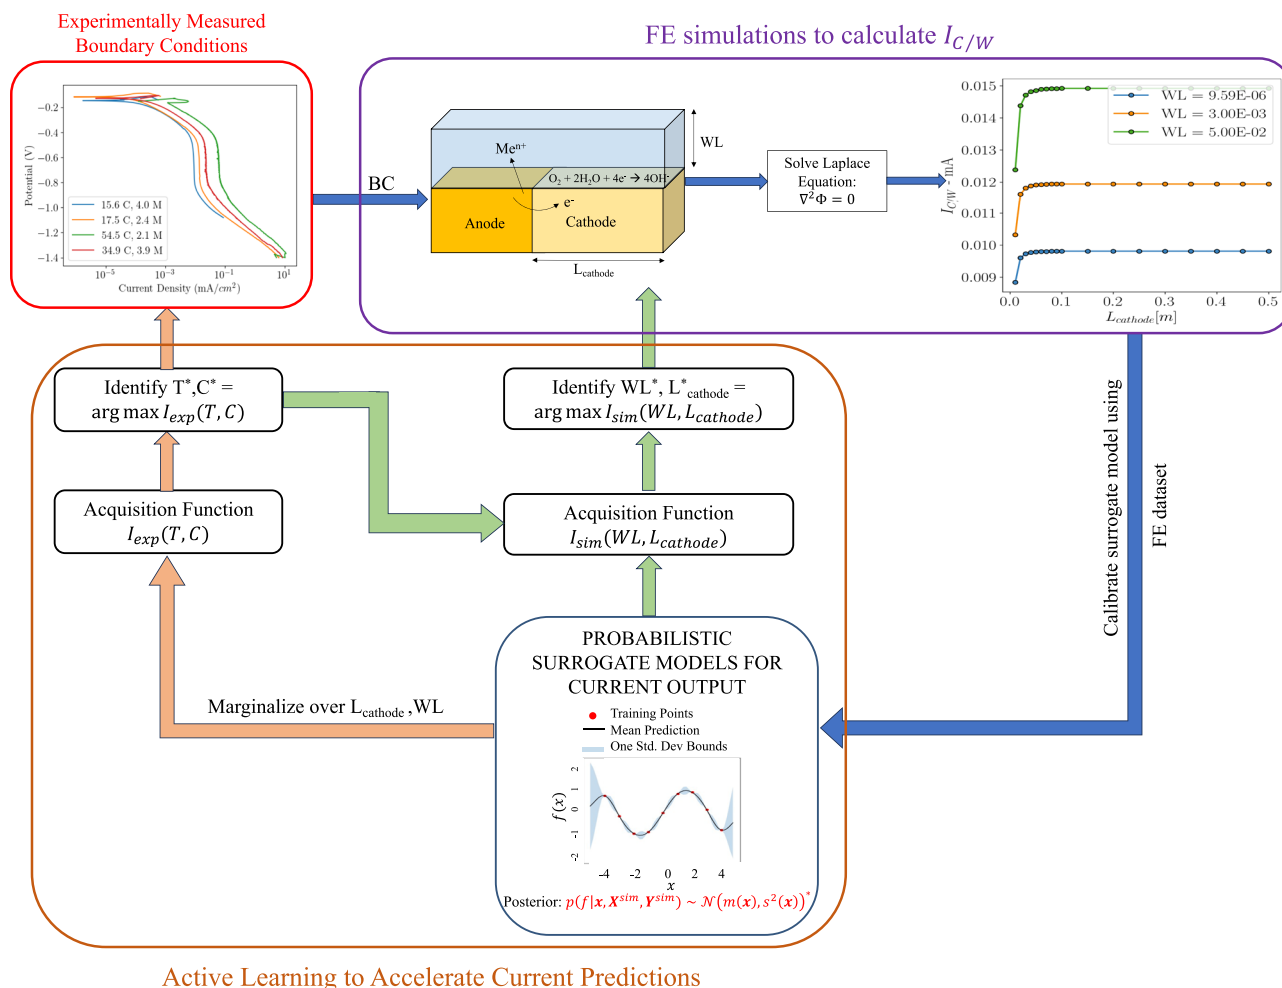
Fig. 1 | Illustration of the active learning protocol utilized in this study. The blue arrows indicate the exploration step where additional data is acquired with FE simulations. The green and orange arrows represent the exploitation step where the acquisition function is used to identify optimal input configurations for further exploration.

deviation and the expected information gain. The maximization of the chosen acquisition function identifies the best configurations of inputs for performing additional experiments. Once the experimental polarization curves are gathered, the acquisition function is used to sequentially identify additional inputs for the simulations. Weighted K-means clustering[6] is utilized in both steps to identify multiple sets of input configurations to drive a balance between exploring the input domain and exploiting the GP surrogate model. The efficacy of this protocol is demonstrated by substantially refining the predictions for the current output of an SS304-AA7050 galvanic couple. As a result, this work will demonstrate a significant (almost 50%) increase in model performance after incorporating the optimal input configurations. In addition, the optimal choice for the acquisition function is identified based on the improvements of the current predictions. The workflow utilized in this study is schematically illustrated in Fig. 1. While the scope of the current study is restricted to galvanic corrosion, the implications of our framework extend far beyond this specific scenario, offering a versatile tool for understanding and mitigating materials degradation across a myriad of environments and applications.

Results
Active learning protocol reduces training costs by 50%
First, the validity and efficacy of the AL framework developed in this work to leverage a previously generated dataset is demonstrated. This dataset consists of 2520 FE simulations of galvanic corrosion and can be found in Montes de Oca Zapiain et al.[30]. The dataset from Montes de Oca Zapiain

et al. calibrated the FE models to the work in Katona et al.[16] and uses a Laplace equation-based FE model to assess the cathodic current of a galvanic couple over a wide range of geometrical and environmental parameters and evaluated the current per width ($I_{C/W}$). The parameter ranges explored using the FE simulations are shown in Table 1 which were sampled using a hypercube sampling strategy. This hypercube specified by Table 1 is denoted as $\mathcal{C}_1$.

The AL framework developed in this work was used to establish an accurate surrogate model for predicting $I_{C/W}$ within the input domain delineated in Table 1 in an accurate and computationally efficient manner using an optimal number of simulation results (see Methods section).

Randomly splitting the FE dataset into training and testing datasets consisting of 2000 and 520 datapoints respectively initializes the AL framework. Once a test/train split was performed, an "initial" batch of training

Table 1 | Parameter ranges for the FE simulations used to calculate the cathodic current output of a galvanic couple

Parameter	Range
Temperature	[25,45] °C
Salt concentration	{1,3,5} M
WL thickness	[7E-06,0.05] m
Cathode Length	[0.01,0.5] m

Fig. 2 | Illustration of the progression of the 10-fold CV errors with the addition of input datapoints to the training dataset from the set of candidates, using the predictive standard deviation and expected information gain acquisition functions. The predictive standard deviation outperforms the expected information gain since the saturation value is reached with fewer candidates.
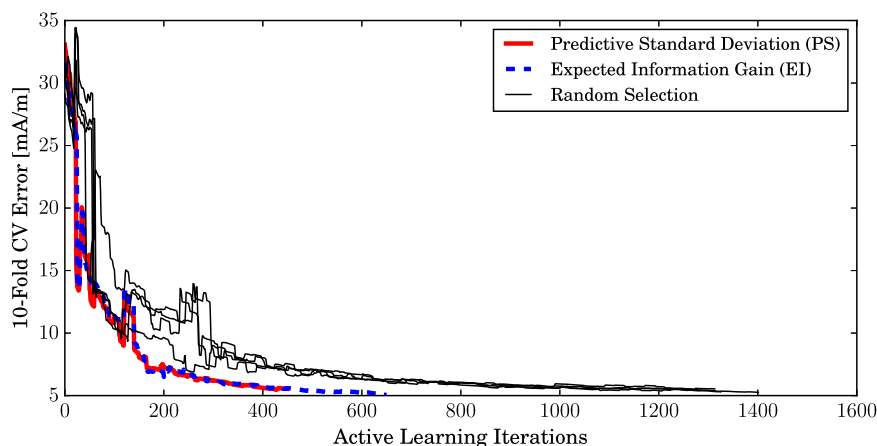

Fig. 3 | Schematic description of model performance. a Parity plots showing the performance of the fused surrogate models. The training datasets were designed using the predictive standard deviation (PS) and expected information gain (EI) respectively. The 10-fold cross validation predictions are shown for the training dataset, whereas the fused model predictions are shown for the test dataset. **b** Cumulative probability of absolute error for the K-fold cross validation predictions and the predictions of the test set using the predictive standard deviation (PS) and expected information gain (EI) respectively.
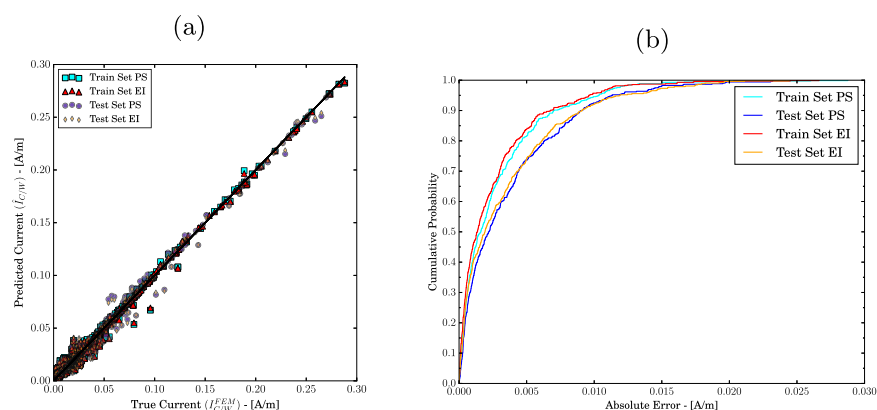
data was chosen by selecting 50 points from the training dataset using the MaxPro approach[35,36]. This approach maximizes the uniformity of the projections of the design space onto any lower dimensional manifold. In other words, the MaxPro approach enables the selection of a batch of 50 datapoints (out of a total of 2000) that best resembles a random uniform design in any of 14 lower dimensional hypercubes ($^4C_1 + {}^4C_2 + {}^4C_3$) of the four dimensional parameter space. This represents the best choice for a space filling design in the absence of supervision (*i.e.*, $I_{C/W}$ predictions from the FE simulations)[35].

Subsequently, the initial batch of training data was used to calibrate an initial GP surrogate model while the remaining datapoints in the training dataset (1950) are treated as candidates. The initial GP surrogate model is used to calculate the acquisition function on the candidates. The candidate datapoint that exhibits the highest value of the acquisition function is then concatenated to the training dataset, and the GP model is recalibrated. This is an iterative process where the datapoints at which the model can be improved the most (determined by the value of the acquisition function) are sequentially added to the training set. This process is repeated until a set number of iterations or a determined threshold is met.

The improvements to the current predictions obtained using this active learning strategy are assessed with the help of the K-fold cross validation (K-fold CV) error, with $K = 10$. To calculate the K-fold CV error, the training dataset is first split into K subsets. The surrogate model is calibrated on the all but one of these subsets, and then used to obtain predictions for the subset that was not used for calibration, also called the 'out-of-fold' subset. This process is repeated for all subsets to generate the K-fold CV predictions and errors. It is noted that the K-fold cross-validation predictions offer a more effective means of evaluating the performance of the GP surrogate model on the training dataset. This is attributed to the exact interpolation of

the training data effected by the GP model, leading to negligible training errors when used for direct prediction.

The progression of the K-fold cross validation error for both acquisition functions is shown in Fig. 2. It is clear that both acquisition functions cause the K-fold cross validation error to saturate to the same value of 5 $mA\ m^{-1}$. Furthermore, it is worth noting that the saturation value obtained with the two different acquisition functions is similar to the K-fold cross validation error obtained using the entire training dataset. In other words, a GP model trained with the 2000 points exhibits a K-fold cross validation error of 5 $mA\ m^{-1}$ showing that the systematic addition of high value observations results in an accurate model.

The predictive standard deviation (PS) acquisition function is more effective for selecting suitable datapoints from the candidate dataset given that the PS required 450 points to saturate the K-fold cross validation. In contrast, the expected information gain (EI) needed 650 points. Despite this contrast, similar values of the cross-validation errors as compared to the GP model trained on the entire dataset were achieved, while utilizing fewer training samples, regardless of the choice of the acquisition function. More precisely, a percentage reduction of at least 50% in computational costs using these acquisition strategies was achieved.

To demonstrate the effectiveness of the AL method, we conducted a series of trials without utilizing an acquisition function. Instead, inputs were randomly selected from the candidate dataset. Five different instances were generated by adjusting the random seed. The surrogate models trained using this random selection approach required over twice as many datapoints, on average, to reach the same level of accuracy as the actively learned surrogate model.

Table 2 | Comparison of the improvements obtained to the predictions for the initial dataset using the acquisition functions reveals that the expected information gain (EI) performs better than predicted standard deviation (PS)

Acquisition Function	Training MAE (10-fold CV)	Testing MAE
Predictive Standard Deviation	$5\,mA\,m^{-1}$	$3.7\,mA\,m^{-1}$
Expected Information Gain	$5\,mA\,m^{-1}$	$3.2\,mA\,m^{-1}$

Surrogate models accurately predict cathodic current for unknown input configurations

The AL framework enabled the development of a surrogate model for predicting $I_{C/W}$ with a significantly smaller fraction of available data. In this section the accuracy of the predictions obtained with the trained model are evaluated and assessed. We note that the optimal datasets identified in the previous step were used to re-train the GP and the NN surrogate models, after which the responses from the respective models were fused. The performance of the fused model is first demonstrated on the training and test sets. Figure 3 schematically presents the accuracy of the developed model. Specifically, Fig. 3a shows the scatter plot of the predictions for $I_{C/W}$ obtained using the fused model (denoted as $\hat{I}_{C/W}$) and $I_{C/W}$ predictions obtained from the FE simulations (shown as $I_{C/W}^{FEM}$) for both acquisition functions considered in this study. The predictions are compared for the test set (*i.e.*, the remaining 520 points on which the model was not trained) as well as the training set. Note that the K-fold CV predictions are used for the training dataset. Fig. 3b shows the cumulative probability of absolute error for the K-fold cross validation predictions and the predictions of the test set. From Fig. 3b it is evident that the error performance of both acquisition functions is similar. Finally, the mean absolute errors of the test set predictions are $3.7\,mA\,m^{-1}$ and $3.2\,mA\,m^{-1}$ for the PS and the EI acquisition functions respectively. Therefore, the error performance, in conjunction with the plots shown in Fig. 3, demonstrate that the model is capable of accurately predicting $I_{C/W}$ for scenarios not captured within the training data. The model performance is summarized in Table 2.

The model was further evaluated through comparison of the predictions to a completely new set of FE results, whose input parameters were obtained from Liu and Kelly[17] and Katona et al.[16]. This new dataset is referred to as the "extrapolation dataset", and it consists of 2885 datapoints. The temperatures and salt concentrations explored in the extrapolation dataset were (i) {25°C, 0.6 mol}., (ii) {35°C, 1mol}, (iii) {35°C, 3mol}, and (iv) {35°C, 5.3 mol}. Notice that the temperatures values lie within the input domain delineated in Table 1, however, the maximum values of the cathode length and WL thickness in the extrapolation dataset are greater than the corresponding maximum values shown in Table 1. Fig. 4 graphically depicts that the tails of the distribution of the cathode lengths and the WL

thicknesses (see Fig. 4b and 4b respectively) lie outside $\mathcal{C}_1$. A similar note is made for the minimum values of salt concentration given that it lies outside of the domain shown in Table 1.

Consequently, the fused model predictions are expected to be less accurate for the extrapolation dataset. Figure 5 graphically depicts the performance of the trained model on the new extrapolation dataset. Specifically, Fig. 5a shows the scatter plot of the predictions for $I_{C/W}$ obtained using the fused model (denoted as $\hat{I}_{C/W}$) and $I_{C/W}$ predictions obtained from the FE simulations (shown as $I_{C/W}^{FEM}$) for the extrapolation set for the model trained with the EI acquisition function and the PS acquisition function. It is evident that model performance significantly decreased. A conclusion that is reinforced with Fig. 5b, which shows the cumulative probability of absolute error for the predictions on the K-fold cross validation predictions, the predictions of the test set, and the predictions of the extrapolation set. Improvements to the model are necessary to enhance the accuracy of the out-of-distribution predictions.

Active learning framework drastically improves out-of-distribution current predictions

Previously shown results highlight that the developed model's performance decreases when tasked to predict the results for parameters outside of the support of the training data. The performance of the surrogate model can be improved by incorporating more data outside the original domain. The input domain is a 4-dimensional space, which makes it non-trivial to select points to be added. Furthermore, it is important to acknowledge that an experimental procedure will need to be performed to solve FE simulations to obtain accurate results for new temperatures and molar salt concentrations. Therefore, there is a critical need for a protocol capable of identifying the points that will have the highest impact in the performance of the surrogate model without performing an extensive search. Fortunately, the AL protocol developed in this work (and validated previously) enables this. The following section demonstrates the efficacy of the developed AL protocol to improve the performance of the model by incorporating a set of points that will yield the highest improvement. The AL workflow utilized to improve the out-of-distribution current predictions is schematically illustrated in Fig. 6.

In order to improve the GP model predictions for the extrapolation dataset, an expanded input domain $\mathcal{C}_2$, whose limits are shown in Table 3, is considered. It is noted that the range of WL thicknesses in Table 3 span all reaction control regimes delineated by Tomashov[14] from the bulk immersion regime to the thin-film regime. Subsequently, the input dataset that was originally used to calibrate the surrogate model in conjunction with the MaxPro augmentation approach[36] is used to sequentially add candidate inputs within the larger input domain shown in Table 3. This set of candidate inputs is denoted as $\mathcal{D}_{ex}$.

Fig. 4 | Probability Density Plots of the input configurations of the extrapolation dataset compared to the Initial Dataset. Probability Density Functions of the (**a**) Cathode Length and (**b**) WL thickness showing that the extrapolation dataset extends beyond the input domain $\mathcal{C}_1$ in Table 1.
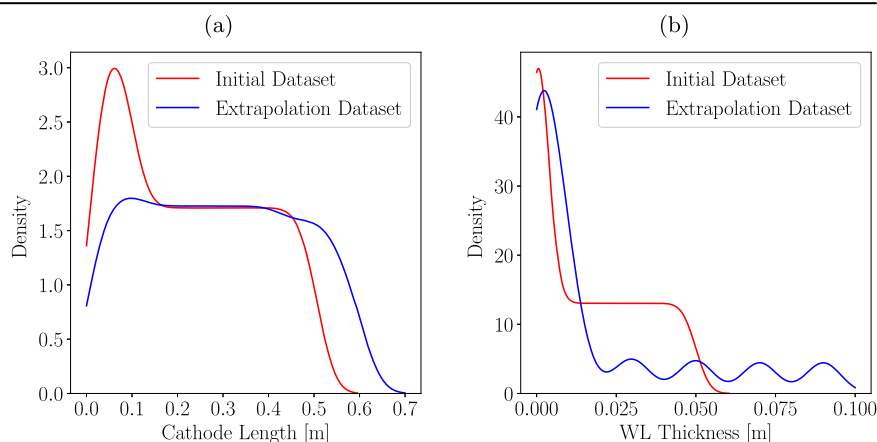
Fig. 5 | Graphic representation of model performance on the extrapolation dataset. a Parity plots showing the performance of the fused surrogate models trained using the predictive standard deviation (PS) and expected information gain (EI) on the extrapolation dataset. **b** Proportion of absolute error for the *K*-fold cross validation predictions, the predictions of the test set and the extrapolation set using the predictive standard deviation (PS) and expected information gain (EI) respectively.

(a) (b)
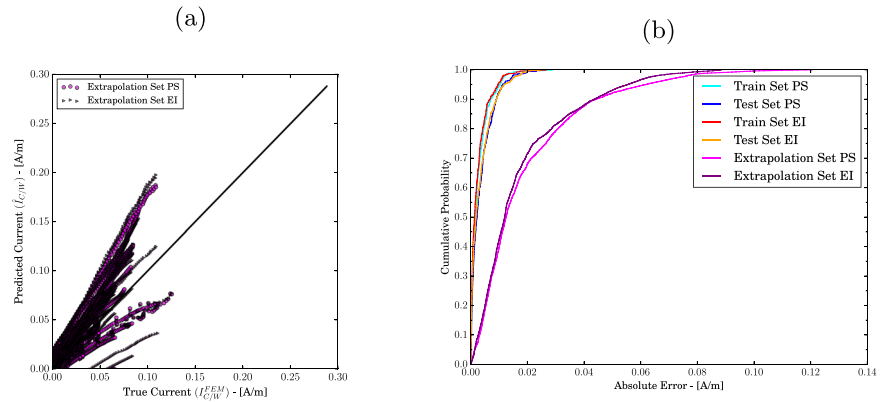
Fig. 6 | Illustration of the workflow used to improve out-of-distribution predictions and dataset enlargement. The purple arrows indicate data acquisition steps, including experiments and simulations. The orange arrows indicate AL loops. The green arrows indicate the dataset expansion by utilizing the staggered AL protocol on a larger domain ($\mathcal{C}_2$) to identify promising candidates.
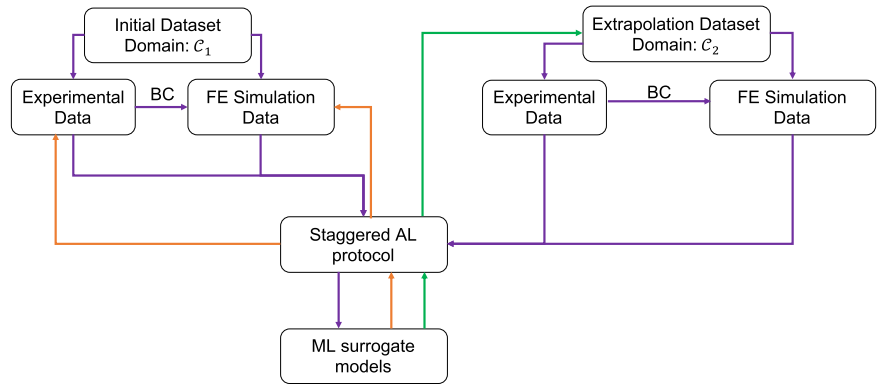
The AL protocol is leveraged to select a batch of input configurations $\mathcal{D}_{ex}^{opt}$ from $\mathcal{D}_{ex}$ for further exploration. To perform this selection, temperatures and salt concentrations at which polarization curves need to be obtained are identified. To ascertain these values, the influence of the geometric parameters from the GP model are marginalized and are obtained through a staggered GP model. This staggered GP model was used to obtain estimates of the acquisition function, which only depends on the temperature and salt concentration. Next, the weighted *K*-means clustering algorithm is used to identify four distinct combinations of temperatures and salt concentrations, by separating $\mathcal{D}_{ex}$ into four clusters, and subsequently calculating the cluster centers. The polarization curves are experimentally collected at these values.

Once the polarization curves are measured, the optimal batch of geometric parameters (cathode length and water layer thickness) for the FE simulations needs to be identified. To obtain these values, the temperature and salt concentration values are set equal to the cluster centers identified in the previous step. The GP model thus retains its dependency on the geometric parameters. The acquisition function is recalculated in terms of the geometric parameters. The optimal batch of geometric parameters, for a given temperature and salt concentration, are identified using another application of the weighted K-means clustering algorithm. 10 distinct

configurations of the geometric parameters for each combination of the temperature and salt concentration are used. Therefore, a single AL iteration with this strategy adds 40 datapoints (10 unique geometric configurations at each of the four environmental conditions) to the training dataset.

Figure 7 shows the optimal design points identified for both acquisition functions. One can clearly see that optimal design $\mathcal{D}_{ex}^{opt}$ for both acquisition functions is not uniform in the two-dimensional projections. This is due to the clustering algorithm ultimately prioritizes design points in $\mathcal{D}_{ex}$ that exhibit the highest value of the acquisition function. Most of the new input configurations are sampled from $\mathcal{C}_2 - \mathcal{C}_1$, mainly because the potential for improving the surrogate model is significantly lower inside the domain where the acquisition function values have saturated.

The surrogate models were re-calibrated after concatenating the newly acquired dataset to the initial dataset. The updated fused model is then used to obtain predictions for the extrapolation dataset (the dataset from ref. 16). Figure 8 shows the performance of the updated model on the extrapolation dataset. Specifically, Fig. 8b compares the predictions of $I_{C/W}$ obtained using the updated fused model and the FE simulations respectively for the extrapolation dataset. Figure 8a shows a qualitative improvement since the points appear closer to the diagonal line (which represents the true values). Figure 8b further reinforces this assessment as it clearly shows that the cumulative probability of absolute error for the extrapolation set improved significantly after the AL iteration. Recall that this improvement was obtained using only 40 new observations. A comparison of the mean absolute errors is presented in Table 4. The results shown in Fig. 8 as well as in Table 4 clearly show that both acquisition functions have provided overall improvements to the quality of predictions for the extrapolation dataset. This highlights the the efficiency of the proposed framework in searching a high-dimensional space for input configurations with the highest potential to improve the performance of the surrogate model. However, in the case of the PS, some predictions are seen to be less accurate than the ones provided by the originally calibrated surrogate model. In contrast, using the EI as the

Table 3 | Expanded input domain selected for exploration using the active learning strategy compared against the input domain used for the generation of the initial dataset

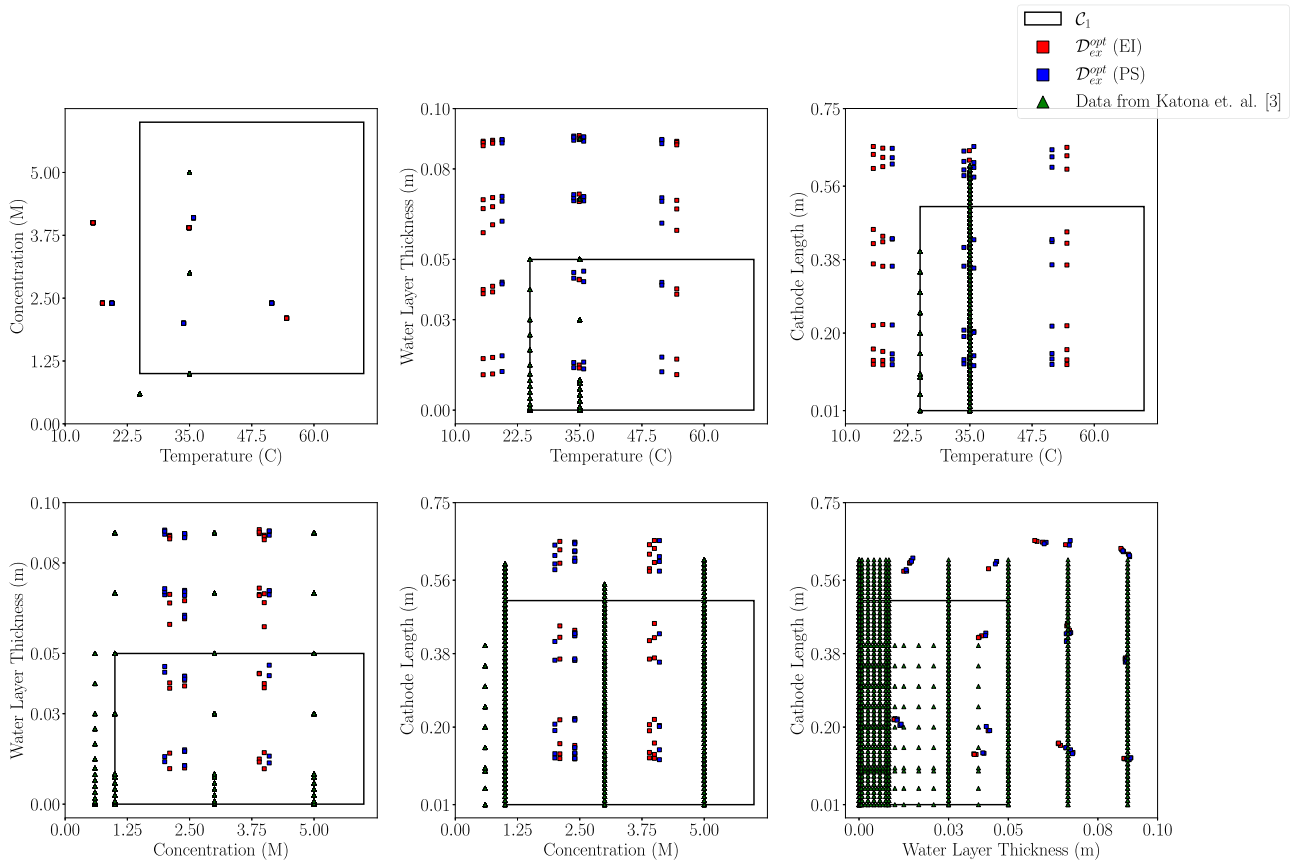| Parameter	Range (Initial Dataset)	Range (Extrapolation Dataset)
Temperature	[25,45] °C	[10,60] °C
Salt concentration	{1,3,5} M	[0.001,5] M
WL thickness	[7E-06,0.05] m	[7E-06,0.1] m
Cathode Length	[0.01,0.5] m	[0.01,0.75] m

Fig. 7 | $\mathcal{C}_1$ is indicated with a black rectangle in each projection (from ref. 16). The green triangles represent the input configurations explored in an alternate study (see refs. 17,16). The blue and red squares represent the points selected using the batch acquisition function for exploring a larger domain has resulted in uniformly active learning approach from the expanded input domain $\mathcal{C}_2$, with the predictive standard deviation and expected information gain being used as weights in the clustering algorithm, respectively. Note that 'CL' refers to cathode length.

Fig. 8 | Graphic representation of model performance on the extrapolation dataset. a Parity plots showing the performance of the GP surrogate models trained using the predictive standard deviation (PS) and expected information gain (EI) on the extrapolation dataset. **b** Proportion of absolute error for the predictions of the extrapolation set using updated model trained on the points identified using the predictive standard deviation (PS) and expected information gain (EI) respectively compared to the original models.
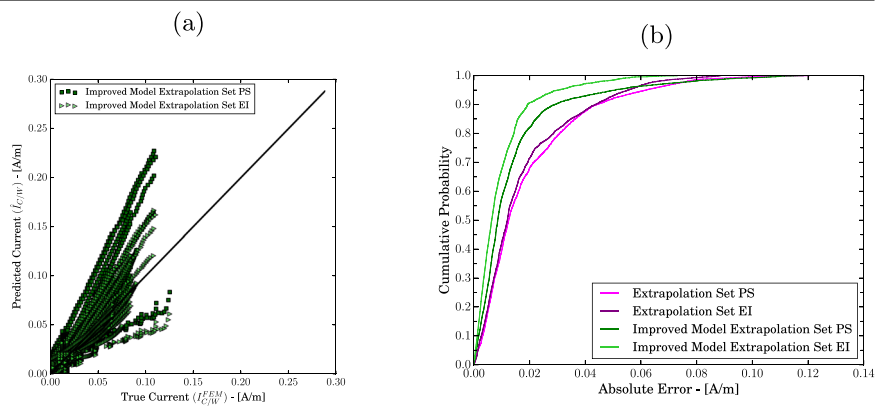
acquisition function for exploring a larger domain has resulted in uniformly better predictions for the extrapolation dataset, as revealed by a point-wise comparison of the absolute errors obtained using the original and refined surrogate model respectively. Indeed, using the EI acquisition reduces the

mean absolute errors on the extrapolation dataset by nearly 50% from 0.017mA to 0.009mA (see Table 4).

Discussion

In this study, a two-step active learning protocol was developed and deployed to calibrate a low-cost surrogate model for the available current in a galvanic couple for a wide range of in-service conditions. The influence of geometric, environmental and chemical factors on the current are considered. The surrogate model is trained on a dataset of FE simulations, and experimental polarization curves are used to inform the BCs in the FE simulations. The framework calibrates an initial surrogate model using a sparse FE dataset. The two-step protocol identifies the optimal configuration(s) for gathering additional polarization curves by marginalizing the output of the surrogate model with respect to the geometric parameters.

Table 4 | Comparison of the improvements obtained to the predictions for the extrapolation dataset using the acquisition functions reveals that the expected information gain (EI) performs better than predicted standard deviation (PS)

Acquisition Function	Before active learning	After active learning
Predictive Standard Deviation	0.019 A/m	0.014 A/m
Expected Information Gain	0.017 A/m	0.009 A/m

Once the polarization curves are gathered, FE simulations are used to gather current data, and the surrogate model is updated using the new datapoints. This iterative process rapidly establishes a highly accurate surrogate while minimizing the total number of simulation runs and experiments. Exploration versus exploitation trade-offs are taken into account by identifying multiple sets of input configurations, with the help of weighted K-means clustering, to collectively refine the current predictions to the maximum possible extent. Two popular choices are considered for the acquisition function, and the rate at which they refine the surrogate model with additional data are compared. Moreover, multiple forms of surrogate models are optimally combined using information fusion approaches to improve the robustness of the current predictions. The efficacy of this protocol is demonstrated for interpolative and extrapolative case studies. In both cases, large gains in performance (both in terms of error and parsimony of data) are achieved. Since the establishment of the surrogate model is a one-time cost,this approach reduces the time taken to obtain current predictions under alternative environmental conditions or geometric configurations by roughly four orders of magnitude. The protocols presented in this study are highly relevant for the rapid screening of various alloys under corrosive conditions.

The AL framework applied a single iteration using two different acquisition functions. The EI acquisition function demonstrated superior performance gains compared to the PS acquisition function. It's important to note that EI considers the domain-wide average of the GP model as the Quantity of Interest (QoI) and seeks new design points with the greatest potential to influence this average value. The domain-wide averaging step regularizes the QoI, making it less sensitive to performance variations exhibited by the GP model, particularly outside the original input domain (C_1, see Fig. 7). This feature is especially valuable in our study due to the numerous physical and chemical phenomena inherent in the corrosion process, requiring robust models and metrics. While the error reduction in the EI framework is significant, further reductions may be possible in subsequent iterations of this process. Limited prior research has compared EI and PS acquisition functions in terms of the performance gains per additional data point[7] and has confirmed the generally greater efficacy of information gain-based criteria.

The framework also helps to overcome problems related to the sparsity and high dimensionality of data in the corrosion field. The current work showcased four input variables (i.e., chloride concentration, temperature, water layer, and cathode length), however, such a problem could also consider galvanic couple pair, elemental composition of the alloys, corrosion parameters such as corrosion potential, among others. Applying an AL framework optimizing corrosion response helps with current materials exposed to various environments but could also be applied to materials discovery in corrosion[37–40]. Such work has also been highlighted in other electrochemical fields such as batteries due to the high-dimensional data[41]. The remarkable efficacy of this framework arises from the usage of an information metric (the acquisition function) to guide data acquisition. Indeed, GP models along with AL strategies have been utilized to refine predictions for functions of up to 40 variables[42]. Further, surrogate models can be used to efficiently screen input variables that do not have a discernible effect on the corrosion property of interest[43], which is especially useful when the inputs are high-dimensional.

As previously noted[25,29], FEM modeling incurs various costs specifically when incorporating experimentally derived boundary conditions (i.e., polarization scans). The developed AL framework allows for the separate selection of environmental and geometric parameters that will reduce the overall error in the trained surrogate model. The framework could have easily been applied only to select geometric parameters for FE simulations eliminating the need to perform experiments. However, the goal in the present work was to show that the framework could reduce errors by adding more experiments and geometries. In this regard, the staggered approach developed in our work allows for a seamless and efficient assimilation of additional experimental and simulation data into the existing dataset. In addition, our approach is especially useful in scenarios where the

experimental and simulation response are controlled by different sets of factors, which is often the case when simplifying assumptions are used to construct the physics-based model. It is noted that this approach is inspired by the Bayesian calibration framework[44]. More specifically, while the experimental and simulation responses utilized in the calibration process are considered to be a function of controllable inputs, the simulation response additionally depends on the material properties. Predictions for new input configurations can be obtained by marginalizing the simulation response w.r.t the probability distribution function associated with the material properties.

Machine Learning (ML) serves as a powerful tool for discerning the optimal scenarios for additional Finite Element (FE) simulations within a dataset, presenting broader implications for the materials degradation community. Acquisition function-based AL protocols offer a quantitative method to assess the potential benefits of adding further simulations, such as exploring geometric variations, to the candidate dataset, thereby enhancing the fidelity of the surrogate ML model (see Fig. 2). In the course of our study, a total of 2520 simulations were conducted. Notably, the minimum error value was attained only 650 optimally selected simulations. Thus, the integration of ML with Finite Element Modeling (FEM) allows users to limit parametric sweeps while ensuring efficient exploration of the domain of interest.

This efficiency in optimizing computational resources while maintaining model accuracy holds promise not only for corrosion-related work but also for a broader spectrum of materials degradation investigations. Beyond its direct application, our approach has the potential to inspire new experiments aimed at reducing uncertainty in ML predictions, offering a systematic means to forecast future events with an informed understanding of prediction uncertainties. Consequently, the protocols developed in this work have the potential to address longstanding scientific gaps across diverse materials degradation challenges[40].

Methods
Experimental setup
A three-electrode cell was used with SS304L ($0.196\,\mathrm{cm}^2$) as the working electrode, a silver/silver chloride reference electrode, and a Platinum-coated Niobium mesh as the counter electrode. Samples were immersed in the solution (ACS grade) and temperature of interest in a 150 mL water-jacketed cell and sparged with lab air to maintain a solution saturated in O_2. The solution temperature was controlled to ± 0.1 °C. Polarization measurements were performed at scan rates of 0.167 after 1 h of open circuit potential (OCP) equilibration and stabilization. The scans were performed from OCP to $-1.4\,\mathrm{V_{SCE}}$.

FE simulations of galvanic corrosion
The FE models created in this study are described in more details in previous works[16]. Briefly, the modeling framework is based on the Laplace Equation[45]. COMSOL Multiphysics v6.1 (COMSOL, Inc., Burlington, MA) software was utilized. The potential and current distributions along a SS304L cathode coupled galvanically with AA7050-T7451 were modeled using a simplified, two-dimensional geometry. The anode length was fixed at 0.01 m, and cathode length and WL thicknesses were dictated by the active learning framework. The lower limit of WL was set by the natural convection boundary layer. The values used for solution conductivity (κ) were obtained from OLI Studio Analyzer 9.5 (OLI Systems, Inc., Cedar Knolls, NJ). The anodic boundary conditions (B.C.) were the anodic kinetics of AA7050 (anodic B.C.) obtained from Liu et al.[17] and were assumed to be constant across all chloride concentrations and temperatures. All calculations were assumed to be steady state and the polarization scans were not dependent upon pH. Both anodic and cathodic currents were measured, and it was ensured currents were identical. The temperatures and salt concentrations shown in Table 1 are representative of commonly encountered environmental conditions. The range of WL thicknesses explored using the FE simulations includes the ohmic resistance-controlled regime as well as the M-T kinetics controlled regime.

Neural Network (NN) ensembles for prediction and uncertainty quantification of cathodic current output

Neural networks (NN) are universal function approximators[6] which map input to output using a combination of linear transformations and non-linear activations. Each linear transformation results in a new set of features called 'hidden units', whereas the non-linear activation allows us to capture complex input-output relationships. A linear transform together with its associated non-linear activation is called a 'layer'.

Neural network parameters, also known as weights, are typically initialized randomly, and adjusted during a training phase of the network which measures the error (also called loss) between network predictions and ground truth values of a dataset[46]. During this process, the gradients of the loss function are used to update the weights. This iterative process results in a trained network capable of approximating the function that defines the input-output pairs in the dataset

In our study, each neural network surrogate model has three layers, with 6464 and 32 hidden units respectively in each layer. In addition, we use the Rectified Linear Unit (ReLU) activation function to account for non-linear relationships between the input configurations and the galvanic current. However, we utilize different random weight initializations to make five independent models. Each random weight initialization allows the surrogate model to have a different approximation of the ground truth data distribution, which in turn allows us to quantify the uncertainty in the predicted response. The optimal architecture of the neural network is identified by maximizing the K-fold cross-validation performance. This approach facilitates achieving an optimal balance between model accuracy and robustness. Integrating ensemble techniques with K-fold cross-validation enables training a neural network model without generalization concerns, despite its high dimensionality.

Gaussian process regression models for cathodic current output

Gaussian Process Regression is a nonlinear, probabilistic interpolation approach for emulating the response obtained from complex FE models[5]. In this approach, the dataset of responses is collectively represented using a multivariate Gaussian distribution, according to

$$\mathbf{Y}^{sim} \sim \mathcal{N}\big(\mathbf{0}, \sigma^2 \mathbf{R}\big(\mathbf{X}^{sim}, \mathbf{X}^{sim}\big)\big) \tag{1}$$

Where $\mathbf{Y}^{sim}$ is the dataset of simulation responses, $\mathbf{X}^{sim}$ is the dataset of inputs, σ^2 is a scaling factor for the variance exhibited by the GP, and $\mathbf{R}$ is a positive-definite covariance matrix. The entries in this matrix are specified with the help of a kernel function, which is typically a stationery, decaying function of the L^2 distance between a pair of inputs. The most popular choice for the kernel function is the Automatic Relevance Determination - Squared Exponential (ARD-SE) kernel[5,7,8,12], which is specified according to

$$R(\mathbf{x}_i, \mathbf{x}_{i'}) = \exp\left\{ -\sum_{j=1}^{D} \frac{(x_{ij} - x_{i'j})^2}{l_j^2} \right\} \tag{2}$$

Where x_{ij} refers to the j-th feature of the input $\mathbf{x}_i$, and $\{l_j\}_{j=1}^{D}$ are the length-scales associated with each feature. These length-scales capture the distance along each feature axis beyond which outputs are considered to be uncorrelated.

GP models consider the joint Gaussian distribution of of the training and test datapoints in order to make predictions. Mathematically, the joint distribution is specified as

$$\begin{bmatrix} \mathbf{Y}^{sim} \\ \hat{y} \end{bmatrix} \sim \mathcal{N}\left(\mathbf{0}, \sigma^2 \begin{bmatrix} \mathbf{R}(\mathbf{X}^{sim}, \mathbf{X}^{sim}) & \mathbf{R}(\mathbf{X}^{sim}, \mathbf{x}^*) \\ \mathbf{R}(\mathbf{x}^*, \mathbf{X}^{sim}) & R(\mathbf{x}^*, \mathbf{x}^*) \end{bmatrix} \right) \tag{3}$$

The predictions for test inputs are obtained by conditioning the joint distribution using the existing simulation dataset (see Eq. (1)). This results in

a normal distribution for the predictions given by

$$\hat{y}(\mathbf{x}^*) \sim \mathcal{N}\big(m(\mathbf{x}^*), s^2(\mathbf{x}^*)\big) \tag{4}$$

$$m(\mathbf{x}^*) = \mathbf{R}(\mathbf{x}^*, \mathbf{X}^{sim}) \mathbf{R}^{-1}(\mathbf{X}^{sim}, \mathbf{X}^{sim}) \mathbf{Y}^{sim} \tag{5}$$

$$s^2(\mathbf{x}^*) = \sigma^2 \big\{ 1 - \mathbf{R}(\mathbf{x}^*, \mathbf{X}^{sim}) \mathbf{R}^{-1}(\mathbf{X}^{sim}, \mathbf{X}^{sim}) \mathbf{R}(\mathbf{X}^{sim}, \mathbf{x}^*) \big\} \tag{6}$$

The optimal values of the lengthscales is estimated by minimizing the log-likelihood of the normal distribution that represents the training dataset (see Eq. (1)).

As mentioned in Sec. 2, the GP surrogate model is used to calculate the two acquisition functions. The first acquisition function (the predictive standard deviation) is calculated as the square-root of the variance in Eq. (6). The second acquisition function (the expected information gain) is calculated following the approach developed by Pandita et al.[7]. The expected information gain is calculated as

$$\mathcal{G}\left(\mathbf{x}^*, \{l_j\}_{j=1}^{D}\right) = \log\left(\frac{\sigma_1}{\sigma_2(\mathbf{x}^*)}\right) + \frac{1}{2}\left(\frac{\sigma_2(\mathbf{x}^*)}{\sigma_1}\right)^2 + \frac{1}{2}\left(\frac{\nu(\mathbf{x}^*)}{\sigma_1 s(\mathbf{x}^*)}\right)^2 \tag{7}$$

Where σ_1^2 is the integral of the variance in Eq. (6) across the entire input domain for a given dataset, $\sigma_2^2(\mathbf{x}^*)$ is the integral of the variance in Eq. (6) *after augmenting the initial dataset with the prediction* $\hat{y}(\mathbf{x}^*)$. The quantity $\nu(\mathbf{x}^*)$ is specified as

$$\nu(\mathbf{x}^*) = \epsilon_n(\mathbf{x}^*) - \boldsymbol{\epsilon}_n^T(\mathbf{x}^*) \mathbf{R}^{-1} \mathbf{R}(\mathbf{X}^{sim}, \mathbf{x}^*) \tag{8}$$

With

$$\epsilon_{n,i} = \sigma^2 \left(\frac{\pi}{2}\right)^{d/2} \prod_{j=1}^{D} \left\{ l_j \left(\mathrm{erf}\left(\frac{1 - x_{ij}}{\sqrt{2}l_j}\right) - \mathrm{erf}\left(\frac{-x_{ij}}{\sqrt{2}l_j}\right) \right) \right\} \tag{9}$$

Fusion of current predictions from multiple surrogate models

Once the surrogate models are established with sufficient fidelity, we consider fusing their predictive responses to build a more robust and accurate predictor. We build a fused predictor that is a linear combination of the GP and NN surrogate model predictions according to

$$\hat{y}^f = \mathbf{w}^T \hat{\mathbf{y}}, \quad \hat{\mathbf{y}} = [\hat{y}_{GP}, \hat{y}_{NN}] \tag{10}$$

Where $\hat{y}^f$ is the fused prediction, $\mathbf{w}^T$ is the weight assigned to each model (which varies with the input configuration), and $\hat{y}_{GP}$ and $\hat{y}_{NN}$ are the GP and NN surrogate model predictions respectively. We note that the fused predictor is unbiased, i.e. $\mathbf{w}^T\mathbf{1} = 1$. The variance associated with the fused predictor is specified according to

$$\mathrm{Var}(\hat{y}^f) = \mathbf{w}^T \boldsymbol{\Sigma} \mathbf{w} \tag{11}$$

Where $\boldsymbol{\Sigma} = \mathrm{Cov}(\hat{\mathbf{y}})$. The optimal value of the weight vector $\mathbf{w}$ for each input configuration is estimated by minimizing the variance in Eq. (11), subject to the constraint $\mathbf{w}^T\mathbf{1} = 1$. This optimization problem is solved with the help of Lagrange multipliers[47]. Further details regarding this approach can be found in refs. 33,34.

Data availability

The data supporting the findings in this study and the compiled models can be found at this Zenodo repository - https://doi.org/10.5281/zenodo.10655559.

Code availability

The codes used to generate the data in this study can be found at this Zenodo repository - https://doi.org/10.5281/zenodo.10655559.

References

1. Ren, Z., Ren, Z., Zhang, Z., Buonassisi, T. & Li, J. Autonomous experiments using active learning and AI. *Nat. Rev. Mater.* **8**, 563–564 (2023).
2. Parente, M. T. et al. Active learning-assisted neutron spectroscopy with log-gaussian processes. *Nat. Commun.* **14** https://doi.org/10.1038/s41467-023-37418-8 (2023).
3. Keane, A., Forrester, A. & Sobester, A. *Engineering Design via Surrogate Modelling: A Practical Guide* (AIAA, 2008). https://doi.org/10.2514/4.479557.
4. Santner, T. J., Williams, B. J. & Notz, W. I. *Some Criterion-based Experimental Designs*, 163–187 (Springer New York, 2003). https://doi.org/10.1007/978-1-4757-3799-8_6.
5. Rasmussen, C. E. & Williams, C. K. I. *Gaussian Processes for Machine Learning* (The MIT Press, 2005). https://doi.org/10.7551/mitpress/3206.001.0001.
6. Bishop, C. M. *Pattern Recognition and Machine Learning (Information Science and Statistics)* (Springer-Verlag, Berlin, Heidelberg, 2006).
7. Pandita, P., Bilionis, I. & Panchal, J. Bayesian optimal design of experiments for inferring the statistical expectation of expensive black-box functions. *J. Mech. Des.* **141** https://doi.org/10.1115/1.4043930 (2019).
8. Pandita, P., Tsilifis, P., Awalgaonkar, N. M., Bilionis, I. & Panchal, J. Surrogate-based sequential bayesian experimental design using non-stationary gaussian processes. *Comput. Methods Appl. Mech. Eng.* **385**, 114007 (2021).
9. Ashenafi, Y., Pandita, P. & Ghosh, S. Reinforcement learning-based sequential batch-sampling for bayesian optimal experimental design. *J. Mech. Des.* **144**, 091705 (2022).
10. Jones, D. R., Schonlau, M. & Welch, W. J. Efficient global optimization of expensive black-box functions. *J. Glob. Optim.* **13**, 455–492 (1998).
11. Khatamsaz, D. et al. Efficiently exploiting process-structure-property relationships in material design by multi-information source fusion. *Acta Mater.* **206**, 116619 (2021).
12. Venkatraman, A., McDowell, D. L. & Kalidindi, S. R. Bayesian analysis of parametric uncertainties and model form probabilities for two different crystal plasticity models of lamellar grains in $\alpha + \beta$ titanium alloys. *Int. J. Plast.* **154**, 103289 (2022).
13. Weilguny, L. et al. Dynamic, adaptive sampling during nanopore sequencing using bayesian experimental design. *Nat. Biotechnol.* **41**, 1018–1025 (2023).
14. TOMASHOV, N. D. Development of the Electrochemical Theory of Metallic Corrosion. *Corrosion* **20**, 7t–14t (2013).
15. Bard, A. & Faulkner, L. *Electrochemical Methods: Fundamentals and Applications* (Wiley, 2000).
16. Katona, R. M. et al. Editors' choice–natural convection boundary layer thickness at elevated chloride concentrations and temperatures and the effects on a galvanic couple. *J. Electrochem. Soc.* **168**, 031512 (2021).
17. Liu, C., Srinivasan, J. & Kelly, R. G. Editors' choice–electrolyte film thickness effects on the cathodic current availability in a galvanic couple. *J. Electrochem. Soc.* **164**, C845 (2017).
18. Moraes, C. & Kelly, R. A comparison of FEM results from the use of different governing equations in a galvanic cell part i: In the presence of a supporting electrolyte. *Electrochem. Acta* **469**, 143146 (2023).
19. Moraes, C. & Kelly, R. A comparison of FEM results from the use of different governing equations in a galvanic cell part II: Impact of low supporting electrolyte concentration. *Electrochem. Acta* **469**, 143153 (2023).
20. Simillion, H., Dolgikh, O., Terryn, H. & Deconinck, J. Atmospheric corrosion modeling. *Corros. Rev.* **32**, 73–100 (2014).
21. Boyce, B. et al. Machine learning for materials science: Barriers to broader adoption. *Matter* **6**, 1320–1323 (2023).
22. Dolgikh, O. et al. Simulation of the role of vibration on scanning vibrating electrode technique measurements close to a disc in plane. *Electrochem. Acta.* **203**, 379–387 (2016).
23. Palani, S., Hack, T., Deconinck, J. & Lohner, H. Validation of predictive model for galvanic corrosion under thin electrolyte layers: An application to aluminium 2024-cfrp material combination. *Corros. Sci.* **78**, 89–100 (2014).
24. Thébault, F., Vuillemin, B., Oltra, R., Allely, C. & Ogle, K. Modeling bimetallic corrosion under thin electrolyte films. *Corros. Sci.* **53**, 201–207 (2011).
25. Marshall, R. S. et al. Galvanic corrosion between coated al alloy plate and stainless steel fasteners, part 2: Application of finite element method and machine learning to study galvanic current distributions. *Corrosion* **79**, 157–173 (2023).
26. Topa, V., Demeter, A., Hotoiu, L., Deconinck, D. & Deconinck, J. A transient multi-ion transport model for galvanized steel corrosion protection. *Electrochem. Acta* **77**, 339–347 (2012).
27. Venkatraman, M. S., Cole, I. S. & Emmanuel, B. Model for corrosion of metals covered with thin electrolyte layers: Pseudo-steady state diffusion of oxygen. *Electrochem. Acta* **56**, 7171–7179 (2011).
28. Hœrlé, S., Mazaudier, F., Dillmann, P. & Santarini, G. Advances in understanding atmospheric corrosion of iron. ii. mechanistic modelling of wet-dry cycles. *Corros. Sci.* **46**, 1431–1465 (2004).
29. de Oca Zapiain, D. M. et al. Accelerating fem-based corrosion predictions using machine learning. *J. Electrochem. Soc.* **171**, 011504 (2024).
30. Bansal, P. et al. Physics-informed machine learning assisted uncertainty quantification for the corrosion of dissimilar material joints. *Reliab. Eng. Syst. Saf.* **227**, 108711 (2022).
31. Pinsler, R., Gordon, J., Nalisnick, E. & Hernández-Lobato, J. M. Bayesian batch active learning as sparse subset approximation. In Wallach, H. et al. (eds.) *Advances in Neural Information Processing Systems*, vol.32 (Curran Associates, Inc., 2019).
32. Gardner, J., Pleiss, G., Weinberger, K. Q., Bindel, D. & Wilson, A. G. Gpytorch: Blackbox matrix-matrix gaussian process inference with gpu acceleration. In Bengio, S. et al. (eds.) *Advances in Neural Information Processing Systems*, vol. 31 (Curran Associates, Inc., 2018).
33. Allaire, D. & Willcox, K. *Fusing information from multifidelity computer models of physical systems*, 2458–2465 (IEEE, 2012).
34. Goldstein, M. & Rougier, J. Reified bayesian modelling and inference for physical systems. *J. Stat. Plan. Inference* **139**, 1221–1239 (2009).
35. Joseph, V. R., Gul, E. & Ba, S. Maximum projection designs for computer experiments. *Biometrika* **102**, 371–380 (2015).
36. Joseph, V. R., Gul, E. & Ba, S. Designing computer experiments with multiple types of factors: The MaxPro approach. *J. Qual. Tech.* **52**, 343–354 (2019).
37. Yuwono, J. A. et al. A computational approach for mapping electrochemical activity of multi-principal element alloys. *npj Mater. Degrad.* **7**, 87 (2023).
38. Sur, D. et al. An experimental high-throughput to high-fidelity study towards discovering al-cr containing corrosion-resistant compositionally complex alloys. *High. Entr. Alloy. Mater.* **1**, 336–353 (2023).
39. Roy, A. et al. Machine-learning-guided descriptor selection for predicting corrosion resistance in multi-principal element alloys. *npj Mater. Degrad.* **6**, 9 (2022).
40. Scully, J. R. & Balachandran, P. V. Future frontiers in corrosion science and engineering, part iii: The next "leap ahead" in corrosion control may be enabled by data analytics and artificial intelligence. *Corrosion* **75**, 1395–1397 (2019).

41. Liu, Y., Guo, B., Zou, X., Li, Y. & Shi, S. Machine learning assisted materials design and discovery for rechargeable batteries. *Energy Storage Mater.* **31**, 434–450 (2020).

42. Yin, J. & Du, X. Active learning with generalized sliced inverse regression for high-dimensional reliability analysis. *Struct. Saf.* **94**, 102151 (2022).

43. Santner, T. J., Williams, B. J. & Notz, W. I. *Sensitivity Analysis and Variable Screening*, 247–297 (Springer New York, New York, NY, 2018).

44. Santner, T. J., Williams, B. J. & Notz, W. I. *Calibration*, 299–379 (Springer New York, New York, NY, 2018).

45. Liu, C. & Kelly, R. G. A Review of the Application of Finite Element Method (FEM) to Localized Corrosion Modeling. *Corrosion* **75**, 1285–1299 (2019).

46. Zhou, Z.-H., Wu, J. & Tang, W. Ensembling neural networks: Many could be better than all. *Artif. Intell.* **137**, 239–263 (2002).

47. Boyd, S. P. & Vandenberghe, L. *Convex optimization* (Cambridge University Press, Cambridge, UK, 2004).

Author contributions
Aditya Venkatraman: Conceptualization, Data curation, Investigation, Writing - original draft, Writing - review & editing; Ryan Katona: Conceptualization, Data curation, Investigation, Writing - original draft, Writing - review & editing; David Montes de Oca Zapiain: Conceptualization, Data curation, Investigation, Methodology, Supervision, Validation, Visualization, Writing - original draft, Writing - review & editing; Demitri Maestas: Conceptualization, Investigation, Writing - original draft, Writing - review & editing; Matthew Roop: Data curation, Writing - review; Philip Noell: Supervision, Project administration, Funding acquisition, Writing - review & editing.

Competing interests
The Authors declare no competing interests.

Additional information
Correspondence and requests for materials should be addressed to David Montes de Oca Zapiain.

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

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