



Statistical finite elements for misspecified models

Connor Duffin^{a,1}, Edward Cripps^a, Thomas Stemler^{a,b}, and Mark Girolami^{c,d}

^aDepartment of Mathematics and Statistics, The University of Western Australia, Perth, WA 6009, Australia; ^bComplex Systems Group, The University of Western Australia, Perth, WA 6009, Australia; ^cDepartment of Engineering, University of Cambridge, Cambridge CB2 1PZ, United Kingdom; and ^dLloyd's Register Foundation Programme for Data-Centric Engineering, The Alan Turing Institute, London NW1 2DB, United Kingdom

Edited by Nancy M. Reid, University of Toronto, Toronto, ON, Canada, and approved November 18, 2020 (received for review July 23, 2020)

We present a statistical finite element method for nonlinear, time-dependent phenomena, illustrated in the context of nonlinear internal waves (solitons). We take a Bayesian approach and leverage the finite element method to cast the statistical problem as a nonlinear Gaussian state-space model, updating the solution, in receipt of data, in a filtering framework. The method is applicable to problems across science and engineering for which finite element methods are appropriate. The Korteweg–de Vries equation for solitons is presented because it reflects the necessary complexity while being suitably familiar and succinct for pedagogical purposes. We present two algorithms to implement this method, based on the extended and ensemble Kalman filters, and demonstrate effectiveness with a simulation study and a case study with experimental data. The generality of our approach is demonstrated in SI Appendix, where we present examples from additional nonlinear, time-dependent partial differential equations (Burgers equation, Kuramoto–Sivashinsky equation).

Bayesian calibration | finite element methods | model discrepancy

The central role of physically derived, nonlinear, time-dependent partial differential equations (PDEs) in scientific and engineering research is undisputed, as is the need for numerical intervention in order to understand their behavior. The finite element method (FEM) has emerged as the foremost strategy to undergo this numerical intervention, yet when these discretized solutions are compared with empirical evidence, elements of model mismatch are revealed that require statistical formalisms to be dealt with appropriately (1–3). To address this problem of model misspecification, in this paper we introduce stochastic forcing inside the PDE and update the FEM discretized PDE solution with data in a filtering context.

Stochastic forcing is introduced through a random function within the governing equations. This represents an unknown process, which may have been omitted in the formulation of the physical model. For an elliptic linear PDE with coefficients Λ , this can be expressed as

$$\begin{cases} L_{\Lambda} u = f + \xi_{\theta}, & \xi_{\theta} \sim \mathcal{GP}(0, C_{\theta}), \\ u := u(x), f := f(x), x \in \Omega \subset \mathbb{R}^d. \end{cases}$$

The push forward of the Gaussian random field ξ_{θ} , with covariance parameters θ , induces a probability measure over the space of admissible solutions to the above. To embed this into a finite element model, we start with the weak form

$$\mathcal{A}_{\Lambda}(u, v) = \langle f, v \rangle + \langle \xi_{\theta}, v \rangle,$$

where $\mathcal{A}_{\Lambda}(\cdot, \cdot)$ is the bilinear form generated from L_{Λ} and $\langle \cdot, \cdot \rangle$ is the appropriate Hilbert space inner product. Discretizing with finite elements $u(x) = \sum_{i=1}^M u_i \phi_i(x)$, $v(x) = \sum_{i=1}^M v_i \varphi_i(x)$ yields the Gaussian measure over the solution FEM coefficients $\mathbf{u} = (u_1, u_2, \dots, u_M) \in \mathbb{R}^M$:

$$p(\mathbf{u} | \Lambda, \theta) = \mathcal{N}(\mathbf{A}^{-1} \mathbf{b}, \mathbf{A}^{-1} \mathbf{G}(\theta) \mathbf{A}^{-T}),$$

where $\mathbf{A}_{ij} = \mathcal{A}_{\Lambda}(\phi_i, \varphi_j)$, $\mathbf{b}_i = \langle f, \varphi_i \rangle$, and $\mathbf{G}(\theta)_{ij} = \langle \phi_i, C_{\theta} \varphi_j \rangle$. This defines a (finite-dimensional) prior distribution over the

FEM model, which represents all assumed knowledge before observing data. The mean is the standard Galerkin solution, and the covariance results from the action of the discretized PDE operator on the covariance $\mathbf{G}(\theta)$; further details are contained in SI Appendix, section 1. This was first developed in ref. 4, and we demonstrate the generality of such an approach by extending it to nonlinear, time-dependent PDEs.

An area in which nonlinear and time-dependent problems are ubiquitous is ocean dynamic processes, where essentially all problems stem from a governing system of nonlinear, time-dependent equations (e.g., the Navier–Stokes equations). The ocean dynamics community has grown increasingly cognizant of the importance of accurate uncertainty quantification (5, 6), with many possible applications [e.g., rogue waves (7), turbulent flow (8)] for our proposed methodology.

An example process is nonlinear internal waves (solitons), which are observed as waves of depression or elevation along a pycnocline in a density-stratified fluid and are of broad interest to both the scientific and engineering communities (9–13). The classical mathematical model for solitons is the Korteweg–de Vries (KdV) equation (14):

$$\begin{cases} u_t + \alpha u u_x + \beta u_{xxx} + c u_x = 0, \\ u := u(x, t), x \in [0, L], t \in [0, T], \end{cases} \quad [1]$$

where u is the pycnocline displacement. Coefficients α , β , and c are determined by physical parameters. Eq. 1 is readily interpretable: waves propagate at wave speed c , nonlinear steepening results from $u u_x$, and dispersion is due to u_{xxx} . Relative coefficient values determine the dominating regime, and waves can vary from quasilinear to highly nonlinear.

Significance

Science and engineering have benefited greatly from the ability of finite element methods (FEMs) to simulate nonlinear, time-dependent complex systems. The recent advent of extensive data collection from such complex systems now raises the question of how to systematically incorporate these data into finite element models, consistently updating the solution in the face of mathematical model misspecification with physical reality. This article describes general and widely applicable methodology for the coherent synthesis of data with FEM models, providing a data-driven probability distribution that captures all sources of uncertainty in the pairing of FEM with measurements.

Author contributions: C.D., E.C., T.S., and M.G. designed research, performed research, analyzed data, and wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

This open access article is distributed under Creative Commons Attribution License 4.0 (CC BY).

¹To whom correspondence may be addressed. Email: connor.duffin@research.uwa.edu.au.

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2015006118/-DCSupplemental>.

Published December 28, 2020.

Despite KdV being well studied (15) and widely applied (16–18), its relative simplicity makes it prone to model mismatch. To compensate for this mismatch, we update the FEM discretized solution with observations in a filtering context. The resulting statistical FEM (statFEM) is shown using simulated and experimental data to

- 1) Approximate the data-generating process with a statistically coherent uncertainty quantification.
- 2) Synthesize physics and data to give an interpretable posterior distribution.
- 3) Utilize sparsely observed data to reconstruct observed phenomena.
- 4) Enable the application of simpler physical models, updated with observations.

For practitioners faced with data, we believe these benefits are of importance, and we demonstrate the generality of our method in [SI Appendix](#) with further examples. Code to replicate the analysis is freely available online.*

A Nonlinear, Time-Evolving Statistical FEM

A Gaussian process (GP), ξ_θ , is introduced inside of the governing equations, which represent an unknown forcing process in space and time, with time-varying parameters θ . For a general nonlinear PDE, this is given by

$$u_t + L_\Lambda u + F_\Lambda(u) + \xi_\theta = 0, \quad \xi_\theta \sim \mathcal{GP}(0, C_\theta), \quad [2]$$

where L_Λ and F_Λ represent linear and nonlinear differential operators, respectively, with coefficients Λ . The push forward of the Gaussian measure ξ_θ induces a probability measure over the space of admissible solutions to Eq. 2 and characterizes our prior belief in the model based on modeling assumptions. The kernel of the covariance operator C_θ is given by

$$\mathbb{E} [\xi_\theta(x, t)\xi_\theta(x', t')] = k_\theta(x, x') \cdot \delta(t, t').$$

The exact form of this covariance can be decided upon by domain experts so that the uncertainty induced is physically motivated. For example, k_θ can be chosen to be a Matérn covariance function to reflect the unknown forcing having derivatives up to a known order. We assume a white noise process in time to facilitate the application of standard Kalman methods to solve the filtering problem; this is also convention in stochastic differential equations (19). When computing the prior defined by Eq. 2, we use fixed parameters θ . When conditioning on data, we take an empirical Bayes approach and estimate θ through the log-marginal posterior.

Coefficients of Eq. 2 are assumed to be known, and we choose to update the numerical solution to the model, acknowledging that estimating Λ [using, e.g., maximum likelihood methods (20), Markov chain Monte Carlo (21), or inversion methods in general (22)] is also of utmost interest.

Discretizing Eq. 2 using finite elements in space with an implicit or explicit Euler method in time,[†] denote by $\mathbf{u}^n \in \mathbb{R}^M$ the FEM coefficients at time $n\Delta t$, for time step Δt . Further analysis of this discretization will be the focus of future work. For a potentially nonlinear system of equations \mathcal{F}_Λ , the resultant system can be expressed as (full construction is given in [SI Appendix, sections 2 and 3](#))

$$\mathcal{F}_\Lambda(\mathbf{u}^n, \mathbf{u}^{n-1}) + \mathbf{e}_{n-1} = 0.$$

The vector \mathbf{e}_{n-1} represents Galerkin discretized increments of a Brownian motion process, $\mathbf{e}_{n-1} \sim \mathcal{N}(\mathbf{0}, \Delta_t \mathbf{G}(\theta_n))$.

Unlike the elliptic example in the Introduction, for Eq. 2 the induced probability measure on the FEM coefficients is not available in closed form, and we present two approximations, based on the extended Kalman filter (EKF) and the ensemble Kalman filter (EnKF). The first linearizes about the current solution with the Jacobian \mathbf{J} of the nonlinear \mathcal{F}_Λ (evaluated at the current solution) to give a Gaussian approximation of $p(\mathbf{u}^n | \boldsymbol{\theta}_{1:n}, \Lambda) \approx \mathcal{N}(\mathbf{m}_n, \mathbf{C}_n)$. The second uses an ensemble in which a perturbed system is solved, with realizations from \mathbf{e}_n , at each time step. Summary statistics (e.g., mean, covariance) are then computed from this ensemble.

For the prior, the deterministic FEM solution is identically equal to the mean in the EKF approach. However we have found in numerical experiments that the EKF method, due to the use of the Jacobian, inflates the covariance at points of high gradient with reduction at points approaching near-zero gradient, when using large time steps. This does not occur with the EnKF approach.

Algorithm 1: EKF algorithm

```

for  $n \leq n_t$  do
  (Prediction step)
  Solve  $\mathcal{F}(\mathbf{m}_{n|n-1}, \mathbf{m}_{n-1|n-1}) = \mathbf{0}$ .
   $\hat{\mathbf{C}}_{n|n-1} = (\mathbf{J}^{n-1})^{-1} (\mathbf{J}^{n-1} \mathbf{C}_{n-1|n-1} (\mathbf{J}^{n-1})^\top) m b \mathbf{J}^n)^{-\top}$ 
  Estimate:
  arg max $_{\theta_n, \sigma_n} \{ \log p(\mathbf{y}_n | \mathbf{y}_{1:n-1}, \theta_{1:n}, \sigma_{1:n})$ 
                     $+ \log p(\theta_n) + \log p(\sigma_n) \}$ .
   $\mathbf{C}_{n|n-1} = \hat{\mathbf{C}}_{n|n-1} + \Delta_t (\mathbf{J}^n)^{-1} \mathbf{G}(\theta_n) \mathbf{J}^n)^{-\top}$ .
   $\mathbf{S}_n = \mathbf{H}_n \mathbf{C}_{n|n-1} \mathbf{H}_n^\top + \sigma_n^2 \mathbf{I}$ .
  (Analysis step)
   $\mathbf{m}_{n|n} = \mathbf{m}_{n|n-1} + \mathbf{C}_{n|n-1} \mathbf{H}_n^\top \mathbf{S}_n^{-1} (\mathbf{y}_n - \mathbf{H}_n \mathbf{m}_{n|n-1})$ .
   $\mathbf{C}_{n|n} = \mathbf{C}_{n|n-1} - \mathbf{C}_{n|n-1} \mathbf{H}_n^\top \mathbf{S}_n^{-1} \mathbf{H}_n \mathbf{C}_{n|n-1}$ .

```

Conditioning on Data. Data $\mathbf{y}_n \in \mathbb{R}^N$ are observed at time $n\Delta t$ on the grid \mathbf{x}_{obs} . These data are corrupted with noise $\boldsymbol{\eta}_n \sim \mathcal{N}(\mathbf{0}, \sigma_n^2 \mathbf{I})$ independent to the model \mathbf{u}^n to give the data-generating process $\mathbf{y}_n = \mathbf{H}_n \mathbf{u}^n + \boldsymbol{\eta}_n$, where the linear observation operator $\mathbf{H}_n: \mathbb{R}^M \rightarrow \mathbb{R}^N$ maps from the computed solution grid to the observation grid, using the FEM interpolant.

The filtered distribution $p(\mathbf{u}^n | \mathbf{y}_{1:n}, \boldsymbol{\theta}_{1:n}, \sigma_{1:n}, \Lambda)$, where $\mathbf{y}_{1:n} = (\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_n)$, is our primary object of interest.[‡] We take a Bayesian interpretation and refer to this as the filtered posterior distribution or just the posterior, when the context is clear. However, as this is a filtering problem, non-Bayesian methods are perfectly valid. We assume that all distributions are Gaussian, so the posterior can be computed with standard methods in data assimilation (23); we use the EKF (24) and the EnKF (stochastic form) (25).

The initial conditions are known (i.e., they are given a Dirac measure). For time n , we make a tentative prediction step according to the PDE model \mathcal{F}_Λ , propagating uncertainty in the previous time step [described by $p(\mathbf{u}^{n-1} | \mathbf{y}_{1:n-1}, \boldsymbol{\theta}_{1:n-1}, \sigma_{1:n-1})$], to give the prediction measure $p(\mathbf{u}^n | \mathbf{y}_{1:n-1}, \boldsymbol{\theta}_{1:n-1}, \sigma_{1:n-1})$. Parameters (θ_n, σ_n) are then estimated, and the full prediction step is completed to estimate $p(\mathbf{u}^n | \mathbf{y}_{1:n-1}, \boldsymbol{\theta}_{1:n}, \sigma_{1:n})$. Data observed at time $n\Delta t$ are then conditioned on to give the updated filtering distribution $p(\mathbf{u}^n | \mathbf{y}_{1:n}, \boldsymbol{\theta}_{1:n}, \sigma_{1:n})$.

We assume the parameters are independent across time [i.e., $p(\theta_n | \theta_{n-1}) = p(\theta_n)$]. Parameters may also be time constant, which is discussed in the maximum likelihood setting in ref. 26 and in the hierarchical Bayesian setting in ref. 27. [SI Appendix](#)

*It is available at <https://github.com/connor-duffin/statkdv-paper>.

[†]Crank–Nicolson may also be used to ensure stability.

[‡]From here on in, we implicitly condition on PDE coefficients Λ .




Fig. 1. Simulation study results (using EnKF). **A** shows a prior solution with fixed hyperparameters θ_n with the data-generating process (DGP; as labeled) for two times with accompanying 95% credible intervals. Note the degree of model mismatch between the two profiles. **B** shows the DGP and the posterior mean and 95% credible intervals for the same times as the prior. Given the data, the model is highly certain about the updated mean, for which the model mismatch has been corrected. This is now using estimated hyperparameters θ_n for a fully data-driven approach. **C** shows the prior mean, DGP, and posterior mean over the entire simulation grid in space and time, and **D** shows the estimated parameters θ_n , σ_n across time, with the true value for σ_n shown as a dashed turquoise line.

also contains a possible modification of the method that accounts for time-invariant parameters. The following procedure provides an overview of the method, and *Algorithms 1* and *2* give pseudocode versions.

Conditioning procedure. At time n , assume that the measure on the previous time is described by

$$p(\mathbf{u}^{n-1} | \mathbf{y}_{1:n-1}, \boldsymbol{\theta}_{1:n-1}, \sigma_{1:n-1}) = \mathcal{N}(\mathbf{m}_{n-1|n-1}, \mathbf{C}_{n-1|n-1}).$$

Then, proceed as follows:

- 1) Compute the tentative prediction step:

$$\begin{aligned} & p(\mathbf{u}^n | \mathbf{y}_{1:n-1}, \boldsymbol{\theta}_{1:n-1}, \sigma_{1:n-1}) \\ &= \int p(\mathbf{u}^n | \mathbf{u}^{n-1}) p(d\mathbf{u}^{n-1} | \mathbf{y}_{1:n-1}, \boldsymbol{\theta}_{1:n-1}, \sigma_{1:n-1}) \\ &\approx \mathcal{N}(\hat{\mathbf{m}}_{n|n-1}, \hat{\mathbf{C}}_{n|n-1}). \end{aligned}$$

- 2) Maximize the EKF log-marginal posterior to estimate parameters:

$$\begin{aligned} & \arg \max_{\boldsymbol{\theta}_n, \sigma_n} \{ \log p(\mathbf{y}_n | \mathbf{y}_{1:n-1}, \boldsymbol{\theta}_{1:n}, \sigma_{1:n}) \\ & \quad + \log p(\boldsymbol{\theta}_n) + \log p(\sigma_n) \}, \end{aligned}$$

where $\hat{\mathbf{G}}(\boldsymbol{\theta}_n) = \Delta_t (\mathbf{J}^n)^{-1} \mathbf{G}(\boldsymbol{\theta}_n) (\mathbf{J}^n)^{-\top}$,

$$\begin{aligned} & p(\mathbf{y}_n | \mathbf{y}_{1:n-1}, \boldsymbol{\theta}_{1:n}, \sigma_{1:n}) \\ &= \mathcal{N}(\mathbf{H}_n \hat{\mathbf{m}}_{n|n-1}, \mathbf{H}_n \hat{\mathbf{C}}_{n|n-1} \mathbf{H}_n^\top + \mathbf{H}_n \hat{\mathbf{G}}(\boldsymbol{\theta}_n) \mathbf{H}_n^\top + \sigma_n^2 \mathbf{I}). \end{aligned}$$

- 3) Compute the full prediction step:

$$\begin{aligned} & p(\mathbf{u}^n | \mathbf{y}_{1:n-1}, \boldsymbol{\theta}_{1:n}, \sigma_{1:n-1}) \\ &= \int p(\mathbf{u}^n | \mathbf{u}^{n-1}, \boldsymbol{\theta}_n) p(d\mathbf{u}^{n-1} | \mathbf{y}_{1:n-1}, \boldsymbol{\theta}_{1:n-1}, \sigma_{1:n-1}) \\ &\approx \mathcal{N}(\mathbf{m}_{n|n-1}, \mathbf{C}_{n|n-1}). \end{aligned}$$

- 4) Complete the analysis step:

$$\begin{aligned} & p(\mathbf{u}^n | \mathbf{y}_{1:n}, \boldsymbol{\theta}_{1:n}, \sigma_{1:n}) \\ &\propto p(\mathbf{y}_n | \mathbf{u}^n, \sigma_n) p(\mathbf{u}^n | \mathbf{y}_{1:n-1}, \boldsymbol{\theta}_{1:n}, \sigma_{1:n-1}) \\ &= \mathcal{N}(\mathbf{m}_{n|n}, \mathbf{C}_{n|n}). \end{aligned}$$

The prior $p(\mathbf{u}^n | \boldsymbol{\theta}_n, \Lambda)$ is recovered if only the full prediction step (step 3) is completed at each iteration; completing the full sequence gives the posterior. Optimization of the log-marginal posterior is done using the limited-memory Broyden–Fletcher–Goldfarb–Shanno algorithm (28) with starting points set to the previous estimates. This log-marginal posterior is calculable due to the Gaussian assumption made in step 1. Prior information on hyperparameters is incorporated through $p(\boldsymbol{\theta}_n)$ and $p(\sigma_n)$, which regularizes the optimization problem.

Simulation Study. We condition on data generated from an extended Korteweg–de Vries (eKdV) equation with a cubic nonlinear term:

$$w_t + \alpha w w_x + \beta w_{xx} + \varepsilon w^3 w_x = 0,$$

Algorithm 2: EnKF algorithm

```

for  $n \leq n_t$  do
  (Prediction step)
  for  $i \leq N_{ens}$  do
    Solve  $\mathcal{F}(\mathbf{u}_{pred}^{n,[i]}, \mathbf{u}^{n-1,[i]}) = 0$ .
  Compute  $\mathbf{m}_{n|n-1}$  and  $\hat{\mathbf{C}}_{n|n-1}$  from  $\{\mathbf{u}_{pred}^{n,[i]}\}$ 
  Estimate:
  arg max $_{\theta_n, \sigma_n} \{ \log p(\mathbf{y}_n | \mathbf{y}_{1:n-1}, \theta_{1:n}, \sigma_{1:n}) + \log p(\theta_n) + \log p(\sigma_n) \}$ 
  for  $i \leq N_{ens}$  do
    Solve  $\mathcal{F}(\mathbf{u}_{pred}^{n,[i]}, \mathbf{u}^{n-1,[i]}) + \mathbf{e}_{n-1}^{[i]} = 0$ .
  Compute  $\mathbf{m}_{n|n-1}$  and  $\hat{\mathbf{C}}_{n|n-1}$  from  $\{\mathbf{u}_{pred}^{n,[i]}\}$ 
   $\mathbf{S}_n = \mathbf{H}_n \mathbf{C}_{n|n-1} \mathbf{H}_n^\top + \sigma_n^2 \mathbf{I}$ .
  (Analysis step)
  for  $i \leq N_{ens}$  do
     $\mathbf{u}^{n,[i]} = \mathbf{u}_{pred}^{n,[i]} + \mathbf{C}_{n|n-1} \mathbf{H}_n^\top \mathbf{S}_n^{-1} (\mathbf{y}_n + \boldsymbol{\eta}_n^{[i]} - \mathbf{H}_n \mathbf{u}_{pred}^{n,[i]})$ 

```

setting $\alpha = 1$, $\beta = 0.01$, and $\varepsilon = 20$. The misspecified KdV model

$$u_t + \alpha u u_x + \beta u_{xxx} + \xi_\theta = 0 \tag{3}$$

has the same coefficient values, with initial conditions set to a wave of depression $u(x, 0) = -0.3 \operatorname{sech}^2(x - 15)$ on the space-time grid $(x, t) \in [0, 20] \times [0, 50]$. Boundary conditions are periodic. Gaussian random forcing ξ_θ has spatial covariance $k_\theta(x, x') = \tau^2 \exp(-\|x - x'\|^2 / (2\ell^2))$ (we refer to τ as the scale parameter and ℓ as the length parameter).

KdV is discretized following ref. 29, using P_1 trial functions and P_0 testing functions, with a Crank–Nicolson method in time. The data-generating process is simulated using Dedalus (30) with 1,024 grid points in space. This is then down sampled to 20 grid points and jittered with synthetic Gaussian observational noise (mean 0, variance 0.001^2) to give the simulated dataset.

We assume $\mathbf{y}_n = \mathbf{H}\mathbf{u}^n + \boldsymbol{\eta}_n$, where \mathbf{u}^n is the Galerkin discretized solution to Eq. 3 and $\boldsymbol{\eta}_n \sim \mathcal{N}(0, \sigma_n^2 \mathbf{I})$. Hyperparameters $\theta_n = (\tau_n, \ell_n)$ and noise level σ_n are estimated at each step by maximizing the log-marginal posterior, with the weakly informative truncated Gaussian priors $\tau_n \sim \mathcal{N}_+(1, 1^2)$, $\ell_n \sim \mathcal{N}_+(1, 1^2)$, and $\sigma_n \sim \mathcal{N}_+(0, 1^2)$. The EnKF is used in this section, with $n_x = 400$, $n_t = 2,001$, and $N_{ens} = 400$. Results are presented in Fig. 1.

For a fixed set of hyperparameters $\tau_n = 0.0025$, $\ell_n = 1$ for all n , the data-generating process and estimated prior are shown in Fig. 1A, appearing visually mismatched in the mean via phase shift, increased oscillations, and increased wave interactions. Note that the stochastic forcing induces an uncertainty about the PDE solution, represented by the 95% credible intervals shown. Note also that the data-generating process is approximately contained within the credible intervals.

Fig. 1B shows that the posterior mean approximates the data-generating process, and the posterior uncertainty bounds have shrunk as a result of conditioning, indicating high certainty about the posterior mean values. Model discrepancy between the data and the statFEM solution has been corrected for. The space-time view of the posterior, shown in Fig. 1C, shows that the posterior has incorporated the complex soliton interactions in the data, not present in the prior.

Parameter estimates (Fig. 1D) indicate that the length and noise parameters are both stable, with the noise being slightly overestimated (i.e., $\sigma_n \approx 0.003 > 0.001$). Times at which the noise is not identified result in it being set to the lower bound. The scale parameter quantifies the accuracy of the model prediction step at each time step. In this case, model predictions vary

in their accuracy and appear approximately bounded to within $(10^{-5}, 10^{-1})$.

Case Study: Experimental Data

We now apply the method to the experimental data collected in ref. 31. Experiments were conducted to study weakly nonlinear models for internal waves in lakes and consisted of generating internal waves in a two-layer stratified system, inside of a clear acrylic tank of dimensions $6 \times 0.3 \times 0.29$ m. The tank contained an upper layer of fresh water and a lower layer of saline water, with a density gradient of $\Delta\rho = 20 \text{ kgm}^{-3}$. The tank was able to rotate in order to establish the initial conditions, which were an inclined plane of angle $\vartheta = 0.5^\circ$. This initial condition mimics the shear induced by strong winds in lakes. At time $t = 0$, the tank is rotated to restore it to the horizontal.

Data were recorded at three spatially equidistant locations in the tank using ultrasonic wave gauges (Fig. 2), taking measurements approximately every 0.01 s, up to $T = 1,000$ s; we use data up to $T = 300$ s. Data are measured in voltages and are post-processed to give pycnocline displacements in meters. These data are plotted in Fig. 3, where the small measurement error is visually apparent. Transient behavior is observed before steepening, and a soliton wave train forms; three such steepening events are observed in the data we analyze. As $T \rightarrow 1,000$ s, dissipation results in the wave profile approaching a flat steady-state profile.

Our physical model is an eKdV equation with— for computational simplicity—a linear dissipation term. We acknowledge that for laminar boundaries, other methods are preferred (31, 32). Including some form of dissipation is important as otherwise, the model becomes impractically mismatched by the end of the simulation. The eKdV is given by

$$u_t + \alpha u u_x + \beta u_{xxx} + c u_x + \nu u + \xi_\theta = 0 \tag{4}$$

for $u := u(x, t)$, $x \in [0, L]$, $t \in [0, T]$, and with coefficients

$$\alpha = \frac{3}{2} \frac{c(h_1 - h_2)}{h_1 h_2}, \beta = \frac{c h_1 h_2}{6}, c = \sqrt{\frac{g' h_1 h_2}{H}}, g' = \frac{\Delta\rho g}{\rho_0}.$$

We set ξ_θ as a GP as described previously, with spatial covariance kernel k_θ set to a squared exponential with scale and length hyperparameters τ and ℓ . For the experiment under consideration, we have $h_1 = 0.232$ m, $h_2 = 0.058$ m, $H = h_1 + h_2 = 0.29$ m, and $\rho_0 = 1,000 \text{ kgm}^{-3}$. The dissipation coefficient ν is an inverse timescale, which is set to $3 \times 10^{-3} \text{ s}^{-1}$.

Incorporation of reflective boundary conditions is done by solving the eKdV equation across the extended domain $[0, 2L]$ with periodic boundary conditions and summing solutions in the (reflected) subdomains $[0, L]$, $[L, 2L]$:

$$u_{\text{exact}}(x, t) := u(x, t) + u(2L - x, t), \quad x \in [0, L].$$

Details on the derivation are in ref. 31. Solutions to the deterministic version of Eq. 4 at the locations of the wave gauges are




Fig. 2. Schematic diagram of the experimental apparatus. Wave gauges (WGs) are labeled WG1, WG2, and WG3, and the initial conditions are shown as a gray line, labeled with initial angle ϑ .




Fig. 3. Experimental data results; posterior computed by the EnKF. **A** shows the observed data, the deterministic solution to the misspecified model, and the posterior mean at the wave gauges with 95% credible intervals up to time $T = 300$ s. Model mismatch between the model and data is visually apparent and has been offset in the posterior mean. **B** shows the posterior mean wave profile inside of the tank, with 95% credible intervals and the data, for three times. The wave profile has been reconstructed by conditioning on the observed data. **C** shows the estimated hyperparameters θ_n , and **D** shows the estimated posterior mean over the space–time grid.

shown in Fig. 3. We show the deterministic solution instead of the prior due to accumulation of errors for large simulation times.

The eKdV model does not capture the observed behavior exactly. The model waves have higher velocity than the observations, and model amplitudes are slightly larger than observed amplitudes. It is conjectured that this is due to misspecification of dissipation, but in any case, the model is misspecified. Rather than estimating the eKdV parameters using inversion techniques, we sequentially update the model with observations to give the posterior $p(\mathbf{u}^n | \mathbf{y}_{1:n}, \theta_{1:n}, \sigma_{1:n})$.

As before, we assume $\mathbf{y}_n = \mathbf{H}\mathbf{u}^n + \boldsymbol{\eta}_n$ with known noise $\boldsymbol{\eta}_n \sim \mathcal{N}(0, 1.3588 \times 10^{-8}\mathbf{I})$. As we solve on the extended domain $[0, 2L]$ and sum solutions, the observation operator \mathbf{H} is taken to be the sum of the appropriate function values given by our FEM interpolant (a linear operation). The observation points are unchanging, as is the solution mesh, so \mathbf{H} is constant in time. The hyperparameters of ξ_θ , $\theta_n = (\tau_n, \ell_n)$, must be estimated at each iteration by maximizing the log-marginal posterior. Due to small data in space (three observations each time step), we use a projection method to estimate hyperparameters. This linearly projects the predicted mean $\mathbf{m}_{n|n-1}$ forward, estimated from the data points: $y_i^n = a_n + b_n \mathbf{m}_{n|n-1}(x_i)$. Parameters a_n, b_n are estimated to give the best least squares linear projection from the prediction to the data. This gives a projected dataset, $\tilde{\mathbf{y}}_n$, using the linear shift: $\tilde{\mathbf{y}}_n = a_n + b_n \mathbf{m}_{n|n-1}$. The estimated hyperparameters are then given by $\arg \max_{\theta_n} \{\log p(\tilde{\mathbf{y}}_n | \mathbf{y}_{1:n-1}, \theta_{1:n}, \sigma_{1:n}) + \log p(\theta_n)\}$, in which the observed data \mathbf{y}_n are replaced with the projected data $\tilde{\mathbf{y}}_n$. We project to a grid of 100 points uniformly spaced across the solution grid. Note that this is only for the parameter estimation step, and we do not use this $\tilde{\mathbf{y}}_n$ as the data in the analysis step.

We set weakly informative priors: $\tau_n \sim \mathcal{N}_+(1, 1^2)$ and $\ell_n \sim \mathcal{N}_+(1, 1^2)$. The posterior is computed using the ensemble

method with $n_x = 200$, $n_t = 1,001$, and $N_{ens} = 2,048$. Results are shown in Fig. 3. The posterior mean values of $p(\mathbf{u}^n | \mathbf{y}_{1:n}, \theta_{1:n}, \sigma_{1:n})$ at the wave gauges are shown in Fig. 3A and offer a close fit to the data in comparison with the eKdV solution. The credible intervals shrink about the data (compare Fig. 3B) and are not seen on the figure.

Posterior wave profiles are shown in Fig. 3B and demonstrate that given the data, the method is able to yield a sensible estimate for the underlying wave profile and is hence able to reconstruct the wave profile given sparse observations in space (Fig. 3D). Furthermore the provided uncertainty quantification is physically sensible, with bounds contracting about the data and expanding near the boundaries.

The hyperparameters, θ_n , are shown in Fig. 3C. The scale parameter is seen to vary between two distinct levels, indicating that the model predictions vary in their accuracy. The amplitude of these mismatch scales shows that in this case, model mismatch is a cumulative effect that takes some time before it is obviously occurring (Fig. 3A, eKdV solution). Repeated conditioning on data helps to mitigate these long timescale effects due to continual updating. A space–time view of the posterior mean wave profile is shown in Fig. 3D, demonstrating that the general behavior of the flow (e.g., reflective boundary conditions, dissipation, wave train formation) is indeed captured.

Conclusions

We present a data-driven approach to the FEM that assimilates observations into nonlinear, time-dependent PDEs by embedding model misspecification uncertainty in the governing equations and sequentially updating the discretized equations with observations in a filtering context. Examples presented using the KdV equation (and the additional systems studied in *SI Appendix*) demonstrate that the method can approximate

the data-generating process to give an interpretable posterior distribution, which reconstructs the observed phenomena.

This work sets the foundation for future studies of embedding data within FEM models. The use of the underlying Kalman framework permits the drawing upon of ideas from high-dimensional data assimilation, which for the systems studied here, were not needed. Techniques from Bayesian inversion can also be used to provide uncertainty quantification for physical quantities of interest, which will also allow for more accurate prediction. Finally, we believe that the development of similar methodology for alternate discretizations (e.g., spectral methods) could also be of great benefit, allowing for even broader application.

Data Availability. Data and code have been deposited on GitHub (<https://github.com/connor-duffin/statkdv-paper>).

ACKNOWLEDGMENTS. We thank the two anonymous referees whose suggestions greatly improved the manuscript. C.D. was supported by a Bruce and Betty Green Postgraduate Research Scholarship and an Australian Government Research Training Program Scholarship at the University of Western Australia. C.D. and E.C. were supported by Australian Research Council Industrial Transformation Research Hub Grant IH140100012. E.C. was supported by Australian Research Council Industrial Transformation Training Centre Grant IC190100031. M.G. was supported by Engineering and Physical Sciences Research Council Grants EP/R034710/1, EP/R018413/1, EP/R004889/1, and EP/P020720/1 and a Royal Academy of Engineering Research Chair.

1. M. C. Kennedy, A. O'Hagan, Bayesian calibration of computer models. *J. Roy. Stat. Soc. B* **63**, 425–464 (2001).
2. K. Judd, L. A. Smith, Indistinguishable states II: The imperfect model scenario. *Phys. Nonlinear Phenom.* **196**, 224–242 (2004).
3. J. O. Berger, L. A. Smith, On the statistical formalism of uncertainty quantification. *Annu Rev Stat Appl.* **6**, 433–460 (2019).
4. M. Girolami, E. Febrianto, G. Yin, F. Cirak, The statistical finite element method (stat-FEM) for coherent synthesis of observation data and model predictions. *Comput. Methods Appl. Mech. Eng.*, 10.17863/CAM.59639 (2020).
5. P. F. J. Lermusiaux, Uncertainty estimation and prediction for interdisciplinary ocean dynamics. *J. Comput. Phys.* **217**, 176–199 (2006).
6. O. B. Fringer, C. N. Dawson, R. He, D. K. Ralston, Y. J. Zhang, The future of coastal and estuarine modeling: Findings from a workshop. *Ocean Model.* **143**, 101458 (2019).
7. M. R. Alam, Predictability horizon of oceanic rogue waves. *Geophys. Res. Lett.* **41**, 8477–8485 (2014).
8. A. J. Majda, M. Branicki, Lessons in uncertainty quantification for turbulent dynamical systems. *Discrete Contin. Dyn. Syst. Ser. A* **32**, 3133–3221 (2012).
9. A. R. Osborne, T. L. Burch, Internal solitons in the Andaman Sea. *Science* **208**, 451–460 (1980).
10. L. Boegman, M. Stastna, Sediment resuspension and transport by internal solitary waves. *Annu. Rev. Fluid Mech.* **51**, 129–154 (2019).
11. D. Cacchione, L. F. Pratson, A. Ogston, The shaping of continental slopes by internal tides. *Science* **296**, 724–727 (2002).
12. Y. H. Wang, C. F. Dai, Y. Y. Chen, Physical and ecological processes of internal waves on an isolated reef ecosystem in the South China Sea. *Geophys. Res. Lett.* **34**, L18609 (2007).
13. X. Huang *et al.*, An extreme internal solitary wave event observed in the northern South China Sea. *Sci. Rep.* **6**, 30041 (2016).
14. D. J. Korteweg, G. de Vries, On the change of form of long waves advancing in a rectangular canal, and on a new type of long stationary waves. *Lond. Edinb. Dubl. Phil. Mag. J. Sci.* **39**, 422–443 (1895).
15. P. G. Drazin, R. S. Johnson, *Solitons: An Introduction* (Cambridge University Press, 1989).
16. K. G. Lamb, L. Yan, The evolution of internal wave undular bores: Comparisons of a fully nonlinear numerical model with weakly nonlinear theory. *J. Phys. Oceanogr.* **26**, 2712–2734 (1996).
17. P. E. Holloway, E. Pelinovsky, T. Talipova, A generalized Korteweg-de Vries model of internal tide transformation in the coastal zone. *J. Geophys. Res. Oceans* **104**, 18333–18350 (1999).
18. K. R. Helfrich, W. K. Melville, Long nonlinear internal waves. *Annu. Rev. Fluid Mech.* **38**, 395–425 (2006).
19. B. Øksendal, *Stochastic Differential Equations* (Springer, 2003).
20. E. L. Ionides, C. Bretó, A. A. King, Inference for nonlinear dynamical systems. *Proc. Natl. Acad. Sci. U.S.A.* **103**, 18438–18443 (2006).
21. M. Girolami, Bayesian inference for differential equations. *Theor. Comput. Sci.* **408**, 4–16 (2008).
22. A. M. Stuart, Inverse problems: A Bayesian perspective. *Acta Numer.* **19**, 451–559 (2010).
23. K. Law, A. Stuart, K. Zygalakis, *Data Assimilation* (Springer, Cham, Switzerland, 2015).
24. A. H. Jazwinski, *Stochastic Processes and Filtering Theory* (Courier Corporation, 2007).
25. G. Evensen, The ensemble Kalman filter: Theoretical formulation and practical implementation. *Ocean Dynam.* **53**, 343–367 (2003).
26. R. H. Shumway, D. S. Stoffer, *Time Series Analysis and Its Applications: With R Examples* (Springer Texts in Statistics, Springer International Publishing, ed. 4, 2017).
27. M. Katzfuss, J. R. Stroud, C. K. Wikle, Ensemble Kalman methods for high-dimensional hierarchical dynamic space-time models. *J. Am. Stat. Assoc.*, 1–43 (2019).
28. J. Nocedal, S. Wright, *Numerical Optimization* (Springer Science & Business Media, 2006).
29. A. Debussche, J. Printems, Numerical simulation of the stochastic Korteweg–de Vries equation. *Phys. Nonlinear Phenom.* **134**, 200–226 (1999).
30. K. J. Burns, G. M. Vasil, J. S. Oishi, D. Lecoanet, B. P. Brown, Dedalus: A flexible framework for numerical simulations with spectral methods. *Phys. Rev. Res.* **2**, 023068 (2020).
31. D. Horn, J. Imberger, G. Ivey, L. Redekopp, A weakly nonlinear model of long internal waves in closed basins. *J. Fluid Mech.* **467**, 269–287 (2002).
32. R. Grimshaw, E. Pelinovsky, T. Talipova, Damping of large-amplitude solitary waves. *Wave Motion* **37**, 351–364 (2003).