34th Symposium on Naval Hydrodynamics
Washington, DC, USA, 26 June - 1 July 2022

Application of Gaussian process multi-fidelity optimal sampling to ship structural modeling

Stephen Guth[1], Bianca Champenois[1], and Themistoklis P. Sapsis[1]
([1]Department of Mechanical Engineering, Massachusetts Institute of Technology)

ABSTRACT

The design of marine structures involves many computationally intensive simulations in order to estimate quantities important for structural performance and safety. In many cases, high fidelity simulations are especially expensive to perform, but we may have access to relatively low fidelity data or simulations. While active search techniques can optimize selection of high fidelity experimental designs, multi-fidelity active search techniques can additionally leverage low fidelity data to further improve design of a limited high fidelity budget. We demonstrate multi-fidelity active search on a ship-wavegroup problem based on the Large Amplitude Motions Program, and show that low fidelity data can improve experimental design and tail matching in recovered vertical bending moment pdfs.
Keywords: Nonlinear ship dynamics and loads; heavy tails and extreme events; Gaussian process regression; wavegroups; reduced-order modeling; Karhunen-Loève expansion; optimal experimental design; active search; multi-fidelity modeling.

INTRODUCTION

The design process for marine structures relies extensively on structural simulations in order to estimate both performance characteristics [Naess and Moan, 2013] and safety characteristics [Belenky et al., 2016, Belenky et al., 2018]. One important quantity to model is the vertical bending moment (VBM) [Sapsis et al., 2020, Sapsis et al., 2021, Belenky et al., 2021], which is important for both structural stability and fatigue lifetimes. While high fidelity simulations that accurately resolve the VBM are expensive to evaluate, lower fidelity simulations are much cheaper.

Gaussian process regression can be used to form surrogate models of black box functions from known data. Furthermore, data from sources of varying levels of fidelity can be combined to form more accurate models. It is be computationally expensive to obtain data from high fidelity models, and low fidelity models are not sufficiently accurate. As such, it is necessary to find optimal sampling techniques so that each additional sample provides the most useful information to the model. Acquisition functions can be optimized to choose the next best point to sample for a single-fidelity model. The goal of this project is to use optimal sampling to build better multi-fidelity models and intelligently minimize acquisition functions to optimally choose the level of fidelity of the data from which a new point is taken from. Combining multi-fidelity modeling with optimal sampling will increase the accuracy of the surrogate model while decreasing the computational cost of building the model. The aforementioned techniques are demonstrated on theoretical systems as well as on a more realistic setup involving the vertical bending moment distribution for a ship in random head seas.

SHIP STRUCTURAL MODELING

We employ the Large Amplitude Motions Program (LAMP version 4.0.5) to characterize marine vessel hydrodynamics and structural responses. LAMP is a numerical solver that computes 3-D potential flow solution of the wave-body interaction problem in order to calculate the time-domain motions and loading of floating bodies [Shin et al., 2003] [Lin et al., 2007b] [Lin et al., 2007a] [Lin et al., 2010]. We follow Guth and Sapsis [Guth and Sapsis, Jpdf] for the model construction–here we briefly discuss the elements of the model required to interface with the multi-fidelity active search.

Representation of Sea State

The LAMP software simulates the motions and dynamics of a vessel as it passes along a specified trajectory through a specified sea state. For our work, we use the ONR Topside series flared variant geometry traveling with constant linear velocity through long crested (unidirectional) head seas with heading

180.0^o. The Froude–Krylov forces (hydrodynamic pressure force) induce a vertical bending moment (VBM), whose distribution previous works have found to have an important asymmetry [Sapsis, 2020],

When setting the design parameters for each simulation run, our primary area of interest is in specifying the sea state. In general, the sea surface elevation $x(\xi,t)$ is a stochastic process, which we assume to be zero mean, statistically stationary, and with Gaussian statistics.

Following [Guth and Sapsis, Jpdf], we use the JONSWAP spectrum [Hasselmann et al., 1973] to describe the time series $x(\xi = 0,t)$ at a fixed spatial location:

$$S_J(\omega) = \frac{\alpha g^2}{\omega^5}\exp\left[-\frac{5}{4}\left(\frac{\omega_p}{\omega}\right)^4\right]\gamma^r \qquad (1)$$
$$r = \exp\left[-\frac{(\omega-\omega_p)^2}{2\sigma^2\omega_p^2}\right]$$

where constants are

$$\alpha = 0.076\left(\frac{U_{10}^2}{Fg}\right)^{0.22}, \quad \omega_p = 22\left(\frac{g^2}{U_{10}F}\right)^{\frac{1}{3}},$$
$$\gamma = 3.0, \quad \sigma = \begin{cases} 0.07 & \omega \le \omega_p \\ 0.09 & \omega > \omega_p \end{cases},$$

and F is the wind fetch. We choose $\alpha = 0.06$ and modal period $T_m = 10s$ (together implying significant wave height of $H_s = 13.2m$). The procedure described may be generalized to other wave spectra without change.

Each simulated experimental run requires a distinct realization of this random process on a finite interval $[0,T]$ [Sclavounos, 2012, Anastopoulos et al., 2016, Anastopoulos and Spyrou, 2016]. We employ the Karhunen-Loève (KL) theorem [Karhunen, 1947, Loève, 1948]:

Theorem 1 (Karhunen Loève) *Consider the stochastic process $x(t)$ which is zero mean and square integrable on the probability space $(\Omega,\mathscr{F},\mathbb{P})$. Define the covariance function*

$$K_{xx}(s,t) = \mathbb{E}\left[x(s)x(t)\right], \qquad (2)$$

with corresponding integral operator over the interval $[0,T]$,

$$T_{K_x}\phi(t) = \int_0^T K_x(t,s)\phi(t)ds, \;\; t\in[0,T]. \qquad (3)$$

Then by Mercer's Theorem for every interval $[0,T]$ the operator T_{K_x} has an orthonormal basis of eigenvectors $\{\hat{e}_{i,T}(t)\}$ and corresponding eigenvalues $\{\lambda_i\}$. Moreover, the coefficients

$$\alpha_i = \int_0^T x(t)\hat{e}_{i,T}(t)dt \qquad (4)$$

are centered orthogonal random variables:

$$\mathbb{E}[\alpha_i\alpha_j] = 0 \;\; for \;\; i\neq j \;\; and \;\; Var(\alpha_i) = \mathbb{E}[Z_i^2] = \lambda_i. \quad (5)$$

Furthermore, we can expand the random process $x(t)$ as

$$x(t) = \sum_{i=1}^{\infty}\alpha_i\hat{e}_{i,T}(t), \;\; t\in[0,T]. \qquad (6)$$

In summary, the eigenvectors of the spatial covariance matrix of the sea surface form an orthonormal basis. The decomposition of $x(t)$ onto this basis produces a set of centered, orthogonal (in the random sense) coefficients. In particular, we can change back and forth between the function representation $x(t), t\in[0,T]$ and the coefficients representation $\alpha_i, i = 1,2,....$

We truncate the KL expansion in equation (6) at a finite number of modes, n. In this way, we may represent the stochastic process on the interval $[0,T]$ as an n-dimensional vector $\vec{\alpha}$ of KL coefficients, each component of which is an orthogonal random variable with variance λ_i. It is *this* design vector that we optimize over using the techniques discussed below.

Finally, we touch on the multi-fidelity represention. Because we introduced a truncation cutoff n, it is natural that we expect a lower choice of n to correspond to samples that are less faithful to the original stochastic process. This suggest that we may choose two different cut-offs, n_{low} and n_{high}, corresponding to the low fidelity model and high fidelity model described *supra*. For this work, drawing on experience with the results from [Guth and Sapsis, Jpdf], we use $n_{\text{high}} = 2-3$ and $n_{\text{low}} = 2$.

We also briefly note that while the KL reconstruction defines the sea surface elevation at $\xi = 0, t\in[0,T]$, we require the sea surface elevation on a wider spatial and temporal interval, both to accommodate the trajectory of the vessel, and for initializing LAMP's numerics. We use the stochastic preluding technique described in [Guth and Sapsis, Jpdf] to extend the wavegroup smoothly backwards (and forwards) in time in such as way that the prelude has statistics consistent with the original stochastic process. Additionally, we use the deep water dispersion relation

$$\omega = \sqrt{gk} \qquad (7)$$

in order to express the time series at $\xi = 0$ as a superposition of traveling monochromatic waves.

Representation of structural dynamics

As LAMP models the structural response of the vessel passing through the prescribed wavegroup, the software records a number of kinematic and dynamic quantities as a time series. We choose the vessel Vertical Bending Moment (VBM) as a representative quantity of interest, as previous works have found that the steady state distribution of VBM has an important asymmetry [Sapsis, 2020].

An important obstacle we encounter is that the VBM is a *time series*, not a scalar quantity. In order to fit the output into the optimal experimental design framework, we first represent the time series as a low dimensional vector, and then handle each component of the vector separately. Specifically, we employ the following KL representation for the VBM:

$$M_y(t|\boldsymbol{\alpha}) = \sum_{i=1}^{n_{out}} Q_i(\boldsymbol{\alpha})\hat{\mu}_{i,T}(t), \quad t \in [0,T]. \tag{8}$$

where $Q_i(\boldsymbol{\alpha})$ are the KL coefficients which are functions of the excitation wavegroup and $\hat{\mu}_{i,T}(t)$ are the VBM KL modes. We represent each $Q_i(\boldsymbol{\alpha})$ as a separate output component. For our presented results, we choose $n_{out} = 6$. We generate both low and high fidelity data by varying the wavegroup parameter n. Of course, in a more realistic setup the high fidelity simulation will not be obtained from a reduced order model, but rather through an experiment or a CFD simulation.

MULTI FIDELITY ACTIVE SEARCH

Gaussian Process Regression

Gaussian process regression can be used to model systems by making a prediction for the posterior based on the prior and the likelihood [Rasmussen and Williams, 2006]. The joint distribution of the observed values and the expected values of the model under the prior is

$$\begin{bmatrix} \mathbf{y} \\ \mathbf{f}_* \end{bmatrix} \sim N(0, \begin{bmatrix} K(X,X) & K(X,X_*) \\ K(X_*,X) & K(X_*,X_*) \end{bmatrix}) \tag{9}$$

$$\mathbf{f}_*|X,\mathbf{y},X_* \sim N(\bar{\mathbf{f}}_*, cov(\mathbf{f}_*)). \tag{10}$$

For economy of space, we abbreviate the kernel matrices $\mathbf{K}$ with the data matrix subscripts. The mean prediction is

$$\bar{\mathbf{f}}(\mathbf{x}_*) = \mathbf{K}_*[\mathbf{K} + \sigma_n I]^{-1}\mathbf{y} \tag{11}$$

and the covariance of the posterior is

$$cov(\mathbf{f}) = K(X_*,X_*) - \mathbf{K}_*[\mathbf{K} + \sigma_n I]^{-1}\mathbf{K}_*^\intercal. \tag{12}$$

The covariance function is taken as the radial basis function (RBF) with automatic relevance determination. σ_f and θ are hyperparameters which represent the signal variance and characteristic length scales.

$$\begin{aligned} k(\mathbf{x_p},\mathbf{x_q}) &= cov(f(\mathbf{x_p}), f(\mathbf{x_q})) \\ &= \sigma_f^2 \exp\left(-\frac{1}{2}(\mathbf{x_q}-\mathbf{x_p})^T\theta(\mathbf{x_q}-\mathbf{x_p})\right). \end{aligned} \tag{13}$$

The hyperparameters for the covariance function are chosen through maximum likelihood estimation.

Multi-Fidelity Modeling

Given s levels of fidelity, the model with the lowest fidelity is denoted by $\mathbf{x_1}, y_1, \bar{\mathbf{f}}_{*1}$, and the model with the highest fidelity is denoted by $\mathbf{x_s}, y_s, \bar{\mathbf{f}}_{*s}$ [Perdikaris et al., 2015]. The prediction for the model with the lowest fidelity follows the Gaussian process regression steps above

$$\bar{\mathbf{f}}_1(\mathbf{x}_*) = \mathbf{K}_{*1}[\mathbf{K}_{11} + \sigma_{n1}I]^{-1}\mathbf{y}_1$$

with covariance

$$cov(\bar{\mathbf{f}}_1) = \mathbf{K}_{**} - \mathbf{K}_{*1}[\mathbf{K}_{11} + \sigma_{n1}I]^{-1}\mathbf{K}_{1*}.$$

Each following model is of the form

$$\bar{\mathbf{f}}_t(\mathbf{x}_*) = \rho_{t-1}\bar{\mathbf{f}}_{t-1} + \delta_t \qquad t = 2,...,s. \tag{14}$$

In this project, we consider only two levels of fidelity, so the prediction for the highest level of fidelity, $s = 2$, can be computed using the following equation

$$\bar{\mathbf{f}}_2(\mathbf{x}_*) = \rho\bar{\mathbf{f}}_1(\mathbf{x}_*) + \mu_d + \mathbf{K}_{*2}[\mathbf{K}_{22} + \sigma_{n2}I]^{-1}(\mathbf{y} - \rho\bar{\mathbf{f}}_1(\mathbf{x}_2) - \mu_d).$$

The corresponding covariance is

$$cov(\bar{\mathbf{f}}_2) = \rho^2 cov(\bar{\mathbf{f}}_1) + \mathbf{K}_{**} - \mathbf{K}_{*2}[K\mathbf{K}_{22} + \sigma_{n2}I]^{-1}\mathbf{K}_{2*}$$

where ρ and μ_d are hyperparameters that are different for each level of fidelity. Like the σ_f and θ of the covariance function, ρ and μ_d can be chosen through maximum likelihood estimation or other optimization techniques.

Optimal Sampling

The maximum response is evaluated at a set number of initial. points. Then, the following points are chosen sequentially through optimization of a chosen acquisition function. In general, acquisition functions optimize the selection of individual samples to globally reduce the uncertainty of the model [Blanchard and Sapsis, 2021a]. Commonly used acquisition functions can be modified to give more weight to points that are more "relevant." Relevant points either have a larger impact on the output or have a

higher probability of occurrence. These output-weighted acquisition functions make use of the likelihood ratio $w(x)$

$$w(x) = \frac{p_x(x)}{p_\mu(\mu(x))}, \tag{15}$$

which significantly improves the convergence rate bringing it close to the optimal [Sapsis and Blanchard, 2022]. In particular, output-weighted sampling is valuable in modeling extreme events because it uses information about the output space from previous observations to decide where to perform future samples. Under these conditions, extreme events are more likely to be found because future samples are chosen based on their impact on the output of the black box function. [Blanchard and Sapsis, 2021b]

For the acquisition function in this work, we highlight the choice of likelihood weighted uncertainty sampling:

$$a_{US-LW}(x) = \sigma^2(x)w(x). \tag{16}$$

Fidelity-Weighted Cost Function

Given that the higher fidelity function is significantly more computationally expensive than the lower fidelity function, the acquisition function can be adjusted to account for the cost of each level of fidelity. By assigning more weight to the higher fidelity model, the acquisition function will tend to choose more points from the lower fidelity function to build the model. The ratio between the cost of the low and high fidelity functions needs to be reasonable to avoid overfitting. The ratio between the additional cost and the original acquisition function also needs to be balanced to make sure that both parts of the total objective are taken into account during the minimization.

The cost-ratio penalty term can be expressed as

$$C_{lo} = \frac{b_1}{b_2}(n_1+1)+n_2$$
$$C_{hi} = \frac{b_1}{b_2}n_1+(n_2+1)$$
$$J = a(x)+\frac{1}{n_{iter}}C_*.$$

where b_1 and b_2 tunable parameters that control the relative costs of sampling from the low fidelity and high fidelity models, respectively. The model selects between sampling from the low and high fidelity function by comparing the optimum value of J for each choice of the penalty term (C_{lo} or C_{hi}).

We can summarize the active search algorithm as the following:

- Initialize $\bar{f}_2$ on initial datasets $\mathscr{D}_1 = \{\mathbf{x}_i, y_i\}_{lo-fi}$, $\mathscr{D}_2 = \{\mathbf{x}_i, y_i\}_{hi-fi}$
- For each step $n \in [1, ..., n_{iter}]$:
  1. Minimize objective function J given low-fidelity model;
  $$J_{lo-fi} = \min_{\mathbf{x}\in\mathscr{X}} J(\mathbf{x}; \bar{f}_1, \mathscr{D}_1, n_1 = n_1+1)$$
  2. Minimize acquisition function J given high-fidelity model;
  $$J_{hi-fi} = \min_{\mathbf{x}\in\mathscr{X}} J(\mathbf{x}; \bar{f}_2, \mathscr{D}_1, \mathscr{D}_2, n_2 = n_2+1)$$
  3. Choose model $\bar{f}_i$ in $\{\bar{f}_1, \bar{f}_2\}$ that minimizes J to select next best point $\mathbf{x}_{n+1}$
  $$\mathbf{x}_{n+1} = \arg\min_{\mathbf{x}\in\mathscr{X}} J(\mathbf{x}; \bar{f}_i, \mathscr{D}_1, \mathscr{D}_2)$$
  4. Evaluate respective simulation f_i at $\mathbf{x}_{n+1}$
  $$y_{n+1} = f_i(\mathbf{x}_{n+1})$$
  5. Augment chosen dataset: $\mathscr{D}_i = \mathscr{D}_i \cup \{\mathbf{x}_{n+1}, y_{n+1}\}$
  6. Update surrogate model $\bar{f}_2$ using multi-fidelity Gaussian process regression

For the marine structure example, we choose initial data sets that include $h = 4$ initial high fidelity points and $l \in \{0, 200\}$ initial low fidelity points. Further, we choose a cost-ratio penalty so that only additional high fidelity points are selected.

Acquisition Functions with Vector Output

For the ship structure problem described above, we note an additional complication: to reconstruct the VBM, we estimate n_{out} distinct mode coefficients. While this representation could be considered in terms of a vector valued Gaussian process, it is simpler to consider a set of n_{out} distinct scalar Gaussian processes. In particular, the KL construction leads to Gaussian posteriors that are nearly uncorrelated.

How can we use the previously described algorithm for picking sample points when we have n_{out} distinct Gaussian processes we would like to 'jointly' optimize? For this work, we use a simple *Round Robin* strategy: at step i, we choose the j^{th} mode to optimize, and we relate i and j using the relation

$$j = i \bmod n_{out}. \tag{17}$$

In other words, each black box evaluation provides a sample for the Gaussian process surrogate associated with *each* output mode. While there are more sophisticated strategies for applying scalar acquisition function to vector functions, we found that the round robin performed adequately for this marine structural problem.

Suitability for Marine Dynamics

Before applying this multi-fidelity framework for actively sampled Gaussian process regression, we take a moment to check whether Gaussian process modeling (kriging) is a good fit for this problem area. There are two main requirements for using Gaussian processes, and another requirement for multi fidelity.

First, Gaussian process regression is best suited for low dimension problems. Capping the dimension of the input space below $n \approx 10$ is an important topological requirement. For this problem, the wavegroup parametrization is very low dimensional–$n \in \{2,3\}$.

Second, kernel matrix inversion for Gaussian processes scales algorithmically with the third power of the sample size. This makes Gaussian process regression a poor choice for many big data applications $n_{samples} > 1000$. For this problem, we always restrict our attention to Gaussian process models with less than 625 samples. Indeed, our goal with active sampling is to *minimize* the training set size.

Finally, multi fidelity is a useful framework for data integration when the low fidelity function and the high fidelity function are highly correlated, so that relationships learned by the low fidelity model can be transferred to the high fidelity model. In the two scenarios we present below (restricted inputs and noisy inputs), relationship partially holds. However, the significant intrinsic noise in this problem may be an important limitation.

DEMONSTRATION AND RESULTS

For this work, we apply the multi-fidelity active search approach described above to the marine structure problem. Our design vector describes the coefficients of parametrized wavegroups with $T = 60$. We choose $s = 2$ levels of fidelity–a high fidelity model we assume to be 'true,' and a low fidelity model.

For our high fidelity model, we choose to focus on two values of n_{high}, a low dimensional example with $n_{\text{high}} = 2$ and a moderate dimensional example with $n_{\text{high}} = 3$. We note that Guth and Sapsis [Guth and Sapsis, Jpdf] found that $n_{\text{high}} = 3$ was generally sufficient to recover accurate steady state statistics. For each problem, we construct a high fidelity "black box" function by collecting 625 wavegroups and associated VBM time series. The wavegroup coefficient vectors $\vec{\alpha}$ are constructed by Latin Hypercube Sampling from a box domain with radius $z^*\sqrt{\lambda_i}$, where λ_i is the KL eigenvalue associated with component i and $z^* = 4.5$ is cutoff z-score that controls how far into the distribution tails sampling takes place.

For our low fidelity model, we use different *wavegroup* models, simulated using the same LAMP code. In one case, we use *noisy* wavegroups, with n fixed coefficients and higher coefficients allowed to vary randomly. In the other case, we use *truncated* wavegroups, where one of the n fixed coefficients was truncated (set to 0).

For active search, we use the likelihood weighted uncertainty sampling from equation (16). This was implemented using both the Python package [Paleyes et al., 2019], as well as the active search package GPSEARCH [Blanchard, 2021].

In the following plots, the thick black line labeled 'true' corresponds to the distribution generated by 10^5 samples from the high fidelity black box. Thus, the 'true' pdf differs slightly between high fidelity models with $n_{\text{high}} = 2$ and $n_{\text{high}} = 3$.

For each problem, we initialized the optimization with $h = 4$ high fidelity samples and $l \in \{0, 200\}$ low fidelity samples. This allows us to estimate the effect of low fidelity data on the sample selection and pdf reconstruction. We plotted the reconstructed pdf for various numbers of additional high fidelity samples, where k is the sum of the $h = 4$ initial samples and the sampled selected via active search. Each reconstructed pdf is generated by 10^5 samples. Due to our particular interest in extreme values of the VBM, we principally evaluate the accuracy of the pdf reconstruction by examining the tail mass.

Low Dimensional Problem ($n_{high} = 2$)

We chose $n_{\text{high}} = 2$ as a simple demonstration. For this case, the low fidelity model is generated by considering $n = 6$ dimensional wavegroups. From the 6-dimensional vector of coefficients $\vec{\alpha}$ we vary only the first two, while the other four are chosen randomly. Because only two input components are controlled, the low fidelity model is a noisier version of the high fidelity model, with somewhat heteroskedastistic noise. In particular, the form of the noise causes the low fidelity model to *generally overestimate* the VBM.

In figure 1 we show the recovery of the VBM pdf both without any low fidelity samples, as well as with $l = 100$ and $l = 200$ low fidelity samples. That is, for the vessel VBM M_y, we show a log scale plot of the probability density function $f_M(s)$. We focus on the tails of this distribution, which for the true model (black, bold) decay faster than Gaussian.

First, for very small amounts of high fidelity data, the overestimate of the tail spread for $l = 0$ is greater than for to $l = 100$ and $l = 200$. This is surprising, because the (true) low fidelity model has heavier tails. This suggests that the multi-fidelity combination is able to model the tails better than high fidelity data alone.

Second, for larger amounts of high fidelity data, the $l = 0$ and $l = 100$ pdfs converge in the tails faster.

One possible explanation for this behavior is that, as the ratio of low fidelity to high fidelity data increases, the hyperparameter tuning necessary associated with the high-low delta kernel are more difficult.

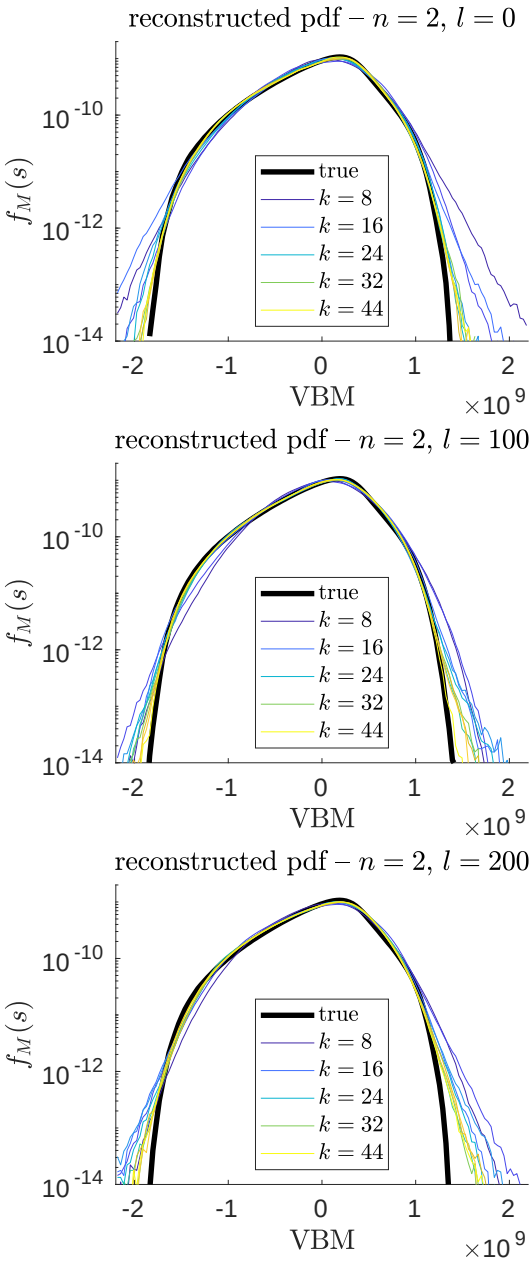

Figure 1: Reconstruction of VBM pdf, for $n = 2$ high fidelity model and (**top**) $l = 0$, (**middle**) $l = 100$, and (**bottom**) $l = 200$ **noisy** low fidelity points. Colored lines correspond to increasing number of sampled high fidelity points.

In figure 2, we show the locations of the initial low fidelity samples (red) and the actively sampled high fidelity samples (blue).

Moderate Dimensional Problem ($n_{high} = 3$)

For the $n_{high} = 3$ case, we consider two different low fidelity models. The first low fidelity model (middle figure 3) is constructed as previously, corresponding to a wavegroup with random higher order components.

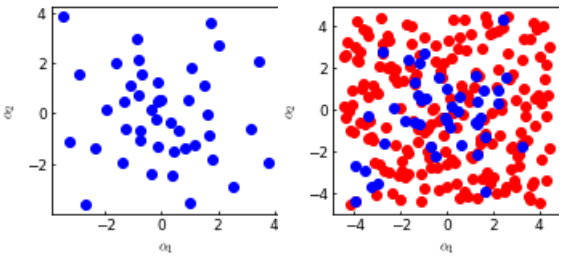

Figure 2: Comparison of sample locations for (*Left*) only high fidelity points and (*Right*) $l = 200$ red low fidelity points.

The second low fidelity model (bottom figure 3) is constructed by truncating the third component of $\vec{\alpha}$, and feeding the resulting $n = 2$ dimensional coefficient vector into the $n = 2$ black box. Unlike the 'noisy' low fidelity model, this model is low fidelity because it ignores one dimension of the input.

We note that for very small numbers of high fidelity samples, the noisy low fidelity model overestimates the variance of the reconstructed pdf. This is not hard to understand–in the absence of better data, the model estimates the true (high fidelity) spread by the low fidelity spread, which was designed to be noisy.

For all cases (high fidelity only, noisy low fidelity, and truncated low fidelity), the reconstructed pdf matches the broad features, such as the asymmetric shoulder, quickly (i.e. with less than 12 high fidelity samples). We see that in the noisy low fidelity case, it appears to converge slightly slower. This can be a result of the overfitting problem mentioned supra. Alternatively, when the low fidelity model contains *more* (noisy) features, it may be more challenging for the multi fidelity model to learn the right balance between between the low fidelity and delta kernels.

We finally note that the truncated low fidelity model appears to best learn the distribution tails–thought only after an early transient while it identifies the way that the low and high fidelity samples treat the truncated dimension differently. For this problem, however, the ultimate difference is not that great. This is likely because, for this problem, the pdf tails are driven by 'extreme' data points near the boundary of the sample domain, which are sharply limited by the total number of high fidelity samples.

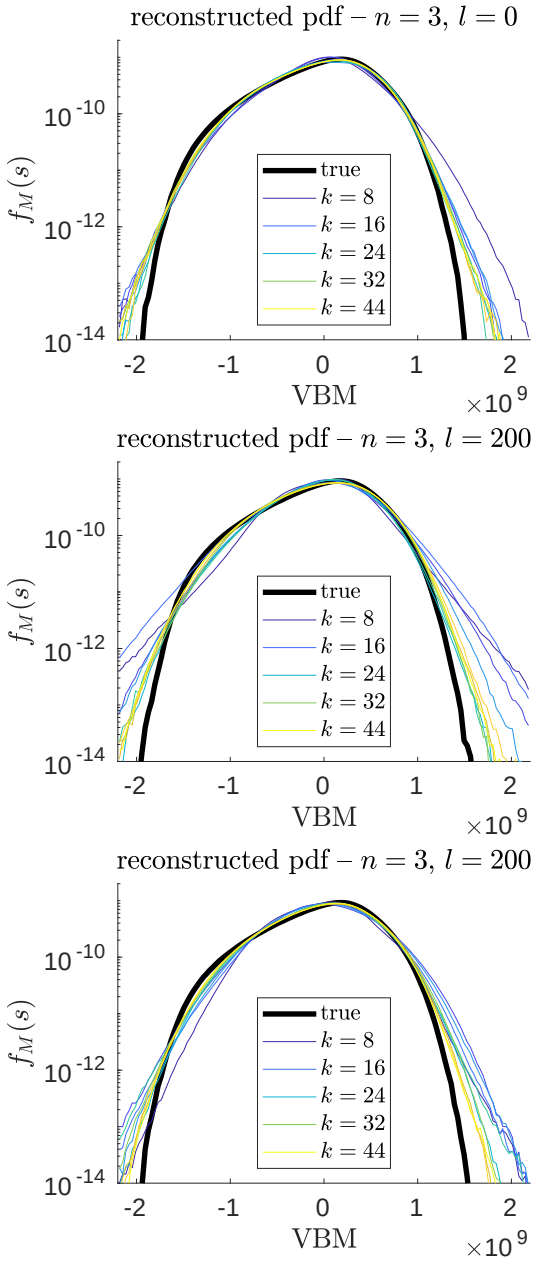

Figure 3: Reconstruction of VBM pdf, for $n = 3$ high fidelity model and (**top**) $l = 0$ low fidelity points, (*middle*) $l = 200$ *noisy* low fidelity points and (*bottom*) *truncated* low fidelity points. Colored lines correspond to increasing number of sampled high fidelity points.

Learning Plots

We quantify the statistical accuracy of the surrogate model by computing an error metric between the reconstructed pdf and the true pdf. Because different error metrics emphasize different aspects of the distribution, we present two: the mean absolute error (MAE, l_1 norm)and the mean absolute log error (MALE, log l_1 norm):

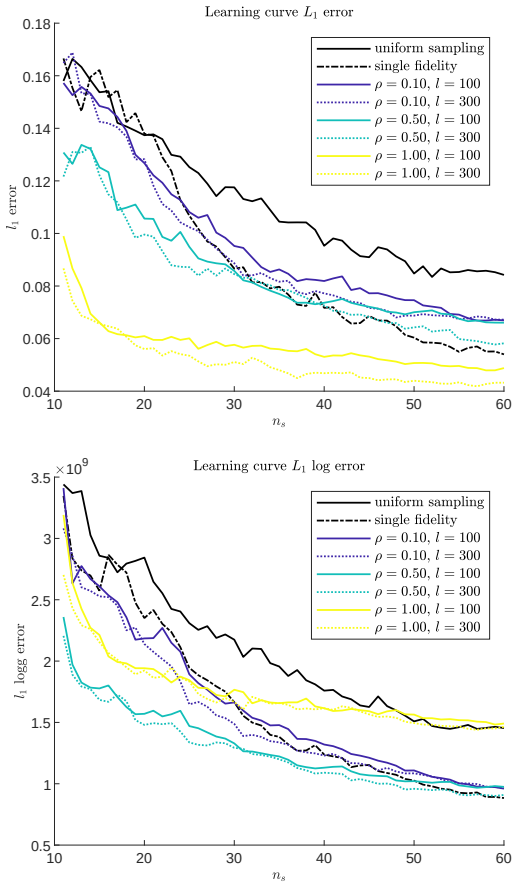

Figure 4: Error curves showing the relationship between size of training set and the reconstructed VBM pdf and the true pdf, using *truncated* low fidelity data. (*Top*) Mean absolute error. (*Bottom*) Mean absolute log error.

$$\varepsilon_1 = \int |f_M^{(i)}(s) - f_M^*(s)|$$

$$\varepsilon_2 = \int |\log f_M^{(i)}(s) - \log f_M^*(s)|$$

We note that for numerical reasons, we must choose a cutoff scale for ε_2, which controls how far into the tail is considered. We used a cutoff corresponding to $f_M(s) \sim 1 \times 10^{-13}$.

Figure 4 presents learning curves for different sized low fidelity sets ($l \in \{100, 300\}$) and different values of ρ ($\rho \in \{0.1, 0.5, 1.0\}$).

We note that *rho* $= 1.0$ behaves very differently when considering MAE (which preferentially detects convergence near the distribution mode) as when

considering MALE (which preferentially detects convergence near the tails). This is likely due to a large value of ρ suppressing the information gained from new high fidelity samples.

We also note that there is an range where the addition of low fidelity data greatly increases accuracy. This improvement is seen both in comparing multi fidelity active search to single fidelity active search (black dot-dashed line), and when comparing small low fidelity sets to larger low fidelity sets (colored continuous lines vs dotted lines).

We interpret this finding to mean that the multi fidelity is most useful when actively sampling in the very-low data regime. As the density of data points increases, the advantage offered by the low fidelity data begins to decrease.

CONCLUSIONS

In this work, we described how to implement a multi-fidelity Gaussian process framework for an active search problem in the context of ship structural response under random waves. Specifically, we applied this multi-fidelity active search to a model in marine structural dynamics, where a low fidelity approximate solution corresponded to either noisy simulation data (corresponding to a noisy wavegroup) or truncated fidelity simulation data (corresponding to a low fidelity wavegroup). We found that in certain cases, the low fidelity data meaningfully improves the recovered pdf tails. However, in other circumstances there is no improvement, or even degraded recovery. One example is situations where the low-fidelity data is over represented and the low-fidelity model contains additional noise features. Because the efficacy of the multi fidelity approach appears to have some dependence on the relationship between low fidelity and high fidelity models, further work should investigate other low-fidelity high-fidelity model pairs beyond the low/high-fidelity wavegroup parametrization.

ACKNOWLEDGEMENTS

The work described in this paper has been funded by the Office of Naval Research (ONR) under Dr. Woei-Min Lin.

REFERENCES

[Anastopoulos and Spyrou, 2016] Anastopoulos, P. A. and Spyrou, K. J. (2016). Ship dynamic stability assessment based on realistic wave group excitations. Ocean Engineering, 120:256–263.

[Anastopoulos et al., 2016] Anastopoulos, P. A., Spyrou, K. J., Bassler, C. C., and Belenky, V. (2016). Towards an improved critical wave groups method for the probabilistic assessment of large ship motions in irregular seas. Probabilistic Engineering Mechanics, 44:18–27. Special Issue Based on Papers Presented at the 7th International Conference on Computational Stochastic Mechanics (CSM7).

[Belenky et al., 2016] Belenky, V., Glotzer, D., Pipiras, V., and Sapsis, T. (2016). On the tail of nonlinear roll motions. In Proceedings of the 15th International Conference on Ship Stability Workshop.

[Belenky et al., 2018] Belenky, V., Weems, K., Pipiras, V., Glotzer, D., and Sapsis, T. (2018). Tail structure of roll and metric of capsizing in irregular waves. In 32nd Symposium on Naval Hydrodynamics.

[Belenky et al., 2021] Belenky, V., Weems, K., Sapsis, T., and Pipiras, V. (2021). Influence of deck submergence events on extreme properties of wave-induced vertical bending moment. In Proceedings of the 1st International Conference on the Stability and Safety of Ships and Ocean Vehicles.

[Blanchard, 2021] Blanchard, A. (2021). gpsearch. https://github.com/ablancha/gpsearch.

[Blanchard and Sapsis, 2021a] Blanchard, A. and Sapsis, T. (2021a). Bayesian optimization with output-weighted optimal sampling. Journal of Computational Physics, 425:109901.

[Blanchard and Sapsis, 2021b] Blanchard, A. and Sapsis, T. P. (2021b). Output-weighted optimal sampling for bayesian experimental design and uncertainty quantification. SIAM/ASA J. Uncertain. Quantification, 9:564–592.

[Guth and Sapsis, Jpdf] Guth, S. and Sapsis, T. (Submitted 2022 (https://sandlab.mit.edu/Papers/22_OEJ.pdf)). Wavegroup-based gaussian process regression for ship dynamics: Between the Scylla of slow Karhunen-Loève convergence and the Charybdis of transient features.

[Hasselmann et al., 1973] Hasselmann, K., Barnett, T., Bouws, E., Carlson, H., Cartwright, D., Enke, K., Ewing, J., Gienapp, H., Hasselmann, D., Kruseman, P., Meerburg, A., Muller, P., Olbers, D., Richter, K., Sell, W., and Walden, H. (1973). Measurements of wind-wave growth and swell decay during the Joint North Sea Wave Project (JONSWAP). Deut. Hydrogr. Z., 8:1–95.

[Karhunen, 1947] Karhunen, K. (1947). Über lineare Methoden in der Wahrscheinlichkeitsrechnung. Ann. Acad. Sci. Fennicae Ser. A. I. Math.-Phys., 1947(37):79.

[Lin et al., 2007a] Lin, W.-M., Collette, M., Lavis, D. R., Jessup, S. D., and Kuhn, J. (2007a). Recent hydrodynamic tool development and validation for motions and slam loads on ocean-going high-speed vessels.

[Lin et al., 2007b] Lin, W.-M., Zhang, S., Weems, K., Jones, P., Meinhold, M., Metcalf, B., and Powers, A. M. (2007b). Numerical simulation and validation study of wetdeck slamming on high speed catamaran. International Conference on Numerical Ship Hydrodynamics [9th].

[Lin et al., 2010] Lin, W.-M., Zhang, S., and Weems, K. M. (2010). Numerical simulations of surface effect ship in waves. In Proceedings of the 2010 Conference on Grand Challenges in Modeling & Simulation, GCMS '10, page 414–421, Vista, CA. Society for Modeling & Simulation International.

[Loève, 1948] Loève, M. (1948). Processus Stochastiques et Mouvement Brownien, chapter Fonctions aléatoires du second ordre. Gauthier-Villars.

[Naess and Moan, 2013] Naess, A. and Moan, T. (2013). Stochastic Dynamics of Marine Structures. Cambridge University Press.

[Paleyes et al., 2019] Paleyes, A., Pullin, M., Mahsereci, M., Lawrence, N., and González, J. (2019). Emulation of physical processes with emukit. In Second Workshop on Machine Learning and the Physical Sciences, NeurIPS.

[Perdikaris et al., 2015] Perdikaris, P., Venturi, D., Royset, J. O., and Karniadakis, G. E. (2015). Multi-fidelity modelling via recursive co-kriging and Gaussian–Markov random fields. Proceedings. Mathematical, physical, and engineering sciences.

[Rasmussen and Williams, 2006] Rasmussen, C. E. and Williams, C. K. I. (2006). Gaussian Processes for Machine Learning. Adaptive Computation and Machine Learning. MIT Press, Cambridge, MA, USA.

[Sapsis et al., 2021] Sapsis, T., Belenky, V., Weems, K., and Pipiras, V. (2021). Extreme properties of impact-induced vertical bending moments. In Proceedings of the 1st International Conference on the Stability and Safety of Ships and Ocean Vehicles.

[Sapsis et al., 2020] Sapsis, T., Pipiras, V., Weems, K., and Belenky, V. (2020). On extreme value properties of vertical bending moment. In Proceedings of the 33rd Symposium on Naval Hydrodynamics Osaka, Japan (Virtual).

[Sapsis, 2020] Sapsis, T. P. (2020). Output-weighted optimal sampling for bayesian regression and rare event statistics using few samples. Proceedings of the Royal Society A, page 476.

[Sapsis and Blanchard, 2022] Sapsis, T. P. and Blanchard, A. (2022). Optimal criteria and their asymptotic form for data selection in data-driven reduced-order modeling with Gaussian process regression. Phil. Trans. Roy. Soc. A, page In Press.

[Sclavounos, 2012] Sclavounos, P. D. (2012). Karhunen–Loève representation of stochastic ocean waves. Proceedings of the Royal Society A, 468(2145):2574–2594.

[Shin et al., 2003] Shin, Y., Belenky, V., Lin, W., Weems, K., Engle, A., McTaggart, K., Falzarano, J., Hutchison, B., Gerigk, M., and Grochowalski, S. (2003). Nonlinear time domain simulation technology for seakeeping and wave-load analysis for modern ship design. Transactions - Society of Naval Architects and Marine Engineers, 111:557–583.