

SPARSE-DYNAMICS FEATURES FOR SAMPLE-EFFICIENT BAYESIAN CALIBRATION OF RESISTANCE SPOT WELDING MODELS

B. VERKENS^{1,2*}, P. LI³, M. FAES^{4,5}, P. VAN RYMENANT¹ AND D. MOENS^{1,2}

¹⁾ KU Leuven, Department of Mechanical Engineering,
Jan Pieter de Nayerlaan 5, 2860 Sint-Katelijne-Waver, Belgium,
email: {bouwe.verkens, patrick.vanrymenant, david.moens}@kuleuven.be

²⁾ FlandersMake@KU Leuven, Belgium

³⁾ Key Laboratory of Urban Security and Disaster Engineering of Ministry of Education, Beijing
University of Technology, Beijing 100124, China, email: li.peipei@bjut.edu.cn

⁴⁾ TU Dortmund, Chair for Reliability Engineering, Leonhard-Euler-Strasse 5, 44227 Dortmund,
Germany, email: matthias.faes@tu-dortmund.de

⁵⁾ Tongji University, International Joint Research Center for Engineering Reliability and Stochastic
Mechanics, Shanghai 200092, China

Key words: Resistance spot welding, Finite element simulation, Uncertainty, Validation, Electro-thermomechanical model.

Summary. *Calibrating a finite element model of resistance spot welding (RSW) requires inferring parameters that cannot be measured directly, such as contact film behaviour and temperature-dependent material properties. Current practice tunes these unknown parameters deterministically, which cannot represent the spread observed between nominally identical welds. Bayesian model updating recovers the parameter distribution instead. However, classical samplers require 10^3 – 10^6 evaluations of a fully coupled electro-thermo-mechanical model, which makes these methods infeasible. A second obstacle is the type of calibration data the RSW finite element model provides. Besides the scalar outputs (nugget geometry), the model also produces time-series signals (dynamic resistance, electrode displacement, ...), whose samples are strongly autocorrelated. Thus entering them directly into an independent-sample likelihood yields an overconfident likelihood and an artificially narrow posterior. This work addresses these two challenges separately. The number of expensive evaluations is reduced by Streamlined Bayesian Active Learning Cubature. This method places a Gaussian-process surrogate on the log-likelihood and selects each new simulation where that surrogate is least informative. The time-series signals are compressed by Sparse Identification of Nonlinear Dynamics with control into a small set of sparse coefficients that act as calibration features. The workflow is validated on a resistively heated wire with known ground truth. The posterior of the unknown parameter(s) is found in tens of model evaluations and is compared against more classical Bayesian model updating samplers.*

1 INTRODUCTION

Resistance spot welding (RSW) is a key joining process for metallic component assembly in the industry. Its field of application ranges from the sealing of nuclear fuel rod end-plugs to the interconnection of cells in electric-vehicle (EV) battery packs [6, 12]. However, the dominant application of this process remains the assembly of the automotive body-in-white, in which the majority of the joints are produced by RSW, with approximately 3000 spot welds per vehicle [4]. The diameter of the weld nugget is highly correlated with the strength of the joint and therefore is considered as an important quality metric [10, 27]. However, the geometry of this weld nugget is typically assessed through destructive or ultrasonic testing, making the evaluation of weldability time-consuming and material-intensive. This has motivated the use of numerical simulations to study the coupled electro-thermal-mechanical mechanisms governing this process. Since the work of Nied [17], finite element (FE) simulations have been extended considerably to capture nugget growth [7, 24], melt pool dynamics [14], expulsion [26], electrode degradation [18], dissimilar welding [23], process signals [15, 25]. Meanwhile, the growing trend of using machine learning (ML) for process monitoring and optimisation of the process [28] has introduced further motivation for numerical models that can mimic the actual process in order to generate labelled training data at a lower cost. The accuracy of these simulations depends strongly on model parameters can't be measured directly. The electrical and thermal contact behaviour at the electrode-sheet and sheet-sheet interfaces dominate the heat generation in the beginning of the process [22]. Yet there is no generally accepted contact-resistance model that takes into account all the different parameters [9] influencing the resistance. These unknown model parameters are in current literature estimated through deterministic optimisation from experimental data. However, in practice, variation in spot welding quality is seen despite using nominally identical welding conditions on the same machine [8, 21]. This is because electrode wear, coating variability, and part fit-up all affect the contact-interface parameters, so a single deterministic point-estimate cannot represent this weld-to-weld variability. Bayesian model updating instead infers a posterior distribution over the parameters, capturing both the parameter values supported by the data and the uncertainty associated with them. However, classical Bayesian methods rely on sampling the posterior directly, such as Markov chain Monte Carlo (MCMC) [1] or transitional MCMC (TMCMC) [5] and require typically 10^3 – 10^6 forward model evaluations, which is infeasible when a single evaluation is an numerically expensive RSW simulation. Building a surrogate of this expensive RSW simulation reduces this cost by replacing it by an inexpensive approximation, typically a polynomial chaos expansion (PCE) or a Gaussian process (GP), trained on a limited number of expensive runs. Recent research in these surrogate-based techniques has provided an elegant solution that minimises the needed forward runs even further, by modelling the likelihood function directly and reducing the amount of forward model evaluations substantially [13].

The RSW FE model generates two types of outputs that can be used as calibration targets. Besides the scalar outputs (nugget geometry), the model also produces time-series signals (dynamic resistance, electrode displacement, ...), whose samples are strongly autocorrelated. This

is because consecutive samples are governed by the same ‘slowly’ varying thermal and mechanical states. Entering such signals into the likelihood, made of independent samples, overstates its information content. The resulting likelihood is overconfident and the posterior is artificially narrow, which defeats the purpose of the stochastic calibration [20]. Therefore, the time-series observations must be reduced to a lower dimensional set of features that preserve the sensitivity to the unknown parameters while removing the temporal redundancy.

In this work, a surrogate-based Bayesian calibration workflow is developed that combines SINDYc feature compression of time-series observations with SBALC inference, as a first step towards stochastic calibration of the RSW FE model. The workflow is demonstrated on an analytical benchmark sharing the Joule-heating mechanism of RSW. An uninformative prior is assumed and the posterior is calibrated against synthetic temperature histories with known ground truth. The proposed workflow is motivated by RSW, but it could be applied to any calibration problem combining an expensive forward model with autocorrelated time-series observations.

The remainder of this paper is organized as follows. Section 2 formulates Bayesian model updating with time-series observations and states the autocorrelation problem and how this is tackled. It gives a brief introduction to the methods of SINDy and SBALC. Section 3 defines the benchmark used and establishes the analytical dependence of the SINDy coefficients on the unknown parameters. Section 4 reports the joint posterior recovered by SBALC and verifies it against MCMC and TMCMC. Section 5 concludes and discusses the approach for implementation to the RSW FE model.

2 METHODOLOGY

Let $\theta \in \mathcal{D}_\theta \subset \mathbb{R}^d$ denote the vector of unknown parameters that need to be inferred, where $\mathcal{D}_\theta$ is the support of the prior $f(\theta)$. The forward model $\mathcal{M}$ maps a parameter vector onto a set of scalar outputs $y_s(\theta)$ and a sampled time series $y_{1:N}(\theta)$ of with N datapoints or samples. The latter type of input will be focused on in the remainder of this paper. The observations are denoted as $\mathbf{Y}_{\text{obs}}$.

2.1 Bayesian model updating with time-series observations

The calibration framework follows the Bayesian approach described in [11]. Given observations $\mathbf{Y}_{\text{obs}}$, Bayes’ theorem updates the prior distribution $f(\theta)$ to the posterior,

$$f(\theta \mid \mathbf{Y}_{\text{obs}}) = \frac{L(\mathbf{Y}_{\text{obs}} \mid \theta) f(\theta)}{c}, \quad c = \int_{\mathcal{D}_\theta} L(\mathbf{Y}_{\text{obs}} \mid \theta) f(\theta) d\theta, \quad (1)$$

where $L(\mathbf{Y}_{\text{obs}} \mid \theta)$ is the likelihood of the observations under the parameter vector θ , $f(\theta)$ is the prior density, and c is the evidence, which serves as an normalising constant of the posterior. Every evaluation of L requires one forward run of the model $\mathcal{M}$, and the computational cost of the updating procedure is therefore reported throughout this work as the number of forward runs. Classical samplers such as MCMC and TMCMC typically require 10^3 - 10^6 runs, which makes these methods infeasible to use when dealing with computationally expensive models.

The likelihood is constructed from the discrepancy between the observations and the model prediction assuming a Gaussian discrepancy with zero mean. For N_{obs} independent observations, the log-likelihood is

$$\mathcal{L}(\boldsymbol{\theta}) = \ln L(\boldsymbol{\theta}) = -\frac{1}{2} \sum_{k=1}^{N_{\text{obs}}} \mathbf{r}_k(\boldsymbol{\theta})^\top \boldsymbol{\Sigma}^{-1} \mathbf{r}_k(\boldsymbol{\theta}) + \text{const}, \quad (2)$$

with $\mathbf{r}_k(\boldsymbol{\theta}) = \mathbf{Y}_{\text{obs}}^{(k)} - \mathcal{M}(\boldsymbol{\theta})$ being the residual vector of the k -th observation, $\boldsymbol{\Sigma}$ the covariance matrix of the observation discrepancy.

However, if all the individual datapoints of the time series $y_{1:N}(\boldsymbol{\theta})$ are inserted into (2) as an N -dimensional observation vector, and the independent-sample assumption $\boldsymbol{\Sigma} = \sigma_T^2 \mathbf{I}_N$ is adopted, the log-likelihood becomes

$$\mathcal{L}_{\text{naive}}(\boldsymbol{\theta}) = -\frac{1}{2} \sum_{i=1}^N \frac{(\tilde{y}_i - y_i(\boldsymbol{\theta}))^2}{\sigma_T^2} + \text{const}, \quad (3)$$

where $\tilde{y}_i$ is the i -th noisy measurement, $y_i(\boldsymbol{\theta})$ is the corresponding model output, and σ_T is the standard deviation of the measurement noise. This would treat N strongly autocorrelated samples as independent evidence and the information is consequently overstated. This would make the posterior artificially narrow and its width no longer reflects the actual physical variability.

A solution to this would be the dimensional reduction of the time series into a small set of meaningful features. A feature map $\mathcal{F} : \tilde{y}_{1:N} \mapsto \boldsymbol{\xi} \in \mathbb{R}^{n_\xi}$ with $n_\xi \ll N$ replaces the N autocorrelated samples by n_ξ calibration targets. It's important that this dimension reduction $\mathcal{F}$ is a function of $\boldsymbol{\theta}$ alone, so that $\mathcal{L}(\boldsymbol{\theta})$ is well defined.

2.2 Feature extraction by sparse dynamics

Sparse identification of nonlinear dynamics (SINDy) recovers the governing equations of a dynamical system from measured state time-series by sparse regression onto a library of candidate functions [2]. As in this research, the model has a time-varying external input that is known, this known time-series input can be seen as a control variable that is added to the same library. This method is known as SINDy with control (SINDYc) [3]. Let $\mathbf{X}$ collect the measured states, $\dot{\mathbf{X}}$ their time derivatives, and $\mathbf{U}$ the sampled control input. The dynamics are then expressed as

$$\dot{\mathbf{X}} = \boldsymbol{\Phi}(\mathbf{X}, \mathbf{U}) \boldsymbol{\Xi}, \quad (4)$$

where $\boldsymbol{\Phi}(\mathbf{X}, \mathbf{U})$ is the library of candidate functions and $\boldsymbol{\Xi}$ is the coefficient matrix. The coefficients are obtained by sparse regression

$$\hat{\boldsymbol{\Xi}} = \arg \min_{\boldsymbol{\Xi}} \|\dot{\mathbf{X}} - \boldsymbol{\Phi}(\mathbf{X}, \mathbf{U}) \boldsymbol{\Xi}\|_2^2 + \lambda \|\boldsymbol{\Xi}\|_1, \quad (5)$$

where λ is the sparsity threshold that controls the number of retained terms. In this work the coefficients in $\boldsymbol{\Xi}$ are used as the calibration features.

2.3 Streamlined Bayesian Active Learning Cubature (SBALC)

Streamlined Bayesian Active Learning Cubature (SBALC) is a surrogate-based calibration recently developed by Li et al. which reduces the number of forward evaluations by placing a Gaussian-process surrogate on the log-likelihood $\mathcal{L}(\boldsymbol{\theta})$ rather than on the model response [13]. This is because $\mathcal{L}(\boldsymbol{\theta})$ is a scalar field over $\mathcal{D}_{\boldsymbol{\theta}}$ irrespective of the dimension of the model output, so that a single surrogate suffices where a response surrogate would require one model per output component. After n expensive evaluations the surrogate provides a mean $m_{\mathcal{L}_n}(\boldsymbol{\theta})$ and a variance function $\sigma_{\mathcal{L}_n}^2(\boldsymbol{\theta})$. The evidence of (1) is then estimated by Monte Carlo integration over the prior,

$$\hat{c} = \mathbb{E}_{\boldsymbol{\theta}} [\exp(m_{\mathcal{L}_n}(\boldsymbol{\theta}))], \quad (6)$$

where the expectation is evaluated over $N_M = 2 \times 10^4$ samples drawn from $f(\boldsymbol{\theta})$.

Each new expensive evaluation is placed where it is most informative for the evidence. The learning function

$$\text{LF}(\boldsymbol{\theta}) = \sigma_{\mathcal{L}_n}^2(\boldsymbol{\theta}) \cdot (e^{m_{\mathcal{L}_n}(\boldsymbol{\theta})+b\sigma_{\mathcal{L}_n}(\boldsymbol{\theta})} - e^{m_{\mathcal{L}_n}(\boldsymbol{\theta})-b\sigma_{\mathcal{L}_n}(\boldsymbol{\theta})}), \quad b = 1, \quad (7)$$

is evaluated over $\mathcal{D}_{\boldsymbol{\theta}}$. The learning function (7) concentrates the sampling in regions where simultaneously the surrogate is the most uncertain and where the bounds on the evidence estimation would shrink. From the surrogate variance, upper (c_{upper}) and lower bounds (c_{lower}) on the evidence can be derived, which are used as a stopping criterion. The procedure is terminated when the relative width of the evidence bounds satisfies

$$\frac{|c_{\text{upper}} - c_{\text{lower}}|}{\hat{c}} < 0.1 \quad (8)$$

over two consecutive iterations. Posterior samples are subsequently generated from the final surrogate by Sampling-Importance-Resampling (SIR) [19], which requires no further forward evaluations. The full derivation of the learning function and of the evidence bounds is given in [13].

3 ANALYTICAL MODEL

The framework of section 2 is first tested on an inexpensive benchmark, for which the ground truth is known. The dependence of the extracted features on the unknown parameters is available in closed form and a posterior from transitional Markov chain Monte Carlo (TMCMC) remains computable at an acceptable cost in order to compare it with the proposed framework.

A copper wire of diameter d and length L_0 is subjected to a time-varying current pulse $I(t)$ with amplitude I_{amp} (figure 1). The wire heats up through Joule-heating and loses heat to its environment at temperature T_{∞} by convection and radiation. The energy balance is

$$\rho c_p V \frac{dT}{dt} = R I(t)^2 - hA(T - T_{\infty}) - \epsilon \sigma A(T^4 - T_{\infty}^4), \quad (9)$$

where ρ is the density, c_p the specific heat capacity, V the volume of the wire, R the electrical resistance, A the lateral surface area, and $\sigma = 5.670 \times 10^{-8} \text{ W/(m}^2\text{K}^4)$ is the Stefan–Boltzmann




Figure 1: Schematic of the heated wire benchmark, with the time-variant input (left) and the time-variant output (right).

constant. In here the geometric and electrical quantities are derived from the wire dimensions as

$$A = \pi d L_0, \quad V = L_0 \frac{\pi d^2}{4}, \quad R = \frac{\rho_e L_0}{\pi d^2 / 4}, \quad (10)$$

with ρ_e being the electrical resistivity. The wire parameters are listed in Table 1. The temperature of the wire over time $T(t)$ is obtained by integrating the equation (9) with an adaptive Runge-Kutta scheme from the initial condition $T(0) = T_\infty$ over the full pulse duration. The unknown parameters are the convective heat transfer coefficient h and the surface emissivity ε , with true values $h_{\text{true}} = 8000 \text{ W/m}^2\text{K}$ and $\varepsilon_{\text{true}} = 0.8$.

Table 1: Parameters used in the analytical model

Symbol	Value	Symbol	Value
d	0.5 mm	T_∞	25 °C (298.15 K)
L_0	1.0 m	t_{start}	10 ms
ρ	8960 kg m ⁻³	t_{rise}	5 ms
c_p	380 J kg ⁻¹ K ⁻¹	t_{dur}	30 ms
ρ_e	17 nΩ m	I_0	2.5 kA
N	5000	σ_T	0.5 K

Noise model Synthetic observations are generated by evaluating the forward model at the true parameter values and assuming additive Gaussian noise as follows

$$\tilde{T}_i = T_i + \eta_i, \quad \eta_i \sim \mathcal{N}(0, \sigma_T^2), \quad i = 1, \dots, N. \quad (11)$$

where $\tilde{T}_i$ is the noisy temperature measurement at the i -th time sample and $\sigma_T = 0.5 \text{ K}$ is the noise standard deviation. The noise is applied to the temperature time-series only, since the current is known exactly.

3.1 Feature extraction

As described in section 2.2, SINDy is used as a dimensional reduction technique where the coefficients Ξ obtained by sparse regression are the extracted features from the time-series signals and are used as calibration targets. The time-series output of the temperature is the observed state and is expressed in Kelvin rather than in degrees Celsius. In Kelvin, the radiative term in equation (9) corresponds to a single library column, T^4 . If the temperature is expressed in Celsius, the same term becomes $(T_C + 273.15)^4$, whose binomial expansion populates every power T_C^0 through T_C^4 . Expressing the temperature in Celcius wouldn't fail SINDy. SINDy would just retrieve a denser system, which is still a dimensional reduction compared to the N initial datapoints of the timer-series signals.

Non-dimensionalisation The state and the control are non-dimensionalised prior to constructing the library. The dimensionless state and control are defined as

$$T^* = \frac{T_K}{T_{\text{ref}}}, \quad u^* = \frac{I}{I_{\text{ref}}}, \quad (12)$$

where T_K is the temperature in Kelvin, I the current, and T_{ref} and I_{ref} the reference scales. After scaling, the coefficients lie within two to three orders of magnitude of one another, and a single threshold becomes meaningful. Substituting equation (12) into equation (9) and dividing by $mc_p T_{\text{ref}}$ yields the dimensionless dynamics

$$\dot{T}^* = \xi_1 + \xi_2 T^* + \xi_3 T^{*4} + \xi_4 u^{*2}, \quad (13)$$

in which the four coefficients are

$$\xi_1 = \frac{hAT_{\infty,K} + \epsilon\sigma AT_{\infty,K}^4}{mc_p T_{\text{ref}}}, \quad \xi_2 = -\frac{hA}{mc_p}, \quad \xi_3 = -\frac{\epsilon\sigma AT_{\text{ref}}^3}{mc_p}, \quad \xi_4 = \frac{RI_{\text{ref}}^2}{mc_p T_{\text{ref}}}, \quad (14)$$

where $T_{\infty,K}$ is the ambient temperature in Kelvin and the remaining symbols are as defined in equation (9) and (10).

The scaling factors T_{ref} and I_{ref} are obtained from a sweep over the whole prior range and are then held fixed. This is when estimating the scaling factors for each time-series separately, the identified coefficients would no longer be a function of the unknown parameters θ alone.

A fourth order component is expected in the radiation term, see equation (14), so the candidate library is built from polynomials in state T^* and control u^* up to order four. A library truncated at a lower order would approximate the radiation by the lower-order terms, which biases ξ_2 .

Sparsity threshold The threshold λ is selected once, by performing a sweep over possible values. The amount of active terms and the relative error between the original and reconstructed signal is recorded for each candidate λ , which exhibits a plateau at four terms.

3.2 Parameter dependence and identifiability

The dependence of the features on the unknown parameters (h and ϵ) can be seen in equation (14) and is summarised in Table 2. The convection coefficient h enters ξ_1 and ξ_2 , the radiation coefficient ϵ enters ξ_1 and ξ_3 , and ξ_4 depends on neither.

Table 2: Analytical dependence of the dimensionless coefficients on the unknown parameters

Feature	h	ϵ	Role
ξ_1	linear	linear, through $T_{\infty,K}^4$	constrains a combination, dominated by h
ξ_2	linear	–	identifies h
ξ_3	–	linear	identifies ϵ
ξ_4	–	–	parameter-independent

Since h and ϵ carry different units, their influence on a given feature is compared by their relative sensitivity

$$S_{\xi_i, \theta_j} = \frac{\theta_j}{\xi_i} \frac{\partial \xi_i}{\partial \theta_j}, \quad (15)$$

which computes the relative change in coefficient ξ_i produced when a change in θ_j would be made. As ξ_2 and ξ_3 depend on h or ϵ solely, their relative sensitivity is both $S_{\xi_2, h} = S_{\xi_3, \epsilon} = 1$. The coefficient of the constant term gives

$$S_{\xi_1, h} = \frac{h}{h + \epsilon \sigma T_{\infty, K}^3}, \quad S_{\xi_1, \epsilon} = \frac{\epsilon \sigma T_{\infty, K}^3}{h + \epsilon \sigma T_{\infty, K}^3}, \quad S_{\xi_1, h} + S_{\xi_1, \epsilon} = 1, \quad (16)$$

so that the relative sensitivity of ξ_1 is partitioned between the convective and radiative heat losses. With $h^* = 8000 \text{ W m}^{-2} \text{ K}^{-1}$, $\epsilon = 0.8$ and $T_{\infty, K} = 298.15 \text{ K}$ the relative sensitivities become

$$S_{\xi_1, h} \approx 99.985\%, \quad S_{\xi_1, \epsilon} \approx 0.015\%. \quad (17)$$

This means that ξ_1 is overwhelmingly more sensitive to variations in h than to variations in ϵ , as can be seen in figure 2 as well.




Figure 2: Influence of the convection coefficient h and emissivity ϵ on the dimensionless library candidate variables 1, T^* , u^{*2} , and T^{*4} .

4 RESULTS

4.1 Verification of the feature map

The sparsity threshold λ is found by a logarithmic sweep over $\lambda \in [10^{-3}, 10^0]$ on a time-series data at the true parameter values (h^* and ϵ^*). As shown in figure 3, the number of active terms decreases from the full library of fifteen candidates, and falls to four over the interval $\lambda \in [0.38, 0.61]$. Meanwhile, the relative fit error remains at 2.8×10^{-4} and is almost identical (2.9×10^{-4}) at four active terms and at fifteen. The threshold is chosen to be $\lambda = 0.5$.

As shown in figure 3, the active-term count rises again to ten for $\lambda > 0.61$ while the relative fit error increases by more than one order of magnitude, to 7×10^{-3} . This is because $\lambda > 0.61$ eliminates one of the four true terms, whose contribution must then be approximated by the lower-order polynomial columns of the library.

Reconstruction of the temperature trajectory from the identified coefficients confirms the accuracy of the feature map at the level of the physical signal. As shown in figure 3, the reconstruction follows the reference solution of equation (9) over the heating, plateau and cooling phases, with a deviation of maximum 0.087 K during the heating ramp at $t = 23$ ms. Therefore the compression from $N = 5000$ samples to four features is achieved without measurable loss of the dynamics.

4.2 Calibration setup

The framework is demonstrated on the resistively heated wire of section 3, in which the convective heat transfer coefficient h and the emissivity ϵ are treated as unknown and collected in $\theta = (h, \epsilon)$. Both parameters are assigned uniform priors, $h \sim \mathcal{U}(6500, 10^4)$ W/(m²K) and $\epsilon \sim \mathcal{U}(0.30, 0.95)$, which are deliberately uninformative. Synthetic observations are generated at




Figure 3: Selection of the sparsity threshold and validation of the identified dynamics. Left: active terms and fit error versus λ . Right: reference and reconstructed temperature curves, applied current pulse, and reconstruction error.

the ground truth $\theta^* = (8000, 0.80)$, adding Gaussian measurement noise of standard deviation 0.25 K to the temperature times-eries, and passing each noisy trajectory through the feature map of section 2.2. This gives $N_{\text{obs}} = 25$ observed feature vectors ξ_{obs} of dimension $n_\xi = 4$, so that the calibration data are the sparse-dynamics coefficients rather than the full temperature signal itself.

4.3 Prior and posterior distributions

The posterior distributions of the convective heat transfer coefficient h and the emissivity ϵ , obtained by SBALC and by MCMC, are shown in figure 4 together with the uniform priors. The posterior means recovered by SBALC, 8010.7 W/(m²K) and 0.79971, deviate from the ground truth by +0.13% and −0.04%, both within one posterior standard deviation. Therefore the compressed feature likelihood is demonstrated to recover the true parameters correctly.




Figure 4: Prior and marginal posterior distributions of h (left) and ϵ (right), obtained with SBALC and MCMC. True values marked h^* , ϵ^* .

4.4 Comparison of the Bayesian solvers

The proposed framework is also compared to other Bayesian solvers. SBALC uses an initial design of four points, an exploration constant $b = 1$, and a convergence ratio on the evidence bounds below 0.1 over two consecutive iterations. The Metropolis-Hastings sampler of UQLab is run with two chains of 2500 steps each, of which the first 500 are discarded as burn-in. TMCMC draws 1000 samples per stage with no additional burn-in.

Table 3: Comparison of SBALC, MCMC and TMCMC on the resistively heated wire benchmark. Posterior moments are averaged over $N_{\text{run}} = 10$ independent runs.

Method	$N_{\mathcal{M}}$	t [s]	h [W/(m ² K)]		ϵ [-]	
			mean	s.d.	mean	s.d.
Ground truth	—	—	8000	—	0.8	—
SBALC	34	24.0	8010.7	18.0	0.79971	0.00588
MCMC (UQLab [16])	5001	130.0	7993.5	17.7	0.80266	0.00576
TMCMC	11000	409.4	8010.8	18.2	0.79960	0.00588

Table 3 lists the posterior mean and standard deviation of each parameter, together with the number of forward model evaluations $N_{\mathcal{M}}$ and the wall-clock time t . The three solvers agree to within one posterior standard deviation on every reported moment. SBALC and TMCMC differ by 0.1 W/(m²K) in the posterior mean of h and by 1.1% in its standard deviation, while MCMC differs from the TMCMC reference by 17.3 W/(m²K) in the mean of h and by 0.0031

in the mean of ϵ . Therefore the surrogate-based posterior is verified against two independent sampling schemes.

At equal accuracy, the cost differs by more than two orders of magnitude. SBALC converges in 34 forward evaluations against 5001 for MCMC and 11000 for TMCMC. This is because the Gaussian process surrogate is placed on the log-likelihood rather than on the model response, so that each expensive evaluation is positioned by the learning function (7) instead the random walk of a sampler. The reduction in wall-clock time is smaller, from 409.4 s to 24.0 s, since here we work with an analytical model that is easy to solve. However, one evaluation of the coupled RSW finite element model can easily take hours to compute, so that $N_{\mathcal{M}}$ rather than t governs the applicability of this framework.

5 CONCLUSIONS AND OUTLOOK

A surrogate-based Bayesian calibration workflow combining SINDYc dimensional reduction with SBALC inference has been developed and demonstrated on a resistively heated wire benchmark, in which the convective heat transfer coefficient h and the emissivity ϵ are inferred from synthetic temperature histories. The main findings are the following.

1. The sparse-dynamics feature map compresses the temperature signal from $N = 5000$ samples to $n_{\xi} = 4$ coefficients without measurable loss of the dynamics. At the selected threshold $\lambda = 0.5$, the relative fit error of the four-term model (2.8×10^{-4}) is almost identical to that of the full fifteen-term library (2.9×10^{-4}), and the trajectory reconstructed from the identified coefficients deviates from the reference solution by maximum 0.087 K. The time-series observations can therefore be dimensionally reduced to four calibration targets, which removes the temporal redundancy that makes the independent-sample likelihood overconfident.
2. The compressed feature likelihood recovers the ground truth. The posterior means obtained with SBALC, 8010.7 W/(m²K) and 0.79971, deviate from the true values by +0.13% and -0.04% respectively, both within one posterior standard deviation. Meanwhile, the three solvers agree to within one posterior standard deviation on every reported moment, which verifies the surrogate-based posterior against two independent sampling schemes.
3. The cost differs by more than two orders of magnitude at equal accuracy. SBALC converged in 34 forward evaluations against 5001 for MCMC and 11000 for TMCMC. It should be noted that the reduction in wall-clock time is smaller, from 409.4 s to 24.0 s, since the benchmark is analytical. However for the coupled RSW FE model, in which a single evaluation can take hours, $N_{\mathcal{M}}$ rather than t governs the applicability of the framework.

On the presented benchmark, the candidate library contains the true dynamics, and the dependence of each coefficient on h and ϵ is known in closed form. Neither condition holds for the dynamic resistance and electrode displacement signals of the RSW FE model, for which the

governing dynamics of the measured quantity are not available a priori. Transferability therefore rests on whether a candidate library can be constructed that yields a stable sparse support across the prior range, and on whether the resulting coefficients remain sensitive to the contact and material parameters of interest. This is the subject of ongoing work.

REFERENCES

- [1] James L. Beck and Siu-Kui Au. Bayesian Updating of Structural Models and Reliability using Markov Chain Monte Carlo Simulation. *Journal of Engineering Mechanics*, 128(4):380–391, April 2002.
- [2] Steven L. Brunton, Joshua L. Proctor, and J. Nathan Kutz. Discovering governing equations from data: Sparse identification of nonlinear dynamical systems. 113(15):3932–3937.
- [3] Steven L. Brunton, Joshua L. Proctor, and J. Nathan Kutz. Sparse Identification of Nonlinear Dynamics with Control (SINDYc)*. *IFAC-PapersOnLine*, 49(18):710–715, January 2016.
- [4] Wayne Cai, Glenn Daehn, Anupam Vivek, Jingjing Li, Haris Khan, Rajiv S. Mishra, and Mageshwari Komarasamy. A state-of-the-art review on solid-state metal joining. *Journal of Manufacturing Science and Engineering, Transactions of the ASME*, 141(3):1–35, 2019. Number: 3.
- [5] Jianye Ching and Yi-Chu Chen. Transitional Markov Chain Monte Carlo Method for Bayesian Model Updating, Model Class Selection, and Model Averaging. *Journal of Engineering Mechanics*, 133(7):816–832, July 2007.
- [6] Sooyong Choi, Sooman Lim, Ali Shan, Jinkyu Lee, Tae Gwang Yun, and Byungil Hwang. Tab-to-Busbar Interconnections in EV Battery Packs: An Introductory Review of Typical Welding Methods. *Micromachines*, 17(1):2, January 2026.
- [7] David Dickinson, Chon-Liang Tsai, and Oussama Jammal. Modeling of Resistance Spot Weld Nugget Growth—Applications for the Automotive Industry. *SAE Transactions*, 99:789–802, 1990.
- [8] S Fukumoto, I Lum, E Biro, D R Boomer, and Y Zhou. Effects of Electrode Degradation on Electrode Life in Resistance Spot Welding of Aluminum Alloy 5182. *WELDING JOURNAL*, 2003.
- [9] M. Hamed and M. Atashparva. A review of electrical contact resistance modeling in resistance spot welding. *Welding in the World*, 61(2):269–290, 2017.
- [10] L. Han, M. Thornton, D. Boomer, and M. Shergold. A correlation study of mechanical strength of resistance spot welding of AA5754 aluminium alloy. *Journal of Materials Processing Technology*, 211(3):513–521, March 2011.

- [11] Marc C. Kennedy and Anthony O'Hagan. Bayesian Calibration of Computer Models. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 63(3):425–464, September 2001.
- [12] B. Le Gloannec, O. Doyen, C. Pouvreau, A. Doghri, and A. Poulon-Quintin. Numerical simulation of resistance upset welding in rod to tube configuration with contact resistance determination. *Journal of Materials Processing Technology*, 238:409–422, December 2016.
- [13] Pei-Pei Li, Chao Dang, Cristóbal H. Acevedo, Marcos A. Valdebenito, and Matthias G. R. Faes. Bayesian model updating via streamlined Bayesian active learning cubature.
- [14] Yong Bing Li, Zhong Qin Lin, Qi Shen, and Xin Min Lai. Numerical analysis of transport phenomena in resistance spot welding process. *Journal of Manufacturing Science and Engineering, Transactions of the ASME*, 133(3):1–8, 2011. Number: 3.
- [15] Tian-Le Lv, Yu-Jun Xia, Siva Prasad Murugan, Fernando Okigami, Hassan Ghassemi-Armaki, Blair E. Carlson, and Yongbing Li. Multi-physical process simulation of resistance spot welding available for synthetic data generation. *Journal of Manufacturing Processes*, 141:709–724, May 2025.
- [16] Stefano Marelli and Bruno Sudret. UQLab: A Framework for Uncertainty Quantification in Matlab. In *Vulnerability, Uncertainty, and Risk: Quantification, Mitigation, and Management - Proceedings of the 2nd International Conference on Vulnerability and Risk Analysis and Management, ICVRAM 2014 and the 6th International Symposium on Uncertainty Modeling and Analysis, ISUMA 2014*, pages 2554–2563, Reston, VA, June 2014. American Society of Civil Engineers.
- [17] H A Nied. The Finite Element Modeling of the Resistance Spot Welding Process. April 1983.
- [18] P. Rogeon, P. Carre, J. Costa, G. Sibilia, and G. Saindrenan. Characterization of electrical contact conditions in spot welding assemblies. *Journal of Materials Processing Technology*, 195(1-3):117–124, 2008. Number: 1-3.
- [19] D. Rubin and D. Rubin. Using the SIR algorithm to simulate posterior distributions. 1988.
- [20] Ellen Simoen, Costas Papadimitriou, and Geert Lombaert. On prediction error correlation in Bayesian model updating. *Journal of Sound and Vibration*, 332(18):4136–4152, September 2013.
- [21] Cameron Summerville, Paul Compston, and Matthew Doolan. A comparison of resistance spot weld quality assessment techniques. *Procedia Manufacturing*, 29:305–312, 2019.

- [22] Zixuan Wan, Hui Ping Wang, Min Wang, Blair E. Carlson, and David R. Sigler. Numerical simulation of resistance spot welding of Al to zinc-coated steel with improved representation of contact interactions. *International Journal of Heat and Mass Transfer*, 101:749–763, 2016.
- [23] Jin Wang, Hui Ping Wang, Fenggui Lu, Blair E. Carlson, and David R. Sigler. Analysis of Al-steel resistance spot welding process by developing a fully coupled multi-physics simulation model. *International Journal of Heat and Mass Transfer*, 89:1061–1072, 2015.
- [24] Yu-Jun Xia, Tian-Le Lv, Hassan Ghassemi-Armaki, Yong-Bing Li, and Blair E. Carlson. Collaborative Simulation of Nugget Growth and Process Signals for Resistance Spot Welding. *Welding in the World*, 67(6):1377–1392, June 2023. Company: Springer Distributor: Springer Institution: Springer Label: Springer Number: 6.
- [25] Yu-Jun Xia, Tian-Le Lv, Yong-Bing Li, Hassan Ghassemi-Armaki, and Blair E. Carlson. Quantitative Interpretation of Dynamic Resistance Signal in Resistance Spot Welding. *Welding Journal*, 102(4):69–87, April 2023.
- [26] Yansong Zhang, Jie Shen, and Xinmin Lai. Influence of Electrode Force on Weld Expulsion in Resistance Spot Welding of Dual Phase Steel with Initial Gap Using Simulation and Experimental Method. *ISIJ International*, 52(3):493–498, 2012.
- [27] Dawei Zhao, Yuanxun Wang, Peng Zhang, and Dongjie Liang. Modeling and Experimental Research on Resistance Spot Welded Joints for Dual-Phase Steel. *Materials*, 12(7):1108, January 2019. Number: 7.
- [28] Baifan Zhou, Tim Pychynski, Markus Reischl, Evgeny Kharlamov, and Ralf Mikut. Machine learning with domain knowledge for predictive quality monitoring in resistance spot welding. *Journal of Intelligent Manufacturing*, 33(4):1139–1163, April 2022. Number: 4.