

# **The Naval Seafloor Evolution Architecture: A Platform for Predicting Dynamic Seafloor Roughness**

WILLIAM S. KEARNEY

*American Society for Engineering Education Postdoctoral Fellow*

ALLISON M. PENKO

*Seafloor Sciences Branch  
Ocean Sciences Division*

July 28, 2022

# REPORT DOCUMENTATION PAGE

**PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ORGANIZATION**

<b>1. REPORT DATE</b> 07-28-2022	<b>2. REPORT TYPE</b> Formal Report	<b>3. DATES COVERED</b>	
		<b>START DATE</b> 01 OCT 2019	<b>END DATE</b> 30 SEP 2022

**4. TITLE AND SUBTITLE**  
The Naval Seafloor Evolution Architecture: A Platform for Predicting Dynamic Seafloor Roughness

<b>5a. CONTRACT NUMBER</b>	<b>5b. GRANT NUMBER</b>	<b>5c. PROGRAM ELEMENT NUMBER</b> 0601153N
----------------------------	-------------------------	---

<b>5d. PROJECT NUMBER</b> 100001764031.0010	<b>5e. TASK NUMBER</b>	<b>5f. WORK UNIT NUMBER</b> 1R32
--	------------------------	-------------------------------------

**6. AUTHOR(S)**  
William S. Kearney, Allison M. Penko

<b>7. PERFORMING ORGANIZATION / AFFILIATION NAME(S) AND ADDRESS(ES)</b> US Naval Research Laboratory 1005 Balch Blvd. Stennis Space Center, MS 39529	<b>8. PERFORMING ORGANIZATION REPORT NUMBER</b> NRL/7350/FR--2022/2
---	--

<b>9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)</b> Office of Naval Research One Liberty Center 875 N. Randolph Street, Suite 1425 Arlington, VA 22203-1995	<b>10. SPONSOR / MONITOR'S ACRONYM(S)</b> ONR	<b>11. SPONSOR / MONITOR'S REPORT NUMBER(S)</b>
--	--	---

**12. DISTRIBUTION / AVAILABILITY STATEMENT**  
DISTRIBUTION STATEMENT A. Approved for public release. Distribution is unlimited.

**13. SUPPLEMENTAL NOTES**

**14. ABSTRACT**  
Predicting the temporal and spatial dynamics of seafloor roughness is important for understanding bottom boundary layer hydrodynamics. The Navy Seafloor Evolution Architecture (NSEA) is a platform for modeling the dynamic nature of the seafloor by combining hydrodynamic forcing information and observations from diverse sources. NSEA's three modules include a specification of hydrodynamic forcing, a seafloor evolution model, and a model to generate roughness realizations. It can be run in forward mode to predict seafloor roughness including the uncertainty from forcing information, or in inverse mode to estimate parameters from observed seafloor roughness. The model is demonstrated and shown to have good agreement with a field dataset of observed seafloor roughness. Similarly running in inverse mode, NSEA was demonstrated to predict the observed mean sediment grain size with good agreement. NSEA's modularity allows for a wide range of applications in hydrodynamic and acoustic modeling, and is built within an expandable framework that lends for coupling to such models with minimal effort.

**15. SUBJECT TERMS**

<b>16. SECURITY CLASSIFICATION OF:</b>			<b>17. LIMITATION OF ABSTRACT</b> SAR	<b>18. NUMBER OF PAGES</b> 44
<b>a. REPORT</b> U/U	<b>b. ABSTRACT</b> U/U	<b>c. THIS PAGE</b> U/U		

<b>19a. NAME OF RESPONSIBLE PERSON</b> Allison M. Penko	<b>19b. PHONE NUMBER (Include area code)</b> (228) 688-4823
--	--

This page intentionally left blank

## CONTENTS

1. INTRODUCTION .....	1
1.1 Statistical Characterization of Seafloor Roughness .....	2
2. SYSTEM ARCHITECTURE .....	5
2.1 Forcing .....	5
2.2 Seafloor Evolution.....	6
2.3 Seafloor Synthesis.....	11
2.4 Observations.....	14
3. STATISTICAL INFERENCE.....	18
3.1 Approximate Bayesian Computation in NSEA.....	19
4. IMPLEMENTATION.....	20
4.1 StingrayCore.....	20
4.2 NSEM.....	23
4.3 SyntheticSeafloor.....	25
4.4 Demo Code .....	27
4.5 Stingray .....	28
5. ARCHITECTURE DEMONSTRATION .....	29
5.1 Forward Simulation.....	30
5.2 Statistical Inference .....	32
6. CONCLUSIONS .....	34
6.1 Future Work.....	34
REFERENCES .....	35

# THE NAVAL SEAFLOOR EVOLUTION ARCHITECTURE: A PLATFORM FOR PREDICTING DYNAMIC SEAFLOOR ROUGHNESS

## 1. INTRODUCTION

The vertical and horizontal variation of the seafloor impacts ocean wave attenuation, acoustic propagation, and the scour and burial of objects on the seabed. These elevation variations cover a wide range of spatial and temporal scales, from the grain scales of seafloor sediments ( $O(100\text{ mm})$ ) to the ocean basins ( $O(10,000\text{ km})$ ). However, bathymetric data is available only at a finite spatial and temporal resolution. The General Bathymetric Chart of the Ocean's global gridded bathymetry data, for example, is provided at a 15-arc-second resolution, which is approximately 450 meters at the equator. "Roughness" refers to the variability in seafloor elevation below this grid scale, though the exact delineation between roughness and bathymetry varies with the available data.

Seafloor roughness is particularly important to hydrodynamics, because it enhances turbulence within the bottom boundary layer and exerts friction on the overlying flow [1, 2], and acoustics, because it modifies the scattering of high-frequency sound waves from the sediment-water interface [3]. Roughness also evolves as sediment is transported and bedforms such as ripples and dunes are created or destroyed. Because roughness is not measured by definition, it must be predicted from observed or modeled geological properties and hydrodynamic forcing. High-resolution numerical models [4–8] can directly simulate the interaction of hydrodynamics and sediment transport in the bottom boundary layer, resulting in the formation and evolution of bedforms. However, these models are computationally intensive and are not suitable for embedding in operational hydrodynamic or acoustic models. Embedded models are statistical in nature, predicting an estimate of the roughness such as a length scale or power spectrum. They typically rely on empirical parameterizations developed from data collected in laboratory and field settings. The majority of models used operationally are equilibrium models that predict ripple heights and lengths from instantaneous hydrodynamic information [9–13]. These models are computationally efficient, but they fail to predict roughness out of equilibrium with the forcing. Non-equilibrium phenomena are most evident during ripple formation, as it takes time for ripples to adjust to new forcing [14], and under relict conditions, in which ripples persist under hydrodynamic forcing not strong enough to mobilize sediment [15, 16].

Non-equilibrium roughness models [14, 17–19] simulate the time evolution of roughness under changing forcing conditions and have been used successfully to predict roughness on shallow continental shelves. These models are statistical models, typically predicting roughness power spectra rather than the exact evolution of the seafloor elevation. They rely on empirical parameterizations of both sediment transport [20] and equilibrium ripple geometry, but embed these empirical parameterizations within a physics-based model for evolution.

The Naval Seafloor Evolution Architecture (NSEA) evolved from a non-equilibrium roughness model also known as NSEA (the Navy Seafloor Evolution Archetype [19]).<sup>1</sup> This roughness model now forms the central

---

Manuscript approved October 25, 2022.

<sup>1</sup>Further uses of NSEA within this technical report refer to the Naval Seafloor Evolution Architecture, and the non-equilibrium roughness model derived from [19] will be called the Navy Seafloor Evolution Model (NSEM).

part of a larger framework for characterizing the spatial and temporal evolution of seafloor roughness on operationally relevant time scales. This technical report describes in detail the components of the framework in Section 2, then discusses the application of statistical inference methods within the model in Section 3. Section 4 describes a reference implementation of NSEA in the Julia programming language [21], and Section 5 applies NSEA to seafloor roughness prediction and the estimation of sediment properties using data from a field experiment. Finally, Section 6 summarizes the current state of the architecture and presents future directions for research and development of NSEA.

## 1.1 Statistical Characterization of Seafloor Roughness

Before describing NSEA, it is valuable to describe how to characterize the statistical properties of the seafloor elevation and introducing notation that will be used throughout this technical report. The elevation of the seafloor,  $\eta(\mathbf{x}, t)$ , is a spatiotemporal *random field*, a stochastic process that is indexed by space,  $\mathbf{x} \in \mathbb{R}^2$ , and time,  $t \in \mathbb{R}^+$ . The issues that arise when modeling the seafloor elevation as a stochastic process will be discussed below, but for now, the seafloor elevation field is assumed to be described by a probability distribution that gives the likelihood of observing a particular configuration of the seafloor elevation in space and time. The goal of this analysis is to compute certain statistical characteristics of the random seafloor elevation. These include the mean function

$$\mu(\mathbf{x}, t) = \langle \eta(\mathbf{x}, t) \rangle, \quad (1)$$

the covariance function

$$K(\mathbf{x}_1, t_1, \mathbf{x}_2, t_2) = \langle (\eta(\mathbf{x}_1, t_1) - \mu(\mathbf{x}_1, t_1)) (\eta(\mathbf{x}_2, t_2) - \mu(\mathbf{x}_2, t_2)) \rangle, \quad (2)$$

and higher-order statistics of the form

$$R_n(\mathbf{x}_1, t_1, \dots, \mathbf{x}_n, t_n) = \langle \eta(\mathbf{x}_1, t_1) \dots \eta(\mathbf{x}_n, t_n) \rangle \quad (3)$$

that are called the "moments" of the random field.

Other important statistical properties can be derived from these moments, such as the mean square roughness:

$$\langle \eta(\mathbf{x}, t)^2 \rangle = K(\mathbf{x}, t, \mathbf{x}, t). \quad (4)$$

The operator  $\langle \cdot \rangle$  represents the average (expectation) under the probability distribution for the elevation field. This is an *ensemble average*, and it does not correspond to a spatial or temporal average. As a result, the moments all generally depend on both space and time.

NSEA assumes that the seafloor elevation is wide-sense stationarity in space [22]. A wide-sense stationary process has a constant mean function ( $\mu(\mathbf{x}, t) = \mu(t)$ ), and the equal-time covariance function  $K_t(\mathbf{x}_1, \mathbf{x}_2) = K(\mathbf{x}_1, t, \mathbf{x}_2, t)$  depends only on the difference between the points  $\mathbf{x}_1$  and  $\mathbf{x}_2$ :  $K_t(\mathbf{x}_1, \mathbf{x}_1 + \mathbf{h}) = K_t(\mathbf{h})$ . Stationarity does not hold in the real world, and the statistical properties of the seafloor elevation field will vary spatially as parameters like the hydrodynamic forcing, water depth, and geological properties of the seafloor change [23]. The assumption of stationarity in NSEA limits us to modeling patches of seafloor over which these parameters do not change significantly.

A further simplification results from assuming that the mean elevation within the patch of interest is zero. Since the seafloor elevation is always measured relative to an arbitrary vertical datum, this does not result in any loss of information. The covariance function  $K_t(\mathbf{h})$  is then equivalent to the second moment,  $R_2$ .

A wide-sense stationary process has a translation invariant covariance function, but is not completely characterized by the covariance function. Higher-order moments can be important in generating realistic seafloor statistics [24], but NSEA does not explicitly model higher-order moments, and in practice, they are assumed to be zero. The seafloor elevation is thus modeled as a zero-mean Gaussian process, which is completely characterized by its covariance function,  $K_t(\mathbf{h})$ . A Gaussian process has the property that all the finite-dimensional marginal distributions are multivariate Gaussian [22, 25]. This means that if the seafloor elevation is measured at a set of points  $\{\mathbf{x}_i\}_{i=1}^N$ , the distribution of the vector  $[\eta(\mathbf{x}_1, t) \dots \eta(\mathbf{x}_N, t)]$  is a multivariate Gaussian with mean zero and covariance matrix

$$\Sigma_{ij} = K_t(\mathbf{x}_i - \mathbf{x}_j). \quad (5)$$

Stationary processes can alternatively be characterized by the Fourier transform of the covariance function

$$S_t(\mathbf{k}) = \frac{1}{2\pi} \int K_t(\mathbf{h}) e^{-i\mathbf{k} \cdot \mathbf{h}} d\mathbf{h}, \quad (6)$$

which is called the *spectral density*, or *power spectrum*, of the process [25]. The spectral density for a two-dimensional random field has dimensions of length<sup>4</sup> because the covariance function and the area element,  $d\mathbf{h}$ , both have dimensions of length<sup>2</sup>. The spectral density measures the average squared amplitude of

fluctuations in the seafloor elevation that have a spatial frequency given by the wavenumber,  $\mathbf{k}$ . The integral of the spectral density over all wavenumbers is the mean-square surface roughness

$$\langle \eta(\mathbf{x}, t)^2 \rangle = K(\mathbf{0}) = \int S_t(\mathbf{k}) d\mathbf{k}. \quad (7)$$

The covariance function and the power spectrum are equivalent in the sense that they can be converted between each other using the Fourier transform. In spatial statistics, it is more common to model the covariance function by selecting from families of well-understood covariances [25] or by developing models of spatial correlation between neighboring sites such as Markov random fields [26]. These approaches most often make the further assumption of isotropy, that the covariance function depends only on the distance,  $\|\mathbf{h}\|$ , between the measurement points. The resulting random field is rotation- and translation-invariant. Rippled seafloors are clearly anisotropic, with a preferred direction, and oscillations that are not accurately characterized by isotropic covariance functions. It proves somewhat easier to represent anisotropic and oscillating covariances in the spectral domain, which is one reason NSEA models the power spectrum rather than the covariance function. The other major reason to prefer the spectral representation is that valid covariance functions must be positive definite, and this constraint is hard to ensure while evolving the covariance function in time. However, as long as the power spectrum is positive, the corresponding covariance function is necessarily positive definite via Bochner's theorem [25]. NSEA can evolve the power spectrum forward in time in such a way that it is always greater than zero, and the resulting covariance function is always valid.

The covariance function and power spectrum presented above are computed at a single time. There are equivalent spatiotemporal covariance functions and spectra that characterize the correlation of the seafloor elevation at unequal times. However, modeling these explicitly introduces additional complexity [27], and the spatiotemporal statistical properties of real seafloors are not well understood. NSEA ignores the temporal correlations in the seafloor dynamics, assuming that observations of the seafloor separated by times much longer than the characteristic time scales for the forcing, such as the wave period, will be uncorrelated.

Modeling the seafloor elevation as a stochastic process implies that there is some source of randomness in this system. However, the evolution of seafloor roughness is governed by deterministic physical processes of hydrodynamics and sediment transport. A natural question is what the probability distribution over the seafloor elevation actually represents. Deterministic physics are not resolved at the spatial and temporal resolutions necessary for a complete, deterministic prescription. Forcing information is necessarily averaged over some spatial and temporal scales. For example, NSEA might be driven with a wave model such as SWAN [28]. The output of SWAN is the power spectrum of the sea surface height within a grid cell of the model. Like with NSEA, the actual spatiotemporal evolution of the sea surface height is replaced by a statistical description. There are therefore many potential realizations of the sea surface height within a grid cell consistent with the modeled wave spectrum, and each such realization would drive the evolution of the seafloor in a slightly different way. A statistical description of the forcing therefore requires a statistical description of the seafloor evolution. This is especially true if the seafloor evolution model is forced by reduced statistics of the forcing such as the significant wave height and period, which contain even less information about the detailed hydrodynamic forcing than the power spectrum.

Furthermore, even if one could completely resolve the sea surface height, waves drive sediment transport through a turbulent bottom boundary layer that exhibits deterministic but chaotic dynamics [29]. Small



perturbations in the flow within the bottom boundary layer can lead to very different patterns of sediment transport even under a fully specified wave forcing. The statistical description of seafloor evolution employed by NSEA lumps the uncertainty in the chaotic bottom boundary layer dynamics in with the uncertainty in the forcing. The randomness in the seafloor elevation field ultimately comes from the influence of these unresolved processes. The averaging operator  $\langle \cdot \rangle$  used in the definition of the moments of the seafloor elevation field above represents an average over the random forcing and the chaotic turbulent dynamics.

## 2. SYSTEM ARCHITECTURE

NSEA consists of four modules: forcing, seafloor evolution, seafloor synthesis, and observations. Information naturally flows from one module to the next, so that the seafloor evolution module consumes forcing information and its outputs are, in turn, consumed by the seafloor synthesis module. However, NSEA also enables information to flow in the reverse direction in two ways. First, because NSEA represents the evolution of the seafloor in time, outputs from one module can be fed back into a previous model. The typical use case for this kind of coupling in time would be to use the roughness information output from the seafloor evolution model to parameterize bottom friction in a hydrodynamic model that supplies the forcing. Each module's inputs and outputs are directly accessible to the user to facilitate feedbacks of this nature. The second way to share information in the reverse direction is to use techniques of statistical inference to estimate parameters and state variables of each module from available data. For example, observations of bedforms can constrain the seafloor elevation field or its power spectrum, which are normally simulated by the seafloor synthesis and evolution modules. NSEA requires that each module be constructed as a probabilistic model to enable these techniques, which are discussed in more detail in Section 3.

### 2.1 Forcing

The forcing information required by NSEA consists of time series of wave and/or current forcing. Wave forcing consists of a bottom orbital velocity,  $u_w$  [m · s<sup>-1</sup>], a bottom semi-orbital excursion,  $A_w$  [m], and a wave direction,  $\varphi_w$  [rad], while current forcing consists of a depth-averaged current velocity,  $u_c$  [m · s<sup>-1</sup>], and a current direction,  $\varphi_c$  [rad]. These time series can be sampled at any interval including irregular sampling intervals, but the seafloor evolution model assumes that the forcing is constant over the sampling intervals. It is important to provide forcing information at a high enough resolution to capture the desired scales of temporal variability. NSEA is agnostic as to the source of the forcing information, as long as it is supplied as a time series of the required forcing variables. Forcing data can therefore come from in situ or remotely sensed observations or from hydrodynamic model output.

The predictions and statistical inferences about the seafloor roughness that NSEA makes are always conditional on the supplied forcing information. If the user wishes to propagate uncertainty in the forcing through the model or integrate over the forcing uncertainty in performing statistical inferences, they must run ensemble simulations of NSEA in which each ensemble member is driven by a different forcing.

Wave bottom orbital velocities and excursions are often not directly observed in the field. They can be calculated from surface wave statistics by using linear wave theory. For a monochromatic wave with amplitude  $H$  and period  $T$  in a water depth  $h$ , the wave bottom orbital velocity is [30]

$$u = \frac{H\pi}{T \sinh kh}, \quad (8)$$

and the wave semi-orbital excursion is given by  $A_w = uT/2\pi$ . The wavenumber  $k$  is determined from the dispersion relation

$$\frac{4\pi^2}{T^2} = gk \tanh kh, \quad (9)$$

which must be solved iteratively for  $k$  [30].

For a continuous spectrum of surface waves, there is some ambiguity in the choice of wave bottom orbital velocity. The root-mean-square wave bottom orbital velocity is given by [28]

$$u_{RMS}^2 = \int_0^{2\pi} \int_0^\infty \frac{\sigma^2}{\sinh^2 kh} E(\sigma, \theta) d\sigma d\theta, \quad (10)$$

where  $E(\sigma, \theta)$  is the directional wave power spectral density with relative frequency,  $\sigma$ , and direction,  $\theta$ . For a monochromatic wave,  $u = \sqrt{2}u_{RMS}$ . NSEA expects wave bottom orbital velocity to be provided as  $u_{1/3}$ , the velocity derived from Eq. (8) with a height equal to the significant wave height, but it does not enforce this condition, and it is up to the user to provide the appropriate wave bottom orbital velocities and excursions.

## 2.2 Seafloor Evolution

The seafloor evolution model, [14, 18, 19], approximates the evolution of the amplitude spectrum of the seafloor elevation ( $\sqrt{S}$  [m<sup>2</sup>]):

$$\frac{\partial \sqrt{S_t(k)}}{\partial t} = \frac{1}{T_t(k)} \left( \sqrt{\bar{S}_t(k)} - \sqrt{S_t(k)} \right). \quad (11)$$

Equation (11) models the relaxation of the amplitude spectrum toward its equilibrium value,  $\sqrt{\bar{S}_t(k)}$ , with a time scale of  $T_t(k)$ . The forcing information enters the seafloor evolution model through the relaxation time scale and the equilibrium ripple spectrum. This model can be interpreted as a linearization of a nonlinear evolution equation (e.g., [31]) around the equilibrium ripple spectrum.

Traykovski [14] derives the relaxation time scale from the sediment continuity equation, leading to the result that the relaxation time scale is proportional to the cross-sectional area of the ripple divided by the width-averaged sediment transport rate. For a sinusoidal ripple with a height of  $\eta$  and a wavelength of  $\lambda$ , the cross-sectional area can be approximated by  $A = \eta\lambda/2$ . The width-averaged volumetric sediment transport rate,  $\hat{Q}$ , is scaled by the bed packing density,  $(1 - \phi)$ , where  $\phi$  is the porosity, leading to the relationship

$$T(k) \propto (1 - \phi) \frac{\lambda\eta}{2\hat{Q}} = \frac{1 - \phi}{2\hat{Q}} \alpha \left( \frac{2\pi}{k} \right)^2, \quad (12)$$

where in the second equality, the ripple height is related to the wavelength by a constant steepness,  $\eta = \alpha\lambda$ , and  $\lambda$  is rewritten in terms of the ripple wavenumber,  $\lambda = 2\pi/k$ . The parameter  $\alpha$  represents the typical steepness of ripples. Traykovski gives  $\alpha = 0.16$ , which is generally found to be the case for wave-orbital ripples [11]. However,  $\alpha$  can be combined with the unknown proportionality constant and used as a tuning parameter [19].

The model represents the peak sediment transport rate,  $\hat{Q}$ , as a function of the hydrodynamic forcing. One could use any empirical sediment transport law or a more complex bottom boundary layer model to determine  $\hat{Q}$ . The default transport formula that is used in NSEA is the bedload transport law of Meyer-Peter and Müller [20]:

$$\hat{Q} = \gamma (\theta - \theta_{cr})_+^\xi / (\rho(s-1)gD_{50}), \quad (13)$$

where  $\theta$  is the non-dimensionalized shear stress,  $\theta_{cr}$  is the non-dimensionalized critical shear stress, and the notation  $(x)_+ = \max(x, 0)$ . The empirical parameters  $\gamma$  and  $\xi$  are typically assigned values of 8 and 1.5 in the literature [32].

The equilibrium ripple spectrum,  $\sqrt{S_t(k)}$ , represents the ripple geometry to which the seafloor would evolve under a constant forcing. Following [14, 17–19], the equilibrium power spectrum is modeled as a Gaussian function centered on an equilibrium wavenumber,  $\bar{\mathbf{k}}$ , and with an amplitude proportional to the square of the equilibrium ripple height,  $\bar{\eta}$ :

$$S_t(\mathbf{k}) = \frac{\bar{\eta}^2}{16\pi\sigma^2} e^{-\frac{\|\mathbf{k}-\bar{\mathbf{k}}\|^2}{2\sigma^2}}. \quad (14)$$

The width of the spectrum around the equilibrium wavelength is a parameter  $\sigma$ . This parameter could also vary separately in the x and y directions [18] with appropriate modifications to Eq. (14). The spectrum is scaled so that the mean squared ripple amplitude is  $\bar{\eta}^2/8$ , so that as the spectral width tends to zero, the resulting spectrum models a sinusoidal ripple field with height  $\bar{\eta}$ .

The equilibrium ripple height and the wavenumber (or, equivalently, the wavelength,  $\bar{\lambda} = 2\pi/\|\bar{\mathbf{k}}\|$ ) are obtained by an empirical equilibrium ripple geometry formula [9, 11–13, 17] that predicts the equilibrium ripple geometry as a function of the instantaneous hydrodynamic forcing. NSEA can use a wide variety of ripple predictors, which enables testing of different parameterizations. An additional ripple predictor is included to allow testing the sensitivity of the model to ripple geometry parameterizations. It takes the form of a power law in the mobility number  $\psi = u_w^2/g(S-1)D_{50}$  and the displacement parameter  $\delta = A_w/D_{50}$ :

$$\frac{\bar{\lambda}}{A_w} = e^{\beta_1} \psi^{\beta_2} \delta^{\beta_3} \quad (15)$$

$$\frac{\bar{\eta}}{\lambda} = e^{\chi_1} \psi^{\chi_2} \delta^{\chi_3}. \quad (16)$$

The exponents  $\{\beta_i\}_{i=1}^3$  and  $\{\chi_i\}_{i=1}^3$  are free parameters. Many existing ripple predictors have similar forms to this power law or can be approximated fairly well using the power law (Fig. 1).




Fig. 1—Comparison between the non-dimensionalized ripple height (A) and wavelength (B) for the ripple predictors from the literature and the power law model. The x-axis in each plot shows the ripple geometry predicted by each formula while the y-axis is the geometry given by the power law fitted to the predictions of each ripple formula. Ideally, the predictions would fall along the black 1:1 line.

Under strong forcing conditions, sheet flow conditions are expected to exist, in which case the entire bed is mobilized and any ripples present are washed out. This is accomplished by setting the ripple height to zero and the wavelength to infinity when the nondimensional shear stress is above the washout threshold,  $\theta_{wo}$ . This is accomplished within the model by modifying the equilibrium ripple height and wavelength when above the washout threshold and setting the adjustment time scale to a small number (typically 60 s) so that the spectrum evolves normally, but rapidly decays toward a flat bed.

In the absence of hydrodynamic forcing, it is expected that bioturbation will slowly modify the roughness, leading to a decay of ripples [33]. This is modeled as a diffusive process by adding a term to Eq. (11):

$$\frac{\partial \sqrt{S_t(k)}}{\partial t} = \frac{1}{T_t(k)} \left( \sqrt{\bar{S}_t(k)} - \sqrt{S_t(k)} \right) - Dk^2 \sqrt{S_t(k)}. \quad (17)$$

The diffusion process is always on, regardless of the forcing. Since the diffusion coefficient is typically on the order of  $10^{-9} \text{ m}^2 \cdot \text{s}^{-1}$  [19, 33, 34], it has little impact on the ripple geometry when sediment is actively being transported, but results in a slow decay of ripples during quiescent periods.

### 2.2.1 Alternative Derivation

The spectral evolution equation models the evolution of the amplitude spectrum over a patch of seafloor, not the fine-scale dynamics of the seafloor elevation itself. One weakness of this approach is the difficulty in extending the spectral evolution model to include additional forcings such as combined waves and currents, internal waves, or additional sedimentological processes such as sediment cohesion and bioturbation. Furthermore, the model posed in Eq. (11) is ultimately derived heuristically rather than through a physical model for the flux itself.

Therefore, it is invaluable to explore how fine-scale models for sediment transport lead to statistical roughness models like Eq. (11). It is not generally possible to derive the statistical roughness model corresponding to a particular fine-scale sediment transport model. However, there is one fine-scale model, a stochastic heat equation, that is analytically tractable and that provides some insight into the non-equilibrium spectral evolution equation.

The stochastic heat equation is given by

$$\frac{\partial \eta}{\partial t} = D \frac{\partial^2 \eta}{\partial x^2} + \zeta(x, t). \quad (18)$$

This equation is written as a one-dimensional process for notational simplicity, but the extension to two dimensions is straightforward. Here,  $\zeta(x, t)$  is a Gaussian process with a mean of zero and a correlation function,  $E[\zeta(x, t)\zeta(x', t')] = C(x - x')\delta(t - t')$ . In other words,  $\zeta$  is white noise in time with a stationary spatial correlation function  $C(\cdot)$ . This stochastic partial differential equation can be discretized over a patch of length  $L$  with periodic boundary conditions by a Fourier Galerkin method, expanding  $\eta(x, t) = \sum_k \eta_k(t) e^{-i2\pi kx/L}$  and  $\zeta(x, t) = \sum_k \zeta_k(t) e^{-i2\pi kx/L}$ . Substituting these expressions, multiplying by  $e^{i2\pi kx/L}$ , and integrating over the patch, the evolution equation becomes, when written in the customary notation for stochastic differential equations,

$$d\eta_k = -Dk^2 \eta_k dt + \sigma_k dW_k, \quad (19)$$

where the variance term,  $\sigma_k^2$ , is the Fourier transform of the spatial covariance function,  $C$ :

$$\sigma_k^2 = \int C(x) e^{i2\pi kx/L} dx. \quad (20)$$

The noise term,  $W_k$ , is a complex circular Brownian process (i.e.,  $W_k = U_k/\sqrt{2} + iV_k/\sqrt{2}$ ) with  $U_k$  and  $V_k$  as independent real-valued Brownian motions. The Fourier coefficients,  $\eta_k$ , are complex variables; however, the real and imaginary parts evolve independently because the coefficients  $\alpha_k = -Dk^2$  and  $\sigma_k$  are real, and the complex white noise term consists of independently sampled real and imaginary components. The complex Fourier coefficients,  $\eta_k = a_k + ib_k$ , and  $a_k$  and  $b_k$  follow the stochastic differential equations

$$da_k = -\alpha_k a_k dt + \frac{\sigma_k}{\sqrt{2}} dU_k \quad (21)$$

$$db_k = -\alpha_k b_k dt + \frac{\sigma_k}{\sqrt{2}} dV_k. \quad (22)$$

The power spectrum of the seafloor elevation in this case is given by the expected value of the squared magnitude of the Fourier coefficients,  $S_k = E[|\eta_k|^2] = E[a_k^2 + b_k^2]$ . Performing a change of variables using Itô's formula [35] gives the stochastic differential equation satisfied by  $c_k = a_k^2 + b_k^2$ ,

$$dc_k = 2\alpha_k \left( \frac{\sigma_k^2}{2\alpha_k} - c_k \right) dt + \sigma_k \sqrt{2c_k} d\tilde{W}_k. \quad (23)$$

Equation (23) is a version of the analytically tractable Cox-Ingersoll-Ross model [36], and a deterministic evolution equation given  $E[c_k] = S_k$  is

$$\frac{dS_k}{dt} = 2\alpha_k \left( \frac{\sigma_k^2}{2\alpha_k} - S_k \right). \quad (24)$$

Identifying  $\bar{S}_k = \sigma_k^2/2\alpha_k$  and  $T_k = 1/2\alpha_k$ , the spectral evolution implied by the stochastic heat equation has a form similar to that of NSEA:

$$\frac{dS_k}{dt} = \frac{1}{T_k} (\bar{S}_k - S_k); \quad (25)$$

however, it is the power spectrum that evolves according to a linear relaxation process rather than the amplitude spectrum.

The time scale parameter is related to  $\alpha_k$ :

$$\frac{1}{2\alpha_k} = \frac{1}{2Dk^2} = T_k = \frac{1-\varphi}{2\hat{Q}}\alpha\left(\frac{2\pi}{k}\right)^2, \quad (26)$$

and one can observe that the  $k^2$  dependency of  $\alpha_k$  matches the  $k^{-2}$  dependency of the relaxation time scale according to Traykovski's derivation. The diffusion coefficient,  $D$ , can be described as a function of the scaled volumetric sediment transport rate:

$$D = \frac{\hat{Q}}{4\pi^2\alpha(1-\varphi)}. \quad (27)$$

The equilibrium spectrum, which is obtained through the equilibrium ripple predictor, can easily be converted into the forcing variance,  $\sigma_k^2 = 2\bar{S}_k\alpha_k$ .

This derivation illustrates that a very simple fine-scale sediment transport model consisting of a diffusive relaxation and a random flux of sediment with a spatial pattern determined by the wave forcing leads to a statistical roughness model with dynamics qualitatively similar to that used in NSEA. However, it does suggest that the power spectrum should be evolved forward in time according to Eq. (25) rather than the amplitude spectrum as in Eq. (11). Both of these formulations are models of complex sediment dynamics, and it is not possible to determine from first principles whether the heuristic scaling with the amplitude spectrum used to develop Eq. (11) or the stochastic sediment flux model used to develop Eq. (25) is more appropriate. At present, NSEA uses the amplitude spectrum equation Eq. (11). Further work is needed to compare the predictions of these and other statistical roughness models to field and laboratory data, and the modular nature of NSEA allows different evolution models to be substituted for one another quite easily.

### 2.3 Seafloor Synthesis

The output of the seafloor evolution model is the amplitude spectrum, which characterizes the statistical properties of the random seafloor elevation field (Section 1.1). Useful statistics of seafloor roughness such as the root-mean-square elevation or the peak ripple wavelength can be computed from the spectrum. Many applications, such as estimating bottom friction, use these statistics directly. However, other applications, such as acoustic modeling of seafloor scattering (e.g., [37, 38]), require realizations of the random seafloor elevation field. The seafloor synthesis component of NSEA provides these realizations through a Fourier synthesis approach [37, 39].

The realizations cover a two-dimensional patch with side lengths  $(L_x, L_y)$ . The assumption of periodic boundary conditions enables the representation of any seafloor elevation field with a Fourier series. The Fourier decomposition of the seafloor elevation is

$$\eta(\mathbf{x}, t) = \sum_{m,n} \eta_{m,n}(t) e^{i(2\pi m x_1/L_x + 2\pi n x_2/L_y)}, \quad (28)$$

where  $\eta_{mn}(t) \in \mathbb{C}$  are the complex-valued Fourier series coefficients of the elevation at time  $t$ . The elevation field is real-valued; therefore, the Fourier coefficients obey a complex conjugate symmetry condition,  $\eta_{m,n} = \eta_{-m,-n}^*$ . The sum in Eq. (28) is over an infinite set of frequencies. To make the synthesis computationally tractable, the series is truncated at  $M$  terms in the  $x$  direction and  $N$  terms in the  $y$  direction, leaving

$$\eta(\mathbf{x}, t) = \sum_{m=-M/2}^{M/2-1} \sum_{n=-N/2}^{N/2-1} \eta_{m,n}(t) e^{i(2\pi m x_1/L_x + 2\pi n x_2/L_y)}. \quad (29)$$

Using Eq. (29), the elevation can be evaluated at any point within the patch with a cost of  $O(MN)$ . The Fourier series can be evaluated more efficiently on an  $M \times N$  rectangular grid using the fast Fourier transform (FFT). The FFT evaluates Eq. (29) at the points  $\mathbf{x}_{jk} = [L_x j/M, L_y k/N]$  for  $j \in [0, 1, \dots, M-1]$  and  $k \in [0, 1, \dots, N-1]$ . FFT libraries such as FFTW [40] refer to Eq. (29) as the ‘‘inverse fast Fourier transform,’’ which is computed in Julia or MATLAB as `ifft(eta)`, where the  $M \times N$  complex-valued array, `eta`, stores the Fourier coefficients.

The Karhunen-Loève theorem [41] represents a random field as a randomly weighted sum of basis functions. In other words, a random field model for the seafloor elevation can be constructed using Eq. (29) with random Fourier coefficients  $\eta_{m,n}$ . If the Fourier coefficients are modeled as independent, complex-normally distributed random variables with a mean of zero and a variance equal to the value of the power spectrum at the corresponding Fourier frequency, then the random field synthesized by Eq. (29) will be a stationary Gaussian random field with a covariance function given by the Fourier transform of the power spectrum. The procedure used by the seafloor synthesis module of NSEA to synthesize gridded seafloor representations is summarized as:

---

**Algorithm 1:** Seafloor synthesis algorithm

---

**Input:**  $S_{mn}$ : Power spectrum defined at the Fourier frequencies

**Output:**  $\eta_{jk}$ : A random seafloor realization on the equally spaced  $M \times N$  grid

$\zeta_{mn} \sim CN(0, 1)$ ;

$$\eta_{jk} = \sum_{m=-M/2}^{M/2-1} \sum_{n=-N/2}^{N/2-1} \sqrt{S_{mn}} \zeta_{mn} e^{i2\pi m j/M + 2\pi k/N};$$


---

The final line is implemented in practice using a fast Fourier transform. Details on the the implementation of the fast Fourier transform and the scaling and arrangement of its coefficients are discussed in Section 4.3.

This algorithm is efficient because of its use of the Fourier transform; however, there are several drawbacks to this method. First, the generated random fields are periodic, with periods  $L_x$  and  $L_y$  in the  $x$  and  $y$  directions. This periodicity induces spurious correlations between pixels on opposite sides of the patch. The simplest way to mitigate spurious correlations is to generate realizations over a patch at least twice as large as



that needed to contain all of the observations [39]. For example, the roughness observations described further in Section 5 are acquired with rotary scanning sonars with a range of 6 meters, so a patch with a side length of 12 meters centered on the instrument location contains all of the observations. However, setting  $L_x$  and  $L_y$  to 12 meters will result in spurious correlations at the edges of the patch, so the patch side lengths should be set to at least 24 meters. There are still spurious correlations between the edges of the simulated grid, but the correlations will not affect the region of the grid covered by the observations. Simulating on a larger grid is similar to the circulant embedding method [42, 43] used to simulate Gaussian processes on grids from a given covariance function. However, the circulant embedding method is more complex than the spectral synthesis algorithm, largely because it must modify the covariance function to ensure that it is nonnegative definite [44, 45]. A nonnegative definite covariance function corresponds to a power spectrum that is always nonnegative. Because the power spectrum is modeled directly, it is always nonnegative, and the additional complexity of circulant embedding techniques is not necessary.

The second drawback to the Fourier synthesis approach arises when observations are available on non-Cartesian grids. Rotary scanning sonars, for example, provide data on a polar grid defined by constant spacing in range and azimuth. Equation (29) could be evaluated directly on the polar grid, but the computational efficiency of the fast Fourier transform is lost. To handle non-uniform observations, NSEA generates seafloor realizations on a uniform grid and then interpolates below the grid scale. Interpolation smoothes the elevation field at subgrid scales, so that the power spectrum of the interpolated field will not be identical to that of the true field. This discrepancy diminishes as the resolution of the grid is increased. As long as the resolved grid meets the Nyquist criterion of the dominant ripple field, inferences about the roughness should not be strongly affected by smoothing at high frequencies. Three interpolation schemes are used within NSEA (Fig. 2). The selection of which scheme to employ depends on the model application. Nearest-neighbor interpolation [39] represents the seafloor with flat pixels defined by the grid (Fig. 2a). This method is computationally efficient, but the resulting model of the seafloor is not continuous. Bilinear interpolation linearly interpolates along the x-axis and then along the y-axis, resulting in a continuous elevation model (Fig. 2b). It is more computationally intensive than nearest-neighbor interpolation. However, bilinear interpolation is implemented natively in the texture filtering hardware of graphics processing units, which makes it competitive with nearest-neighbor interpolation at the cost of requiring more complex code to control the GPU. While the elevation is continuous using bilinear interpolation, the spatial derivative of elevation (i.e., seafloor slope) is not. Continuous spatial derivatives are especially useful when working with acoustic scattering models that depend on the local surface slope. Bicubic interpolation can be used in this case, providing continuous first and second derivatives of the elevation field (Fig. 2c). Bicubic interpolation can also be implemented efficiently using texture filtering functions on the GPU [46]. For qualitative comparisons between seafloor realizations and non-Cartesian data, nearest-neighbor interpolation suffices and can be computed efficiently on the CPU. For further modeling using the seafloor realizations, bicubic interpolation is recommended because of its continuous derivatives, but requires a GPU for efficient computation. An alternative to bicubic interpolation is to compute the desired derivatives of the seafloor elevation field using spectral methods [47] and to interpolate the derivatives using nearest neighbors. This method can be very efficient, but does not result in a consistent model of seafloor elevation and may cause spurious artifacts when used in downstream modeling efforts.




Fig. 2—Subgrid-scale interpolation schemes used within NSEA. The black dots represent the original gridded elevation measurements and are identical in each of the three plots. A. Nearest-neighbor interpolation, B. bilinear interpolation, C. bicubic interpolation

## 2.4 Observations

The seafloor elevation field is never observed directly. It is only ever observed through indirect and noisy measurements. Extracting useful information about seafloor roughness and its evolution from these noisy and indirect observations requires a probabilistic observation model that connects the available observations to the seafloor elevation field. Statistical inference uses observations, the observation model, and the seafloor evolution model to estimate the seafloor elevation, statistical properties of the elevation field (e.g., power spectrum), or model parameters from the observations.

To enable statistical inference of seafloor properties from observations, the observation model is formally specified as a conditional probability distribution  $\mathbf{z} \sim p_{\theta}(\mathbf{z}|\eta)$  for the observations,  $\mathbf{z} = [\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_N]$ , given the realization of the seafloor,  $\eta$ . The distribution is parameterized by a collection of parameters  $\theta$  that represent characteristics of the measurement system, such as the signal-to-noise ratio and the point spread function, or geophysical parameters, such as the sediment grain size.

While the observation model is a conditional probability distribution, it is not always possible to evaluate the corresponding density function,  $p_{\theta}(\mathbf{z}|\eta)$ , specifically when the observations depend on additional *latent*, or unobserved, variables,  $\mathbf{u}$ , in addition to the seafloor elevation. The observation probability density function is then determined by integrating the joint distribution of the observed and latent variables over the latent variables:

$$p_{\theta}(\mathbf{z}|\eta) = \int p_{\theta}(\mathbf{z}, \mathbf{u}|\eta) d\mathbf{u}. \quad (30)$$

With the exception of a few special cases, this integral cannot be evaluated analytically, and numerical approximations must be made. However, random samples from the observation distribution can often be drawn easily from this model by first sampling the latent variables,  $\mathbf{u}$ , from the probability distribution,

$p_\theta(\mathbf{u}|\eta)$ , and then sampling the observed variables from the distribution  $p_\theta(\mathbf{z}|\mathbf{u}, \eta)$ . Such a *generative* model enables statistical inference in complex models that do not have tractable likelihoods using specialized inference techniques [48]. Complex multiphysics simulations can therefore be treated within the same statistical framework as simpler, analytically tractable statistical models.

NSEA separates the observation model from the seafloor evolution model to enable many different kinds of observations to be used in estimating properties of the seafloor roughness. As long as the observation model can be represented in the form of a conditional probability distribution described above, statistical inference proceeds identically. Users who work with particular observation types can use NSEA to develop and test their own observation models with realistic seafloor roughness information provided by NSEA's seafloor evolution and synthesis models.

Two observation models built into the reference implementation of NSEA are described in Sections 2.4.1 and 2.4.2. The first model is a simplified observation model appropriate for use with observations from profiling sonars, stereophotography, structure-from-motion, or other measurement systems that return high-resolution bathymetric elevation information. This simplified model illustrates the considerations that go into designing an observation model. The second is an imaging sonar model used to compare acoustic return data from rotary scanning sonars to the output of the seafloor evolution model. The imaging sonar model is not an accurate seafloor acoustic model, but it does enable qualitative comparisons between observations and model outputs, and it demonstrates how more sophisticated acoustic models could be incorporated into NSEA.

#### 2.4.1 A Simplified Observation Model

Profiling sonars are often used to make high-resolution elevation maps in sediment transport studies [14, 49]. These sonars have very narrow beams in both azimuth and elevation, so the ensonified patch of seafloor is relatively small, especially at short ranges. The measurements recorded by the profiling sonar are a time series of backscattered signal strength for each sonar pulse at a given azimuth and elevation. The time axis records the two-way travel time of the transmitted pulse from the sonar to the seafloor, where it is then reflected back to the sonar. The profiling sonar identifies the seafloor by selecting the time of the first strong return in the time series. Knowing the speed of sound in water,  $c$ , the two-way travel time,  $t$ , can be converted to range from the sonar,  $R$ , by the formula  $R = ct/2$ . The processed data from the profiling sonar then consists of a set of azimuth ( $\theta$ ) and elevation ( $\varphi$ ) angles and the range of the first return from the sonar,  $(\theta_i, \varphi_j, R_{ij})$ . Standard trigonometry converts these polar coordinates into Cartesian coordinates, producing a collection of horizontal points and corresponding seafloor elevation,  $(\mathbf{x}_{ij}, \eta_{ij})$ .

$$z_i = \int h(\mathbf{x}_i, \mathbf{x}')\eta(\mathbf{x}') d\mathbf{x}' = (h * \eta). \quad (31)$$

A true point observation of the field has a point spread function equal to a Dirac delta function  $h(\mathbf{x}, \mathbf{x}') = \delta(\mathbf{x} - \mathbf{x}')$ . A more realistic point spread function with a finite width will effectively blur the seafloor elevation field, smoothing out high-frequency variability. This indirect measurement model may not be significant when working with a narrow-beamed profiling sonar close to the seabed, where the point spread function will be narrow enough not to smooth away important roughness information, but in other observational scenarios or modalities, it might not be possible to resolve small-scale roughness because of the inherent properties of the measuring device.

The point spread function represents the “known” measurement process in the sense that it is intrinsic to the instrument and is reproducible and deterministic. However, even with perfect knowledge of the instrument’s point spread function, other processes will contaminate the measurements. For example, the first return may represent scattering from organisms in the water column rather than from the seafloor, or the speed of sound might fluctuate with changes in water temperature. The data will always have “noise” in that the observation,  $z_i$ , acquired from the sonar at a point  $\mathbf{x}_i$  is equal to the seafloor elevation plus some residual error:

$$z_i = (h * \eta)(\mathbf{x}_i) + \varepsilon_i. \quad (32)$$

The measurement error terms are unobserved by definition; however, if they were known, they could be corrected in the data-processing pipeline. A common model for the error term is an assumption of independent, normally distributed measurement errors with zero mean and variance,  $\sigma_\varepsilon^2$ , in which case the conditional distribution of the observation,  $z_i$ , given the seafloor elevation,  $\eta_i$ , is also normal:

$$z_i | \eta \sim \mathcal{N}((h * \eta)(\mathbf{x}_i), \sigma_\varepsilon^2). \quad (33)$$

The observation model parameters are the point spread function,  $h$ , and the noise variance,  $\sigma_\varepsilon^2$ . This model is not necessarily an accurate model of the measurement errors in profiling sonar observations, but it does allow for analytical statistical treatment. For example, since NSEA assumes  $\eta$  to be a Gaussian process with a known power spectrum,  $S(k)$ , the Wiener filter [50] optimally estimates the seafloor elevation from the observations. More complicated measurement models could be considered; for example, errors could depend on the location or could be correlated, or error distributions could have heavy tails that would accommodate outliers.

#### 2.4.2 Imaging Sonar Observation Models

Imaging sonars are widely used due to their ability to observe large areas of the seafloor at a high resolution. However, their measurements are very different from those of the profiling sonars described above. Similar to profiling sonar, an imaging sonar ensonifies an area of the seafloor and records a time series of backscattered signal strength. It differs from profiling sonar in the shape of the ensonified area and in the processing of the time series to the final observations. Profiling sonars have narrow beams of a few degrees around a central bore, whereas imaging sonars have wide beams in the elevation direction. Side-scan and rotary scanning sonars have narrow beams in the azimuthal direction to ensonify narrow strips of seafloor. Synthetic aperture sonars use wide azimuthal beams, but process multiple pulses together to synthesize narrow beam widths. The instrument samples the backscattered signal from a single pulse at discrete times, so each sample corresponds to a particular two-way travel time from the instrument. This travel time is converted into range from the sonar using the same  $R = ct/2$  formula in processing profiling sonar measurements, and the output of the instrument is essentially a series of backscattered signal amplitudes at a set of discretely spaced ranges for each pulse. The instrument then moves through the water, so the next pulse ensonifies a different portion of the seafloor. Side-scan sonars move forward, mapping out strips on either side of the instrument, while rotary-scanning sonars rotate in the azimuthal direction and map an annulus around the instrument position. A complete scan consists of a stack of backscattered amplitude time series. These stacks can then be projected into geographic coordinates using knowledge of the instrument’s location to produce a sonar image.

Sonar images are not maps of seafloor elevation, yet bathymetric features like ripples in sonar imagery can be observed because the backscattered amplitude depends on the orientation of the seafloor surface relative

to the instrument. The simplest model for the angular dependence of the scattering strength is given by Lambert's law for diffuse scattering, which roughly states that the acoustic energy impinging on the seafloor is scattered equally in all directions [3, 51]. Taking into account the relative orientation of the seafloor surface to the instrument gives the well-known cosine-squared form of Lambert's law for the backscattering cross section,  $\sigma$ ,

$$\sigma = \frac{\rho}{\pi} \cos^2 \theta_i, \quad (34)$$

where  $\rho$  is the proportion of incident energy that is scattered rather than absorbed by the seafloor, and  $\theta_i$  is the incident angle relative to the surface normal. This model is often given in the acoustic literature with the grazing angle,  $\theta_g$ , relative to the seafloor surface, in which case  $\sigma \propto \sin^2 \theta_g$ .

The cosine of the incident angle can be written as  $\cos \theta_i = \hat{\mathbf{n}} \cdot \hat{\mathbf{L}}$ , where  $\hat{\mathbf{n}}$  is the unit surface normal vector and  $\hat{\mathbf{L}}$  is the unit vector pointing from the surface to the sonar. If the seafloor surface is given as  $\eta(\mathbf{x})$  and the sonar location is  $\mathbf{o}$ , then the surface normal vector is given by

$$\mathbf{n} = \left[ -\frac{\partial \eta}{\partial x}, -\frac{\partial \eta}{\partial y}, 1 \right] \quad (35)$$

and the observation direction vector is  $\mathbf{L} = [o_x - x, o_y - y, o_z - \eta(x, y)]$ . The unit vectors are obtained by normalizing these vectors.

Lambert's law can form the basis of a simplified imaging sonar observation model within NSEA, as it can simulate the backscattering cross section that would be obtained from a given seafloor roughness realization. For rotary scanning sonars, the natural range and azimuth coordinates of the acquired data can be converted into Cartesian surface coordinates with trigonometry, assuming that the amplitude of seafloor roughness is small relative to the height of the sensor above the seafloor. The required seafloor elevation field and its derivatives can be computed efficiently with the fast Fourier transform and filtering procedures outlined in Section 2.3. Two further points must be considered to turn the Lambertian scattering model into a valid observation model for use in NSEA. First, instrument-specific effects such as the sonar transmit and receive beam patterns create additional variability in the image not due to the seafloor roughness [52, 53]. These effects, however, affect the image only in the downrange, or across-track, direction, so they can be modeled as a range-dependent amplitude,  $B(r)$ , which is multiplied by the backscattering cross section corresponding to the given range and azimuth angle. The image is then formed by the equation,

$$Y(r, \theta) = B(r)\sigma(r, \theta), \quad (36)$$

where  $\sigma(r, \theta)$  is the scattering cross section of the portion of the seafloor surface that corresponds to the given range,  $r$ , and azimuth,  $\theta$ .

The imaging sonar model finally needs to be written probabilistically to fit into NSEA's statistical inference framework. There are many models for the noise statistics of acoustic backscattering from the seafloor [3, 54–56], but the simplest can be derived by assuming that the received sonar signal consists of contributions from a very large number of small scatterers. The central limit theorem suggests that the received signal will be approximately complex normally distributed, so the signal amplitude is Rayleigh distributed [54], and the probabilistic observation model can be written as  $Y(r, \theta) \sim \text{Rayleigh}(B(r)\sigma(r, \theta))$ . The parameters of this

observation model are contained entirely within the range-dependent amplitude function,  $B(r)$ , which must be estimated from the data.

The Lambertian model is not necessarily a model of seafloor acoustics [51, 57], and inferences on seafloor roughness from this observation model should be subjected to extensive testing and validation. However, it does enable qualitative comparisons between sonar imagery and the output of the seafloor evolution model that account for the characteristics of imaging sonar observations, namely the directional filtering of seafloor roughness [51].

### 3. STATISTICAL INFERENCE

The previous section has described NSEA as a set of coupled models: forcing, seafloor evolution, seafloor synthesis, and observations. Each of these models can be described probabilistically, defining a probability distribution over its output conditional on its input. NSEA therefore takes the form of a hierarchical model:

$$\text{parameters} \sim p_{\text{parameters}}(\cdot) \quad (37)$$

$$\text{forcing} \sim p_{\text{forcing}}(\cdot | \text{parameters}) \quad (38)$$

$$\text{spectrum} \sim p_{\text{NSEM}}(\cdot | \text{parameters}, \text{forcing}) \quad (39)$$

$$\text{seafloor} \sim p_{\text{synthesis}}(\cdot | \text{spectrum}, \text{parameters}) \quad (40)$$

$$\text{observations} \sim p_{\text{obs}}(\cdot | \text{seafloor}, \text{forcing}, \text{parameters}). \quad (41)$$

A few examples of inferences one might make using NSEA are:

- estimation of properties of the seafloor (e.g., grain size) from sidescan sonar imagery,
- comparison of multiple formulations of sediment transport processes within the seafloor evolution model to assess their validity, and
- forecast of the evolution of seafloor roughness at a particular location using forcing conditions predicted by a wave model.

These examples can be categorized by three types of statistical inferences: estimation, model comparison, and prediction, respectively. Their commonalities include the use of data to estimate unobserved quantities of interest and the quantification of the uncertainty in those estimates. Bayesian statistics provide a framework for making these inferences in hierarchical models. The goal of Bayesian inference is to use the observed variables to compute the *posterior distribution* of the unobserved or latent variables. The posterior distribution can be obtained formally with Bayes's rule, which states that the posterior is proportional to the *likelihood* (the observation model) times the *prior* (the model for the latent variables). For example, the posterior distribution for the seafloor elevation field could be written as

$$p(\text{seafloor} | \text{observations}) = \frac{p(\text{observations} | \text{seafloor})p(\text{seafloor})}{p(\text{observations})}, \quad (42)$$

where the likelihood is  $p(\text{observations} | \text{seafloor})$  and the prior is  $p(\text{seafloor})$ . Note, however, that the likelihood,  $p(\text{observations} | \text{seafloor})$ , the prior,  $p(\text{seafloor})$ , and the normalizing constant  $p(\text{observations})$

(also called the *marginal likelihood*) do not appear directly in the hierarchical model, Eq. (37). They must be obtained by integrating or *marginalizing* over the other variables in the model: in this case, the parameters, forcing and seafloor spectrum. For example, the prior distribution of the seafloor elevation is formally given by

$$p(\text{seafloor}) = \int p(\text{seafloor}|\text{spectrum}, \text{parameters})p(\text{spectrum}|\text{forcing}, \text{parameters}) \\ p(\text{forcing}|\text{parameters})p(\text{parameters})d\text{spectrum}d\text{forcing}d\text{parameters}. \quad (43)$$

These integrals will not generally have analytical solutions, and they are very high-dimensional, so calculating them using numerical approaches like quadrature scale very poorly. Bayesian statistics offers a wide variety of algorithms for estimating integrals like Eq. (43). The choice of algorithm depends on the kinds of observations that are available, the posterior distributions that are of interest, and the simplifications that the user is willing make to the complete hierarchical model. One algorithm, approximate Bayesian computation [58], that is implemented in NSEA is described.

### 3.1 Approximate Bayesian Computation in NSEA

The inference task considered in this section is to estimate the parameters of the seafloor evolution model from observations of ripple wavelength that have been extracted from sonar imagery. This section describes the methodology, and results from a test of this methodology are presented in Section 5. The Bayesian inference problem is to estimate the posterior distribution of the parameters given the wavelength observations

$$p(\theta|\hat{\lambda}_{1:T}, u_{1:T}, A_{1:T}) = \frac{p(\hat{\lambda}_{1:T}|\theta, u_{1:T}, A_{1:T})p(\theta)}{p(\lambda_{1:T})}, \quad (44)$$

where the parameters are collected in the vector,  $\theta$ , and the observations consist of a time series of estimated ripple wavelengths,  $\hat{\lambda}_{1:T}$ . The case presented in Section 3 uses a fast Fourier transform to process the sonar images and to estimate the peak ripple wavelength [19], but other image processing techniques [14, 59] could also be employed. The forcing information is supplied as a time series of bottom orbital velocity,  $u_{1:T}$ , and semiorbital excursion,  $A_{1:T}$ . The prior distribution for the parameters,  $p(\theta)$ , is ultimately chosen by the user based on assumptions about the ranges of the parameters.

The likelihood,  $p(\hat{\lambda}_{1:T}|\theta, u_{1:T}, A_{1:T})$ , is the probability distribution over the wavelength data induced by the hierarchical model. However, this likelihood cannot be analytically computed from Eqs. (37)–(41) for two reasons. First, it requires a high-dimensional integral over the seafloor realization that cannot be computed analytically. It is possible to estimate this integral using traditional Bayesian inference algorithms like Markov chain Monte Carlo. The more challenging problem is that the morphological processing required to estimate ripple wavelengths from the sonar imagery applies a complicated, nonlinear, and non-invertible function to the images. This additional processing step does not allow for the computation of the probability distribution for the estimated ripple wavelengths conditioned on the seafloor realization,  $p(\hat{\lambda}|\eta)$ , which is needed for Monte Carlo sampling of the seafloor realization.

Approximate Bayesian computation (ABC) [58] enables Bayesian inference for models with intractable likelihoods by replacing direct computation of the likelihood with a simulation. In the context of the parameter inference problem given in Eq. (44), ABC consists of sampling parameter values from the prior distribution,  $p(\theta)$ , running the model forward to simulate ripple wavelengths, and then comparing the

simulated ripple wavelengths with the observed ripple wavelength by calculating a discrepancy measure such as the mean squared difference between the simulated and observed ripple wavelengths. Parameter values with a corresponding discrepancy measure less than a user-defined tolerance are accepted by the algorithm, and the posterior distribution is approximated by the accepted parameter values. Approximate Bayesian computation is exact in the limit as the discrepancy tolerance goes to zero, but a tolerance of zero rejects almost all of the sampled parameter values, which means more samples must be taken to reduce the statistical variance of the posterior estimate. Conversely, a tolerance set too high will accept all of the sampled parameters, and the resulting posterior will more closely resemble the prior distribution than the true posterior distribution. Choosing the tolerance requires making a trade-off between the bias of a too-high tolerance and the variance of a too-low tolerance. A pilot simulation with relatively few samples can help determine the acceptance rate as a function of the tolerance. The tolerance is then set to acquire the desired number of accepted samples given the total number of samples that can be run with the available computational budget.

To simulate ripple wavelengths, one would first run the seafloor evolution model with the sampled parameter value to estimate the seafloor spectrum, synthesize a seafloor realization from that spectrum, apply the sonar imagery model to simulate a sonar image consistent with the spectrum, and finally, use the image processing method of choice to estimate ripple wavelengths from the simulated sonar image. However, this complete simulation is computationally intensive and the inferences depend heavily on the accuracy of the sonar imagery model. NSEA currently simulates ripple wavelengths by calculating the peak ripple wavelength from the maximum of the simulated seafloor spectrum and adding noise to the peak wavelength value. Future work is needed to enhance the computational capacity of NSEA and to enhance the realism of the sonar imagery model to allow ABC to use complete simulations from the hierarchical model.

## 4. IMPLEMENTATION

Development of NSEA currently focuses on an implementation written in the Julia language [21] called `Stingray.jl`. `Stingray` provides a core interface for describing the seafloor elevation within a given localized area using spectral methods (`StingrayCore`, Section 4.1), an implementation of the spectral seafloor evolution model (NSEM, Section 4.2), a set of methods for synthesizing seafloor realizations from spectral information (`SyntheticSeafloor`, Section 4.3), and an imaging sonar model for simulating sidescan sonar images from seafloor realizations (`Stingray`, Section 4.5). Each of these components is implemented as a separate Julia package, but are developed within a single Git repository that facilitates installing compatible versions of each of the packages.

### 4.1 `StingrayCore`

The `StingrayCore` module provides a lightweight interface for representing seafloor elevation realizations within a rectangular patch. The module exports a single concrete type:



```

struct Seafloor{T, D,
    A <: AbstractArray{T,D},
    B <: AbstractArray{Complex{T},D},
    P <: AbstractFFTs.Plan{T}}

    p::P
    F::B
    V::B
    eta::A
    kx
    ky
    dims
    L::Tuple{T,T}
end

```

The fields of this type are a Fourier transform plan,  $p$ , as defined in the `AbstractFFTs.jl` interface, fields for the Fourier transform of the elevation,  $F$ , and a filter impulse response,  $V$ , that is used to generate realizations having a given power spectrum. The field  $\eta$  is a real-valued array that contains a gridded seafloor elevation. These three arrays,  $F, V, \eta$ , are not necessary for running the seafloor evolution model, but preallocating them in the `Seafloor` prevents having to reallocate potentially very large arrays when synthesizing seafloor realizations. The geometry of the patch is defined by the fields  $k_x, k_y, \text{dims} = (M, N)$  and  $L = (L_x, L_y)$ , which are the wavenumber grids in the  $x$  and  $y$  directions, the dimensions of the gridded seafloor representation, and the length of the patch sizes in the  $x$  and  $y$  directions in meters. A `Seafloor` object is created by calling the type constructor `Seafloor(u, L)` with a prototype array,  $u$ , that has the same data type and size as the desired gridded seafloor representation, and the tuple  $L = (L_x, L_y)$  of patch side lengths. The prototype array  $u$  is only used to determine the correct sizes of the various internal fields of the `Seafloor` type. The use of the the prototype array allows diverse array types from the Julia ecosystem to be used as storage for the gridded seafloor representation and its Fourier coefficients. However, a user of `Stingray` most likely would use a standard Julia array with an element type of `Float64` and a size  $(M, N)$ , in which case the desired `Seafloor` object can be created as:

```
sf = Seafloor(zeros(M,N), (Lx,Ly)).
```

Two convenience methods are also provided to allow users to create `Seafloor` objects without worrying about the details of prototype array:

```
sf = Seafloor((M,N), (Lx,Ly))
```

is equivalent to

```
sf = Seafloor(zeros(M,N), (Lx,Ly))
```

and

```
sf = Seafloor(N,L)
```

is equivalent to

```
sf = Seafloor(zeros(N,N), (L,L)).
```

The patch and its grid representation have equal lengths and resolutions in the x and y directions using the latter method.

The only restriction on the prototype array is the requirement for it to be an input to the `plan_rfft(u)` to create a FFT plan. For the regular Julia Array, the FFTW.jl bindings to the FFTW library [40] provide this plan. One must load the desired FFT provider before attempting to create a Seafloor object, or the Seafloor command will result in an error. See the below example:

```
# Load in the required Julia packages
using StingrayCore

# Create the seafloor object
M,N=(256,512)
Lx,Ly=(5.0,10.0)

# Load the FFT provider
using FFTW

# Create the seafloor object with the method for rectangular patches
sf = Seafloor((M,N), (Lx,Ly))
```

or

```
# Create the seafloor object with the method for square patches
sf = Seafloor(M,L)
```

The various models can also be run on CUDA-enabled GPUs by using the CuArray type from the CUDA.jl package [60]. One could create a Seafloor on the GPU with `sf_gpu = Seafloor(CUDA.zeros(M,N), (Float32(Lx), Float32(Ly)))`. Note the casting of the patch sizes to Float32, which is necessary because the default element type for a CuArray is a 32-bit floating point number. If Lx and Ly are another type, such as Float64, type instabilities will result, which can dramatically slow down operations on the GPU. The type parameters of the Seafloor enforce these types to match. If one tries to pass a prototype array with Float32 elements and Float64 patch edge lengths, a method error will result. The `Seafloor((M,N), (Lx,Ly))` convenience methods automatically create a Julia Array-based Seafloor object, so the method with the prototype array must be used to create CUDA-enabled Seafloor objects.

For CuArrays, CUDA.jl provides an interface to FFTs through CUFFT. To employ CUDA-enabled GPUs, the following commands can be used:

```
using CUDA
M,N=(256,512)
Lx,Ly=(5f0,10f0)
sf_gpu = Seafloor(CUDA.zeros(M,N),(Lx,Ly))
```

## 4.2 NSEM

The NSEM module implements the spectral seafloor evolution model described in Section 2.2. The module exports a single function `nsem`, which has multiple methods used to run the model with different configurations. The most general method is `nsem(sf::Seafloor, t, uw, Aw, phiw, uc, phic, params)`, which takes a `Seafloor` object, `sf`, a vector of time steps, `t`, given in seconds since an arbitrary time, and similarly sized vectors of bottom wave orbital velocity (`uw` [ $m/s$ ]), semiorbital excursion (`Aw` [ $m$ ]), wave direction (`phiw` [radians counterclockwise from the x-axis]), bottom current velocity (`uc` [ $m/s$ ]), and bottom current direction (`phic` [radians counterclockwise from the x-axis]). The model parameters, `params`, are passed to the model in a `NamedTuple` with the following fields:

```
params = (;
    D50, # Median sediment grain size
    S, # Sediment specific gravity
    theta_cr, # Critical Shields stress for initiation of motion
    wo_diff, # logarithm of the difference between the
              # critical stress for washout and for initiation
              # of motion (e.g., log(0.168-0.05))
    alpha, # Time scale multiplier
    phi, # Porosity
    sigma, # Equilibrium spectral width
    gamma, # Sediment transport rate multiplier
    xi, # Sediment transport rate exponent
    D, # Diffusion coefficient
    Cd, # Drag coefficient
    rp # Ripple predictor
)
```

Most of these parameters are simply real numbers, though care should be taken to ensure that they have the same type as the underlying array type of the `Seafloor` type (e.g., `Float64` if an `Array{Float64}` is used, `Float32` if a `CuArray{Float32}` is used). The washout threshold must always be greater than the critical Shields stress for initiation of motion, `theta_cr`, so the parameter input for the washout threshold is the logarithm of its excess over the critical Shields stress (`wo_diff`). The washout threshold, `theta_wo`, is computed within NSEM as:

```
theta_wo = params.theta_cr + exp(params.wo_diff).
```

The ripple predictor is an object that is a subtype of `NSEM.EquilibriumRipplePredictor`. Each `EquilibriumRipplePredictor` object is callable with a method,

```
(rp::EquilibriumRipplePredictor)(Ab,psi,theta,D50),
```

that takes the semiorbital excursion, `Ab`, the mobility number, `psi`, the Shields stress, `theta`, and the grain size, `D50`, and returns `(eta, lambda)` the equilibrium ripple height and wavelength. Five equilibrium ripple predictors are currently implemented: `Nielsen` [9], `WibergHarris` [11], `Goldstein` [13], `Nelson` [12], and `ParameterizedRipplePredictor`. The first four of these are empirical formulas from the literature. The relevant ripple predictor objects can be created as `Nielsen()`, `WibergHarris()`, etc., and passed as a model parameter in `params`, to `nsem`. The parameterized ripple predictor represents the power law model of Eqs. (15) and (16). A `ParameterizedRipplePredictor` object is created by passing vectors with the exponents for each of these terms, `beta` and `chi`.

The current-generated ripple predictor from Soulsby and Whitehouse [17] has also been implemented (`SoulsbyWhitehouse`). It requires the sediment specific gravity, `S`, which is passed to the object constructor as `rp = SoulsbyWhitehouse(S)`. Under current-dominated conditions, the semiorbital excursion is undefined, and the method for the Soulsby and Whitehouse ripple predictor only takes three parameters,

```
(erp::SoulsbyWhitehouseCurrent)(psi,theta,D50).
```

The Soulsby and Whitehouse predictor depends only on the grain size and the sediment's specific gravity, but the forcing variables `psi` and `theta` are retained to provide compatibility with future ripple predictor implementations that may depend on the strength of the current forcing. When the shear stress due to currents is stronger than that due to waves, the `SoulsbyWhitehouse` ripple predictor is used to compute the equilibrium ripple height and wavelength. The parameter `params.rp` should always be one of the wave-generated ripple predictors, even for current-dominated systems. The `SoulsbyWhitehouse` predictor will automatically be used for current-generated ripples. Future versions of NSEM might allow the specification of alternative current-generated ripple predictors.

See Table 1 for a list of typical values and units of the model parameters.

Table 1—Typical NSEM Model Parameter Values and Units

Parameter	Units	Typical Value
$D_{50}$	m	0.00023
$S$	-	2.65
$\theta_{cr}$	-	0.05
$\Delta\theta_{wo}$	-	$\log(\theta_{wo} - \theta_{cr})$
$\alpha$	-	10.0
$\varphi$	-	0.4
$\sigma$	m	1.0
$\gamma$	-	8.0
$\xi$	-	1.5
$D$	$\text{m}^2 \text{s}^{-1}$	$5.0 \times 10^{-9}$
Ripple predictor	-	<code>Nelson()</code>
Grid size	-	$512 \times 512$
Patch edge length	m	15

The output of:

```
nsem( sf::Seafloor, t, uw, Aw, phiw, uc, phic, params )
```

is a three-dimensional array containing the simulated amplitude spectrum. The first two dimensions are the wavenumber in the x and y directions and have lengths equal to the lengths of `sf.kx` and `sf.ky`, respectively. The third dimension represents the time step and has a length equal to the length of the vector `t`, which should be equal to that of `uw`, `Aw`, etc.

An `nsem` method for running NSEM with only wave-generated ripples is also provided:

```
nsem( sf::Seafloor, t, uw, Aw, phiw, params ).
```

If the system of interest is wave-dominated or current data are not available, one can provide `nsem` with only wave forcing data. The parameters, `params`, are the same as for the full model, except that the drag coefficient, `Cd`, is not required and will be ignored if provided. To ease integration between NSEM and other sources of forcing information, a method,

```
nsem( sf::Seafloor, df, params ),
```

is provided for running the model from forcing data within a `DataFrame` by reading a CSV file containing the forcing information. The columns of the `DataFrame` should be labeled as “`t`”, “`uw`”, “`Aw`”, “`phiw`”, “`uc`”, and “`phic`”, corresponding to the inputs passed to the complete method of `nsem`. The “`t`” column should consist of `DateTime` objects from the Julia standard library `Dates` module that record the observation times.

### 4.3 SyntheticSeafloor

The `SyntheticSeafloor` module implements Fourier synthesis of seafloor elevation realizations consistent with a given amplitude spectrum. The Fourier synthesis algorithm is discussed in Section 2.3. This section discusses the implementation of this algorithm and its use with other Stingray packages. `SyntheticSeafloor` exports two functions,

```
synthesize( sf::Seafloor, v, u )
```

and

```
synthesize!( sf::Seafloor, v, u ).
```

Both functions expect a `Seafloor` object (discussed in Section 4.1) that defines the geometry of the grid (an  $M \times N$  array) and two arrays, `v` and `u`. Array `v` is the desired amplitude spectrum, organized in the pattern expected by FFTW for real-valued Fourier transforms. The output of NSEM, for example, is an array with dimensions  $(\lfloor M/2 \rfloor + 1, N, T)$  containing the amplitude spectrum at each time step of the simulation. To generate a seafloor elevation realization at time index `t` of the simulation, one would pass the slice `v[:, :, t]` to `synthesize`. Array `u` is an  $M \times N$  array of independent standard normal variables, generated, for example, by `randn(M,N)`. Both `synthesize` and `synthesize!` then compute the gridded realization of the seafloor elevation, `eta` as:

```
eta = irfft(v .* rfft(u), sf.dims[1]).
```

The difference between `synthesize` and `synthesize!` is in how this computation is performed internally. `synthesize` implements the seafloor synthesis algorithm using the Fourier transform plan cached in the `sf` object

```
function synthesize(sf::Seafloor,v,u)
    p = sf.p
    p \ (v .* (p * u))
end.
```

This implementation has the downside of allocating memory for the forward Fourier transform of  $u$ , the filter Fourier coefficients, and the gridded seafloor realization. Allocating memory repeatedly inside iterative algorithms such as Markov chains Monte Carlo inference can harm performance; therefore, the `synthesize!` variant is also provided. Following Julia conventions that functions ending in `!` modify their arguments, `synthesize!` performs in-place Fourier transforms using cached arrays that are created when the `Seafloor` object is created.

```
function synthesize!(sf::Seafloor,v,u)
    p = sf.p
    F,eta = sf.F,sf.eta

    # Perform the in-place forward FFT of u
    mul!(F,p,u)

    # Multiply by the amplitude spectrum in place
    F .*= v

    # Perform the in-place inverse FFT of F
    ldiv!(eta,p,F)

    eta
end
```

This method allocates no memory and is suitable for use in high-performance inner loops. However, because it modifies the arrays stored in the `Seafloor` object, it is not thread-safe: if two threads attempt to use `synthesize!` simultaneously on the same `Seafloor` object, the results will not be correct. If parallel processing is desired, either the non-mutating `synthesize` can be used or multiple identical `Seafloor` objects can be created and passed to each processor.

Both versions work for any fast Fourier transform that implements the `AbstractFFTs` interface. CUDA-accelerated seafloor synthesis requires constructing a `Seafloor` object on the GPU as discussed in Section 4.1 and then passing `CuArrays` for  $v$  and  $u$ . The arrays can either be CPU arrays that are copied to the GPU using the `cu` function provided by CUDA (e.g., `synthesize(sf, cu(v), cu(u))`), or they can be native GPU arrays. Random values for  $u$  can be generated on the GPU using `CUDA.randn(M,N)`, while if NSEM is run with a GPU-based `Seafloor` object, the output amplitude spectra will already be on the GPU and can be accessed using the same slice syntax as the CPU version. Other array types should work similarly, but are not regularly tested during Stingray development.

#### 4.4 Demo Code

The following example commands can be used to run the demo code for NSEM (Section 4.2) and the resulting time series of seafloor elevation realizations generated with the SyntheticSeafloor (Section 4.3). The demo code uses a time series of wave (only) forcing to run NSEM and then uses SyntheticSeafloor with the random noise,  $u$ , correlated in time with a correlation coefficient,  $\rho$ .

```
using SyntheticSeafloor, NSEM, FFTW, DataFrames, CSV, DelimitedFiles

# Create Seafloor object
sf=Seafloor((512,512),(5.0,5.0));

# Load in wave forcing data as a DataFrame
df=DataFrame(CSV.File("NSEM_DEMO_DATA.csv"));

# Set model parameters in a Named Tuple,  $\theta$ 
 $\theta=(D50 = 2.3e-4, S = 2.65, \theta_c = 0.05, \Delta\theta_{wo} = \log(0.168-0.05), \alpha = 10.0,$ 
 $\phi = 0.4, \sigma = 1.0, \gamma = 8.0, \xi = 1.5, D = 5e-9, rp = Nielsen())$ 

# Run the NSEM simulation with the seafloor object (sf),
# forcing dataFrame (df), and model parameters ( $\theta$ )
S=nsem(sf,df, $\theta$ );

# Write the amplitude spectrum to a text file
writedlm("output_spectrum.txt",S)

# Set N equal to the length of the grid and t to the length of time
N=size(S,2);
t=size(S,3);

# Correlate random noise in time with correlation coefficient,  $\rho$ 
u=zeros(N,N,t);
rho=0.99;

for i in 2:t
    u[:, :, i] = rho*u[:, :, i-1] + sqrt(1 - rho^2) * randn(N,N)
end

# preallocate elev
elev=zeros(N,N,t);

# Run SyntheticSeafloor on the NSEM output of the amplitude spectrum
# time series
for i in 1:t
    elev[:, :, i]=fourier_synthesize(sf,S[:, :, i],u[:, :, i])
end

# Write the seafloor elevation realization time series to a file
writedlm("output_elev.txt",elev)
```

## 4.5 Stingray

The Stingray module implements the observation models described in Section 2.4. Currently, Stingray handles only the deterministic part of the observation model, leaving the full probabilistic specification of the model to be handled within the statistical inference process. The observation models are defined as types

```
struct PointSpreadModel{I}
  h # The point spread function
  X # The observation locations
end
```

and

```
struct LambertianModel
  h # The measurement height
  B # The beam pattern
  R # The ranges of the observations
  theta # The azimuth angles of the observations
end
```

that can be applied to a gridded seafloor elevation field

```
Y1 = (m::PointSpreadModel)(sf::Seafloor[,eta])
Y2 = (m::LambertianModel)(sf::Seafloor[,eta])
```

to obtain the deterministic response of the observation model. If an additional array `eta` is passed to these methods, the model uses `eta` as the seafloor elevation field. Otherwise, it uses the cached elevation field in the `Seafloor` object.

The `PointSpreadModel` first applies the point spread function to the gridded seafloor elevation field using a fast Fourier transform-based convolution, then interpolates to the observation locations using the interpolation method (`NearestNeighbor`, `Bilinear`, or `Bicubic`) given as the parameter `I` in the type definition. The `LambertianModel` first computes the derivatives of the gridded elevation field at the observation locations by interpolating with a kernel that is the derivative of the bicubic filter. The derivatives are then used to compute the surface normal and the surface is shaded using the Lambertian scattering model, Eq. (34). Finally, the stored beam pattern is applied to the observations, which are returned as an array with rows corresponding to the range bins and columns corresponding to the azimuth angles. Bicubic interpolation is computationally intensive when run on the CPU, so the `LambertianModel` should almost always be applied to a `Seafloor` object and an `eta` defined on the GPU.



## 5. ARCHITECTURE DEMONSTRATION

To illustrate the application of NSEA to a real-world seafloor roughness prediction problem, the model was employed to simulate with data collected during the Target and Reverberation Experiment (TREX) off the coast of Panama City, Florida, in 2013. Instrumentation on two quadpods collected data for 34 days (April 20 – May 23, 2013). The shallower of the quadpods was at a depth of approximately 7.5 m while the deeper was at a depth of 20 m. The present demonstration focused on data from the shallow quadpod. Wave height, period, and direction were recorded by a Nortek AWAC-AST at 2 Hz for 1,024 s every 30 minutes. Linear wave theory was used to estimate the bottom orbital velocity and the semiorbital excursion from the observed significant wave height. The forcing data input to NSEA included the bottom orbital velocity ( $u$ ), wave semiorbital excursion ( $A$ ), and the wave direction ( $\phi$ ) (Fig. 3). A 2.25 MHz sector-scanning sonar (Imagenex 881b) installed at a height of 1.05 m above the seafloor imaged approximately a  $6\text{ m} \times 6\text{ m}$  area under the quadpods every 12 minutes. A fast Fourier transform of a  $3\text{ m} \times 3\text{ m}$  patch of the the backscatter image was computed and the wavenumber corresponding to the maximum Fourier amplitude was extracted and converted into an estimate of the ripple wavelength.




Fig. 3—(A) Bottom orbital velocity, (B) bottom semiorbital excursion, and (C) wave direction ( $0^\circ$  North) at the shallower quadpod during the TREX13 experiment

## 5.1 Forward Simulation

The predictive ability of NSEA was demonstrated by running the seafloor evolution model (NSEM) with the observed forcing and a set of realistic parameters given in Table 2. The median grain size was directly measured (0.23 mm), while the other parameters were set to common values from the literature. The diffusion coefficient is set to zero; therefore, the modeled ripples do not decay in the absence of forcing. The current-generated ripple module was turned off for this run due to the negligible observed currents. NSEM simulates only wave-generated ripples using the Nelson ripple predictor [12]. From the amplitude spectra output by NSEM, the peak wavelength was computed from the  $x$  and  $y$  wavenumbers corresponding to the maximum of the amplitude spectra. Roughness realizations were generated from the amplitude by the seafloor synthesis algorithm, and the simplified Lambertian scattering model is applied to simulate a sonar image from the roughness realizations. The range-dependent beam pattern,  $B(r)$ , was estimated by computing the average over the azimuthal direction of the backscatter intensity in all of the images and dividing by the predicted backscatter response for a flat bed. The estimated beam pattern was then applied to the backscatter intensity predicted by the Lambertian scattering model so that the simulated images could be directly compared to the uncorrected sonar images from the instrument.

Table 2—Default Parameter Values Used in the TREX NSEA Runs

Parameter	Value
$D_{50}$	0.00023 m
S	2.65
$\theta_{cr}$	0.05
$\theta_{wo}$	0.168
$\alpha$	1.0
$\varphi$	0.4
$\sigma$	1.0 m
$\gamma$	8.0
$\xi$	1.5
$D$	$0 \text{ m}^2\text{s}^{-1}$
Ripple predictor	Nelson [12]
Grid size	$1024 \times 1024$
Patch edge length	10 m

The resulting roughness realizations and simulated sonar images for three times during the experiment are displayed in Fig. 4 alongside the observed sonar images. The predicted ripple wavelengths and directions qualitatively resemble those in the observed sonar data. The roughness realizations and simulated sonar images cannot reproduce the exact pattern of roughness observed in the sonar data. In particular, the phase of the ripples in the simulated roughness realizations will never match the phase in the observed sonar data because the roughness realization is drawn randomly from the Gaussian process implied by the modeled amplitude spectrum without taking into account any of the information provided by the sonar image. Furthermore, temporal correlations between the sonar images are not accurately reproduced by the model because the roughness realizations are independently drawn at each time step. The peak ripple wavelength time series predicted by the model is compared to the observed ripple wavelengths extracted from the imagery by the FFT in Fig. 5. Again, the model predictions agree fairly well with the data. The discrepancy between the model and the observations at the beginning of the time series exists because NSEM initializes the seafloor spectrum to be in equilibrium with the hydrodynamic forcing, even if the forcing is not strong enough to




Fig. 4—Comparison between (A) a roughness realization generated from the NSEA model output, (B) a sonar image simulated from the roughness realization, and (C) the observed sector-scanning sonar image for three different times during the TREX deployment

mobilize sediment and to generate ripples. In the case of the TREX data, long-wavelength relict ripples were present at the site at the beginning of the deployment. These ripples were formed by stronger forcing some time prior to the deployment and existed out of equilibrium with the hydrodynamics. Since NSEM does not have any information about the forcing prior to the deployment, it is unable to predict the geometry of the relict ripples, and the discrepancy persisted until the first sediment mobilization event around April 24, 2013. Once sediment was mobilized, the seafloor ripples reset, and the evolution model predicted a new ripple wavelength within a few centimeters of the observations. The end of the time series illustrates the ability of NSEM to predict the persistence of relict ripples that formed during its simulation period.




Fig. 5—Comparison between the modeled (red dots) and observed (cyan dots) peak ripple wavenumbers. The grayscale contour plotted in the background represents the predicted 1D wavenumber spectrum at extracted from the peak energy direction in the 2D spectrum. The model predicts a peak wavenumber of zero when there is a washout.

## 5.2 Statistical Inference

The parameters of the seafloor evolution model can be estimated from the TREX forcing data and the ripple wavelength data using approximate Bayesian computation. The prior distribution for the parameters was chosen to be uniform with lower and upper bounds selected by expert elicitation and shown in Table 3. The prior distribution of the power law ripple predictor, Eqs. (15) and (16), parameters was determined by fitting the power law model to the compiled ripple data from Nelson et al. [12]. One million parameter samples were taken from the prior distribution and were run through the model using the forcing data shown in Fig. 3. The peak wavelength at each time step was then computed from the wavenumber corresponding to the maximum of the amplitude spectrum. The discrepancy measure used was the mean squared deviation between the predicted and the observed ripple wavelength, and the tolerance was chosen to accept approximately 1,000 parameter samples, based on the acceptance rate of a pilot run.

Table 3—Lower and Upper Bounds of the Uniform Prior Distribution for Each of the Parameters

Parameter	Lower bound	Upper bound
$D_{50}$	0.0001 mm	0.002 mm
S	2.5	2.8
$\theta_{cr}$	0.01	0.07
$\theta_{wo}$	0.01	1.0
$\alpha$	1.0	20.0
$\varphi$	0.3	0.45
$\sigma$	1.0 m	25.0 m
$\gamma$	1.0	15.0
$\xi$	0.0	2.5
$D$	$0 \text{ m}^2\text{s}^{-1}$	$1 \times 10^{-12} \text{ m}^2\text{s}^{-1}$
$\beta_1$	-1.4	-1.0
$\beta_2$	-0.2	0.1
$\beta_3$	-0.2	0.1
$\chi_1$	1.7	2.2
$\chi_2$	-0.6	-0.3
$\chi_3$	-0.2	0.0

The posterior distributions for the median grain size and the critical shear stress are shown in Fig. 6. The peak of the posterior distribution corresponds to the most likely parameter value in the model, while the concentration of the posterior distribution relative to the prior distribution (shown with a black line in Fig. 6) indicates how sensitive the model is to that parameter. The grain size posterior peaks around a value close to the ground truth of 0.23 mm and concentrates relative to the prior. The seafloor evolution model's sensitivity to the grain size is to be expected, since the grain size relates the scales of the flow to the bed response and thus appears as a scaling parameter in many of the empirical formulae used in the seafloor evolution model, such as the bedload transport law (Eq. (13)) and the ripple predictors (Eqs. (15) and (16)). The critical shear stress posterior concentrates around a value of 0.02. While there were no independent measurements of critical shear stress at this site, the estimated value is lower than is typically used in sediment transport modeling studies ( $\sim 0.05$  for quartz sand). The critical shear stress sets the timing of sediment mobilization and the magnitude of sediment transport but does not play a role in the ripple geometry; therefore, the lower estimated critical shear stress indicates that the ripples started evolving sooner and evolved more quickly than has been historically observed. This discrepancy could indicate the importance of presently unmodeled processes in the timing of sediment mobilization; however, it is also possible that the inference algorithm uses the flexibility in the critical shear stress parameter to overcome some other discrepancy between the model and the data. More sophisticated sensitivity analyses are necessary to discriminate between these possibilities and to improve both the model and the inference procedure [61].




Fig. 6—Approximate posterior distributions for the median grain size (A) and the critical shear stress (B). The uniform prior distribution is indicated with the black line. The ground truth grain size is indicated with the red line in plot (A).

## 6. CONCLUSIONS

The Naval Seafloor Evolution Architecture is a flexible platform for forecasting seafloor roughness. NSEA is based on a simplified sediment transport model that predicts the statistical properties of the seafloor elevation under changing hydrodynamic conditions. It provides a Bayesian inference framework to combine the seafloor evolution model with hydrodynamic forcing information and observations from diverse sources. NSEA can be run in forward mode to predict seafloor roughness from wave and current forcing and to quantify the uncertainty in those predictions. Downstream applications such as acoustic or hydrodynamic models can use these roughness forecasts as boundary conditions. NSEA can also be run in an inverse mode using observations to estimate roughness or the parameters of the seafloor evolution model. The Bayesian approach unifies these two formulations into the single task of estimating the conditional probability distribution of the quantity of interest given the available information.

NSEA's modular architecture separates the specification of hydrodynamic forcing, the seafloor evolution model, and the generation of roughness spectra, seafloor elevation, and acoustic sonar response predictions. For example, the seafloor evolution model is agnostic to the source of the hydrodynamic forcing information, so the forcing can be provided from in situ observations, as illustrated in Section 5, or from hydrodynamic models. The seafloor evolution model could also be used independently to feed roughness information back into a hydrodynamic model. Finally, the acoustic observation model can be substituted to take advantage of the complementary information available from different observation modalities. An example of a simplified acoustic model for high-frequency imaging sonars is described in Section 2.4.2. However, with more sophisticated acoustic models, it is possible to constrain seafloor roughness and the parameters of the seafloor evolution model using many different kinds of acoustic observations. NSEA's modularity means that such models, developed by independent teams, can be combined with the seafloor evolution model and Bayesian inference tools with minimal effort.

### 6.1 Future Work

NSEA was originally developed to model the evolution of roughness due to wave-generated ripples on shallow, sandy seafloors, and it has been extensively tested in such environments. For NSEA to realize

the goal of producing forecasts of seafloor roughness globally requires additional research and validation. NSEA must be extended to consider the modification of roughness by processes other than surface waves. Current-generated ripples can be simulated by NSEA, but are less extensively tested, and the wave and current ripple parameterizations do not currently simulate the complex patterns that can be generated by combined wave and current flows. Larger-scale bedforms can also be created by currents. While there are few theoretical limitations in simulating large-scale roughness with NSEA's spectral approach, new empirical parameterizations need to be incorporated for NSEA to generate realistic bedforms at larger scales. In deeper water, surface waves have less influence on the bottom and processes such as geostrophic currents, internal waves and turbidity currents become more important. These can generate a large variety of geomorphological features beyond simple bedforms. Representing these kinds of forcings and their impacts on the seafloor is essential to extend the applicability of NSEA into deeper waters. Applying NSEA globally also requires mechanisms for predicting roughness and its evolution in mud, silt, and gravel sediment and in hard-bottom environments like bedrock outcrops and coral reefs.

Biological processes can also modify roughness. Organisms moving and feeding on the bottom can create and destroy roughness. In addition, the growth of sediment-associated organisms like microphytobenthos can affect the mobility of sediment and therefore the evolution of bedforms. At present, NSEA represents biological processes as a diffusion coefficient that smoothes roughness during quiescent conditions. Even this simplified model poses challenges of estimating the diffusion coefficient in different locations and times of year. To properly represent the effect of biological activity on roughness evolution likely requires coupling NSEA to ecological models.

The major challenge of all of these extensions is to ensure that the complexity required to represent these different processes is justified by increased predictive power. Each new process adds parameters that need to be calibrated, which can increase the overall uncertainty in the model's predictions even if the new processes make those predictions more accurate. An important benefit of framing NSEA as a probabilistic model is the ability to evaluate the model against observations and to make rigorous statistical comparisons between different formulations of the model. As NSEA is further developed and applied to new settings in the future, it should constantly be tested against observations to ensure that new parameterizations create better predictions of roughness evolution.

## REFERENCES

1. P.P. Mathisen and O. S. Madsen, "Waves and currents over a fixed rippled bed: 1. Bottom roughness experienced by waves in the presence and absence of currents," *Journal of Geophysical Research: Oceans* **101**(C7), 16533–16542 (1996), doi:<https://doi.org/10.1029/96JC00954>. URL <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/96JC00954>.
2. R. Styles and S.M. Glenn, "Modeling bottom roughness in the presence of wave-generated ripples," *Journal of Geophysical Research: Oceans* **107**(C8), 24–1–24–15 (2002), doi:<https://doi.org/10.1029/2001JC000864>. URL <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2001JC000864>.
3. D. Jackson and M. Richardson, *High-Frequency Seafloor Acoustics* (Springer Science & Business Media, 2007).

4. V. Marieu, P. Bonneton, D. L. Foster, and F. Arduin, "Modeling of vortex ripple morphodynamics," *Journal of Geophysical Research: Oceans* **113**(C9) (2008), doi:<https://doi.org/10.1029/2007JC004659>. URL <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2007JC004659>.
5. A. Penko, J. Calantoni, S. Rodriguez-Abudo, D. Foster, and D. Slinn, "Three-dimensional mixture simulations of flow over dynamic rippled beds," *Journal of Geophysical Research: Oceans* **118**(3), 1543–1555 (2013), doi:10.1002/jgrc.20120. URL <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1002/jgrc.20120>.
6. M. Mazzuoli, A. G. Kidanemariam, and M. Uhlmann, "Direct numerical simulations of ripples in an oscillatory flow," *Journal of Fluid Mechanics* **863**, 572–600 (2019), doi:10.1017/jfm.2018.1005.
7. A. Salimi-Tarazouj, T. J. Hsu, P. Traykovski, Z. Cheng, and J. Chauchat, "A Numerical Study of Onshore Ripple Migration Using a Eulerian Two-phase Model," *Journal of Geophysical Research: Oceans* **126**(2), e2020JC016773 (2021), doi:<https://doi.org/10.1029/2020JC016773>. URL <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2020JC016773>, e2020JC016773 2020JC016773.
8. A. Salimi-Tarazouj, T. J. Hsu, P. Traykovski, and J. Chauchat, "Eulerian Two-Phase Model Reveals the Importance of Wave Period in Ripple Evolution and Equilibrium Geometry," *Journal of Geophysical Research: Earth Surface* **126**(7), e2021JF006132 (2021), doi:<https://doi.org/10.1029/2021JF006132>. URL <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2021JF006132>, e2021JF006132 2021JF006132.
9. P. Nielsen, "Dynamics and geometry of wave-generated ripples," *Journal of Geophysical Research: Oceans* **86**(C7), 6467–6472 (1981), doi:10.1029/JC086iC07p06467. URL <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/JC086iC07p06467>.
10. L. C. van Rijn, "Sediment Transport, Part III: Bed forms and Alluvial Roughness," *Journal of Hydraulic Engineering* **110**(12), 1733–1754 (1984), doi:10.1061/(ASCE)0733-9429(1984)110:12(1733).
11. P. L. Wiberg and C. K. Harris, "Ripple geometry in wave-dominated environments," *Journal of Geophysical Research: Oceans* **99**(C1), 775–789 (1994), doi:10.1029/93JC02726. URL <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/93JC02726>.
12. T. R. Nelson, G. Voulgaris, and P. Traykovski, "Predicting wave-induced ripple equilibrium geometry," *Journal of Geophysical Research: Oceans* **118**(6), 3202–3220 (2013), doi:10.1002/jgrc.20241. URL <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1002/jgrc.20241>.
13. E. B. Goldstein, G. Coco, and A. B. Murray, "Prediction of wave ripple characteristics using genetic programming," *Continental Shelf Research* **71**, 1–15 (2013).
14. P. Traykovski, "Observations of wave orbital scale ripples and a nonequilibrium time-dependent model," *Journal of Geophysical Research* **112** (2007), doi:10.1029/2006JC003811.
15. A. E. Hay, "Near-bed turbulence and relict waveformed sand ripples: Observations from the inner shelf," *Journal of Geophysical Research: Oceans* **113**(C4) (2008), doi:<https://doi.org/10.1029/2006JC004013>. URL <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2006JC004013>.
16. C. B. DuVal, A. C. Trembanis, and D. C. Miller, "A Regime-State Framework for Morphodynamic Modeling of Seabed Roughness," *Journal of Geophysical Research: Oceans* **126**(5), e2020JC016769 (2021), doi:<https://doi.org/10.1029/2020JC016769>. URL <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2020JC016769>, e2020JC016769 2020JC016769.



17. R. Soulsby, R. Whitehouse, and K. Marten, “Prediction of time-evolving sand ripples in shelf seas,” *Continental Shelf Research* **38**, 47–62 (2012), doi:10.1016/j.csr.2012.02.016.
18. T. R. Nelson and G. Voulgaris, “A spectral model for estimating temporal and spatial evolution of rippled seabeds,” *Ocean Dynamics* **65**, 155–171 (2015), doi:10.1007/s10236-014-0801-y.
19. A. Penko, J. Calantoni, and B. T. Hefner, “Modeling and observations of sand ripple formation and evolution during TREX13,” *IEEE Journal of Oceanic Engineering* **42**(2), 260–267 (2017).
20. E. Meyer-Peter and R. Müller, “Formulas for bed-load transport,” Proceedings of the IAHSR 2nd meeting, Stockholm, appendix 2 (IAHR), 1948.
21. J. Bezanson, A. Edelman, S. Karpinski, and V. B. Shah, “Julia: A Fresh Approach to Numerical Computing,” *SIAM Review* **59**(1), 65–98 (Jan 2017), ISSN 1095-7200, doi:10.1137/141000671. URL <http://dx.doi.org/10.1137/141000671>.
22. P. Guttorp, *Stochastic Modeling of Scientific Data* (Chapman and Hall/CRC, 1995).
23. K. Holland and P. Elmore, “A review of heterogeneous sediments in coastal environments,” *Earth-Science Reviews* **89**(3), 116–134 (2008), ISSN 0012-8252, doi:<https://doi.org/10.1016/j.earscirev.2008.03.003>. URL <https://www.sciencedirect.com/science/article/pii/S0012825208000329>.
24. D. Tang, F. S. Henyey, B. T. Hefner, and P. A. Traykovski, “Simulating realistic-looking sediment ripple fields,” *IEEE Journal of Oceanic Engineering* **34**(4), 444–450 (2009).
25. C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning* (The MIT Press, 11 2005).
26. F. Lindgren, H. Rue, and J. Lindström, “An explicit link between Gaussian fields and Gaussian Markov random fields: the stochastic partial differential equation approach,” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **73**(4), 423–498 (2011), doi:<https://doi.org/10.1111/j.1467-9868.2011.00777.x>. URL <https://rss.onlinelibrary.wiley.com/doi/abs/10.1111/j.1467-9868.2011.00777.x>.
27. T. Gneiting, M. G. Genton, and P. Guttorp, “Geostatistical space-time models, stationarity, separability and full symmetry,” in B. Finkenstadt, L. Held, and V. Isham, eds., *Statistical methods for spatio-temporal systems* (Taylor and Francis, 2007).
28. T. S. team, “SWAN Scientific and Technical Documentation,” Delft University of Technology, 2021.
29. C. Leith, “Stochastic models of chaotic systems,” *Physica D: Nonlinear Phenomena* **98**(2), 481–491 (1996), ISSN 0167-2789, doi:[https://doi.org/10.1016/0167-2789\(96\)00107-8](https://doi.org/10.1016/0167-2789(96)00107-8). URL <https://www.sciencedirect.com/science/article/pii/0167278996001078>, Nonlinear Phenomena in Ocean Dynamics.
30. P. L. Wiberg and C. R. Sherwood, “Calculating wave-generated bottom orbital velocities from surface-wave parameters,” *Computers & Geosciences* **34**(10), 1243–1262 (2008), ISSN 0098-3004, doi:<https://doi.org/10.1016/j.cageo.2008.02.010>. URL <https://www.sciencedirect.com/science/article/pii/S009830040800054X>, Predictive Modeling in Sediment Transport and Stratigraphy.

31. J. P. Davis, D. J. Walker, M. Townsend, and I. R. Young, “Wave-formed sediment ripples: Transient analysis of ripple spectral development,” *Journal of Geophysical Research: Oceans* **109**(C7) (2004), doi:10.1029/2004JC002307. URL <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2004JC002307>.
32. M. Wong and G. Parker, “Reanalysis and Correction of Bed-Load Relation of Meyer-Peter and Müller Using Their Own Database,” *Journal of Hydraulic Engineering* **132**(11), 1159–1168 (2006), doi:10.1061/(ASCE)0733-9429(2006)132:11(1159). URL <https://ascelibrary.org/doi/abs/10.1061/%28ASCE%290733-9429%282006%29132%3A11%281159%29>.
33. D. R. Jackson, M. D. Richardson, K. L. Williams, A. P. Lyons, C. D. Jones, K. B. Briggs, and D. Tang, “Acoustic observation of the time dependence of the roughness of sandy seafloors,” *IEEE Journal of Oceanic Engineering* **34**(4), 407–422 (2009).
34. A. M. Penko, S. Johnson, and J. Calantoni, “Simulation of measured seafloor roughness spectrum time series using a coupled ripple-bioturbation model,” *Proceedings of the Seabed Sediment Acoustics Measurement Modeling* **37**, 365 (2015).
35. B. Øksendal, *Stochastic Differential Equations* (Springer, 2003).
36. J. C. Cox, J. E. Ingersoll, and S. A. Ross, “A Theory of the Term Structure of Interest Rates,” *Econometrica* **53**(2), 385–407 (1985), ISSN 00129682, 14680262. URL <http://www.jstor.org/stable/1911242>.
37. E. I. Thorsos, “The validity of the Kirchhoff approximation for rough surface scattering using a Gaussian roughness spectrum,” *The Journal of the Acoustical Society of America* **83**(1), 78–92 (1988), doi:10.1121/1.396188. URL <https://doi.org/10.1121/1.396188>.
38. D. R. Olson and D. Jackson, “Scattering from layered seafloors: Comparisons between theory and integral equations,” *The Journal of the Acoustical Society of America* **148**(4), 2086–2095 (2020), doi:10.1121/10.0002164. URL <https://doi.org/10.1121/10.0002164>.
39. C. J. Paciorek, “Bayesian smoothing with Gaussian processes using Fourier basis functions in the spectralGP package,” *Journal of Statistical Software* **19**(2), nihpa22751 (2007).
40. M. Frigo and S. Johnson, “The Design and Implementation of FFTW3,” *Proceedings of the IEEE* **93**(2), 216–231 (2005), doi:10.1109/JPROC.2004.840301.
41. E. Wong and B. Hajek, *Stochastic Processes in Engineering Systems* (Springer, 1985).
42. G. Chan and A. T. Wood, “Simulation of stationary Gaussian vector fields,” *Statistics and Computing* **9**(4), 265–268 (1999).
43. C. R. Dietrich and G. N. Newsam, “Fast and Exact Simulation of Stationary Gaussian Processes through Circulant Embedding of the Covariance Matrix,” *SIAM Journal on Scientific Computing* **18**(4), 1088–1107 (1997), doi:10.1137/S1064827592240555. URL <https://doi.org/10.1137/S1064827592240555>.
44. M. L. Stein, “Fast and Exact Simulation of Fractional Brownian Surfaces,” *Journal of Computational and Graphical Statistics* **11**(3), 587–599 (2002), doi:10.1198/106186002466. URL <https://doi.org/10.1198/106186002466>.

45. T. Gneiting, H. Ševčíková, D. B. Percival, M. Schlather, and Y. Jiang, “Fast and Exact Simulation of Large Gaussian Lattice Systems in  $\mathbb{R}^2$ : Exploring the Limits,” *Journal of Computational and Graphical Statistics* **15**(3), 483–501 (2006), doi:10.1198/106186006X128551. URL <https://doi.org/10.1198/106186006X128551>.
46. C. Sigg and M. Hadwiger, “Fast third-order texture filtering,” in M. Pharr, ed., *GPU Gems 2*, pp. 313–329 (Addison-Wesley, 2006).
47. J. P. Boyd, *Chebyshev and Fourier Spectral Methods* (Courier Corporation, 2001).
48. K. Cranmer, J. Brehmer, and G. Louppe, “The frontier of simulation-based inference,” *Proceedings of the National Academy of Sciences* **117**(48), 30055–30062 (2020), doi:10.1073/pnas.1912789117. URL <https://www.pnas.org/doi/abs/10.1073/pnas.1912789117>.
49. L. Brakenhoff, M. Kleinhan, G. Ruessink, and M. van der Vegt, “Spatio-temporal characteristics of small-scale wave–current ripples on the Ameland ebb-tidal delta,” *Earth Surface Processes and Landforms* **45**(5), 1248–1261 (2020), doi:<https://doi.org/10.1002/esp.4802>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/esp.4802>.
50. N. Wiener, *Extrapolation, Interpolation, and Smoothing of Stationary Time Series: With Engineering Applications* (The MIT Press, 08 1949), ISBN 9780262257190, doi:10.7551/mitpress/2946.001.0001. URL <https://doi.org/10.7551/mitpress/2946.001.0001>.
51. J. Bell, M. Chantler, and T. Wittig, “Sidescan sonar: a directional filter of seabed texture?,” *IEE Proceedings - Radar, Sonar and Navigation* **146**, 65–72(7) (February 1999). URL [https://digital-library.theiet.org/content/journals/10.1049/ip-rsn\\_19990266](https://digital-library.theiet.org/content/journals/10.1049/ip-rsn_19990266).
52. P. Cervenka and C. de Moustier, “Sidescan sonar image processing techniques,” *IEEE Journal of Oceanic Engineering* **18**(2), 108–122 (1993), doi:10.1109/48.219531.
53. J. M. Bell, “Application of optical ray tracing techniques to the simulation of sonar images,” *Optical Engineering* **36**(6), 1806–1813 (1997), doi:10.1117/1.601325. URL <https://doi.org/10.1117/1.601325>.
54. A. P. Lyons and D. A. Abraham, “Statistical characterization of high-frequency shallow-water seafloor backscatter,” *The Journal of the Acoustical Society of America* **106**(3), 1307–1315 (1999), doi:10.1121/1.428034. URL <https://doi.org/10.1121/1.428034>.
55. D. Abraham and A. Lyons, “Novel physical interpretations of K-distributed reverberation,” *IEEE Journal of Oceanic Engineering* **27**(4), 800–813 (Oct 2002), ISSN 0364-9059, doi:10.1109/joe.2002.804324. URL <http://dx.doi.org/10.1109/joe.2002.804324>.
56. A. Lyons, S. Johnson, D. Abraham, and E. Pouliquen, “High-Frequency Scattered Envelope Statistics of Patchy Seafloors,” *IEEE Journal of Oceanic Engineering* **34**(4), 451–458 (Oct 2009), ISSN 1558-1691, doi:10.1109/joe.2008.2008410. URL <http://dx.doi.org/10.1109/joe.2008.2008410>.
57. J. Folkesson, H. Chang, and N. Bore, “Lambert’s Cosine Law and Sidescan Sonar Modeling,” *Proceedings of the 2020 IEEE/OES Autonomous Underwater Vehicles Symposium (AUV)*, 2020, pp. 1–6. doi:10.1109/AUV50043.2020.9267946.
58. M. Sunnåker, A. G. Busetto, E. Numminen, J. Corander, M. Foll, and C. Dessimoz, “Approximate Bayesian Computation,” *PLOS Computational Biology* **9**(1), 1–10 (01 2013), doi:10.1371/journal.pcbi.1002803. URL <https://doi.org/10.1371/journal.pcbi.1002803>.

59. A. Skarke and A. C. Trembanis, “Parameterization of bedform morphology and defect density with fingerprint analysis techniques,” *Continental Shelf Research* **31**(16), 1688–1700 (2011), ISSN 0278-4343, doi:<https://doi.org/10.1016/j.csr.2011.07.009>. URL <https://www.sciencedirect.com/science/article/pii/S0278434311002603>.
60. T. Besard, C. Foket, and B. De Sutter, “Effective Extensible Programming: Unleashing Julia on GPUs,” *IEEE Transactions on Parallel and Distributed Systems* (2018), ISSN 1045-9219, doi:10.1109/TPDS.2018.2872064.
61. A. Gelman, J. B. Carlin, H. S. Stern, D. Dunson, A. Vehtari, and D. B. Rubin, *Bayesian Data Analysis* (CRC Press, 2014).