Bayesian Inference of Drag Parameters Using AXBT Data from Typhoon Fanapi

IHAB SRAJ,* MOHAMED ISKANDARANI,+ ASHWANTH SRINIVASAN,+ W. CARLISLE THACKER,#
JUSTIN WINOKUR,* ALEN ALEXANDERIAN,@ CHIA-YING LEE,+
SHUYI S. CHEN,+ AND OMAR M. KNIO*

* *Duke University, Durham, North Carolina*
+ *University of Miami, Miami, Florida*
# *University of Miami, and Atlantic Oceanographic and Meteorological Laboratory, Miami, Florida*
@ *University of Texas at Austin, Austin, Texas*

(Manuscript received 8 August 2012, in final form 9 January 2013)

ABSTRACT

The authors introduce a three-parameter characterization of the wind speed dependence of the drag coefficient and apply a Bayesian formalism to infer values for these parameters from airborne expendable bathythermograph (AXBT) temperature data obtained during Typhoon Fanapi. One parameter is a multiplicative factor that amplifies or attenuates the drag coefficient for all wind speeds, the second is the maximum wind speed at which drag coefficient saturation occurs, and the third is the drag coefficient's rate of change with increasing wind speed after saturation. Bayesian inference provides optimal estimates of the parameters as well as a non-Gaussian probability distribution characterizing the uncertainty of these estimates. The efficiency of this approach stems from the use of adaptive polynomial expansions to build an inexpensive surrogate for the high-resolution numerical model that couples simulated winds to the oceanic temperature data, dramatically reducing the computational burden of the Markov chain Monte Carlo sampling. These results indicate that the most likely values for the drag coefficient saturation and the corresponding wind speed are about 2.3×10^{-3} and $34\,\mathrm{m\,s^{-1}}$, respectively; the data were not informative regarding the drag coefficient behavior at higher wind speeds.

1. Introduction

The present work focuses on investigating the behavior of the drag coefficient at high wind speeds—a regime that is particularly vexing to observe, and that is critical to understanding the interplay between tropical cyclones and oceans. It has been known for some time that the drag coefficient saturates at high wind speeds (rather than continuing to increase), and this saturation affects the rate at which a tropical cyclone loses momentum to the ocean (Powell et al. 2003; Donelan et al. 2004; French et al. 2007; Shay and Jacob 2006; Sanford et al. 2007; Jarosz et al. 2007). The particular wind speed and drag values at which this transition occurs are not well known because of scatter in the observational data.

The extreme environmental conditions during tropical cyclones make answering these questions with direct observations very difficult. This is compounded by the fact that observations of the upper-ocean temperature field under high-wind conditions are rare. Fortunately, however, the intense observational program, Impact of Typhoons on the Ocean in the Pacific (ITOP), dropped a large number of airborne expendable bathythermographs (AXBTs) in the anticipated track of Typhoon Fanapi 2010. The central objective of the present paper is to investigate whether the temperature dataset resulting from ITOP's Fanapi observations can inform us on the behavior of the drag coefficient. In particular, we attempt to answer the following questions: What is the drag coefficient saturation value $C_D^{\max}$ and at what wind speed $V_{\max}$ does it occur? Does the drag coefficient decrease for wind speeds greater than $V_{\max}$ or does it stay constant?

As further described below, we will adopt an inverse modeling approach toward the objective above. To this end, we take advantage of recent developments in uncertainty quantification (UQ) methods, which make it feasible to probe the sensitivity of complex and realistic ocean models to uncertain parameters,

Corresponding author address: Omar Knio, Department of Mechanical Engineering and Materials Science, Duke University, Durham, NC 27708.
E-mail: omar.knio@duke.edu

DOI: 10.1175/MWR-D-12-00228.1

and to assess the amount of information that can be gained from specific measurements and observations. Specifically, we will rely on polynomial chaos (PC) expansions to construct a faithful surrogate of the response of the large and complex geophysical model simulations. The availability of the surrogate enables us to efficiently implement a Bayesian inference formalism to the inverse problem, namely, because it dramatically reduces the cost of sampling the posterior distribution.

PC expansions have been developed and applied in the engineering community to quantify uncertainties in numerical simulations; their principal use is in propagating input uncertainties through large, complex, and nonlinear models to compute the ensuing output uncertainties. The present authors have been engaged in investigating the applicability of PC expansions for UQ in oceanic simulations. Thacker et al. (2012) explored the impact of uncertainty in southern boundary conditions on the Gulf of Mexico circulation as simulated by the Hybrid Coordinate Ocean Model (HYCOM). Srinivasan et al. (2010) applied PC expansions to quantify the uncertainties in oil droplet size distribution and oil composition on the oil fate discharged from the Deepwater Horizon blow out. Alexanderian et al. (2012) investigated the impact of parametric uncertainty in the ocean mixed layer parameterization and drag coefficient on the sea surface temperature (SST) cooling in the wake of Hurricane Ivan (2004).

Implicit in these aforementioned applications is the assumption that the input uncertainties and their statistical distributions are well known when in fact very little is known about them, and one is often forced to make practical choices based on the best available information. Most often, one resorts to the common choices of Gaussian distributions of the uncertain parameters, or uniform distributions (commonly referred to as noninformative priors in the Bayesian literature). In the present paper we use PC expansions and Bayesian inference to correct the probability distribution of *input* uncertainties, and in the process illustrate how observational data can be effectively incorporated in the UQ framework.

The inverse problem addressed here relies on indirect observational data (temperature) to infer quantities of interest (drag coefficient parameters), and on a forward model (HYCOM) that can predict the observations given a set of drag parameters. Numerous methods have been developed for such inverse modeling (Vogel 2002; Sivia 2006; Tarantola 2005). Variational techniques akin to maximizing likelihood have been used to estimate parameters for numerical models of the ocean and atmosphere (e.g., Smedstad and O'Brien 1991; Zedler et al. 2009, 2012), as have closely related methods such as the ensemble Kalman filter (e.g., Annan et al. 2005; Kondrashov et al. 2008; Aksoy et al. 2006a,b; Hu et al. 2010). As the parameters for biological systems are numerous, model specific, and poorly known, their estimation from observations receives a great deal of attention (e.g., Fasham and Evans 1995; Fennel et al. 2001; Losa et al. 2004; Jones et al. 2010; Dowd 2011). Here we employ the same Bayesian framework for parameter estimation that these approaches use while reducing the computational burden through the use of PC expansions. Mattern et al. (2012) have similarly exploited the virtues of such polynomial expansions for examining the response of ecosystem models to finite perturbations of their uncertain parameters, but without Bayesian inference. Similarly seeking computational efficiency, Olson et al. (2012) used a Gaussian emulator rather than polynomial expansions, combining a Bayesian framework with Markov chain Monte Carlo (MCMC) sampling to estimate the probability densities that describe the uncertainties in the parameters of a climate model.

The major hurdle in applying Bayesian inference is the computational cost associated with sampling the uncertainty space using a large, complex, and CPU-intensive model like HYCOM. Here we resort to PC expansions to build a faithful surrogate (also known as an emulator) that can be efficiently sampled and used in lieu of the model during the MCMC phase (Marzouk et al. 2007; Marzouk and Najm 2009). An ensemble-like calculation is required to construct the surrogate and this represents the most CPU-intensive part of the calculations. In the present work, we rely on an adaptive approach to construct the surrogate in order to mitigate the computational cost while minimizing the surrogate's approximation errors. Several error metrics will be presented to ascertain the surrogate's validity.

The behavior of the drag coefficient at high wind speeds over the ocean has been the subject of considerable research. Shay and Jacob (2006) used airborne expendable current profiler (AXCP) data deployed during Hurricane Gilbert (1988) to calculate the downward energy flux from the mixed layer into the thermocline and equated it to the applied wind stress; their calculations show the drag coefficient leveling off at a value of about 3.2×10^{-3} for wind speeds upward of $30\,\mathrm{m\,s^{-1}}$. Sanford et al. (2007) compared SST and transport measured by profiling floats under Hurricane Frances (2004) to their model counterparts and concluded that drag parameterizations consistent with drag saturation and decrease produce a better fit to observations. Chen et al. (2007, 2013) used a fully

coupled atmosphere–wave–ocean model to explicitly compute the surface wave-induced stress (form drag), the most dominant contributor to wind stress in high wind conditions; it was found that the surface stress was highly asymmetric around the storm center in Hurricane Frances (Chen et al. 2013). Jarosz et al. (2007) estimated the drag coefficient by combining barotropic linear dynamics over the continental shelf in the northeastern Gulf of Mexico with in situ current measurements under Hurricane Ivan (2004); their analysis produced a quadratic fit to the drag-law dependency that peaks around 2.2×10^{-3} at a wind speed of $32\,\mathrm{m\,s^{-1}}$, and that suggests a decrease in the drag for higher wind speed. Zedler et al. (2009) performed a model sensitivity analysis on the ocean response to Hurricane Frances and found that a constant drag coefficient of 1.2×10^{-3} reduced misfits between model results and drifter and buoy measurements. More recently, Zedler et al. (2012) conducted a feasibility study of inferring drag parameters from temperature data in an idealized process-oriented setting, with the aims of optimizing the observations' location and investigating the impact of measurement errors on the inferred parameters.[1] Walsh et al. (2010) simulated the cooling under several tropical cyclones with varying drag and enthalpy coefficients, and mixed layer parameterization. They concluded that drag formulations with saturation explain the satellite-observed SST better than those without saturation.

The layout of the manuscript is as follows. First, we motivate the UQ problem by discussing the uncertainties of the drag parameters in section 2. Section 3 summarizes the ITOP observational data and their model counterparts. Section 4 gives a brief overview of the Bayesian formalism that brings together observations and model results. Section 5 gives a brief overview of the PC expansions that were used to build the surrogate model. Section 6 describes the main results in three subsections: first, we present evidence that the surrogate is a faithful representation of the HYCOM model; second, we use the surrogate to explore the response surfaces of simulated model temperatures, and to compute their sensitivity to the control parameters; and third, we present the posterior distribution obtained via MCMC sampling and contrast it to the prior distributions. A discussion of our findings and methodology is presented in section 7.

[1] Our goals are similar to those of Zedler et al. (2012), but we have relied on existing observations and on data-assimilated ocean and atmospheric general circulation models.

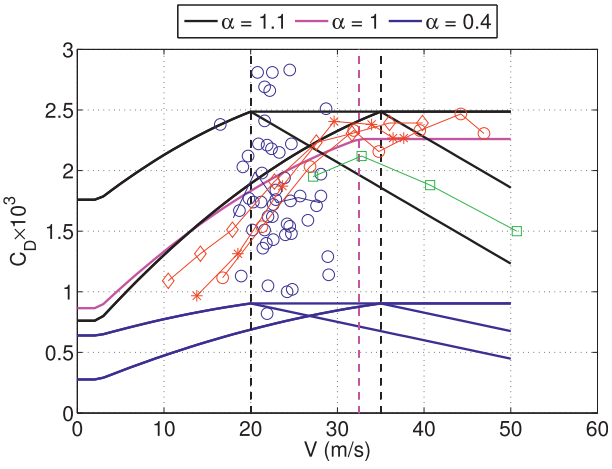

FIG. 1. Variations of drag coefficient C_D with wind speed. The thin blue circles refer to aircraft observations (French et al. 2007), the red lines with symbols refer to laboratory experiments (Donelan et al. 2004), and the green line refers to dropsonde measurements (Powell et al. 2003). The blue and black curves illustrate the eight cases determined by setting the control parameters to the extremes of their allowed ranges, and which are shown in blue ($\alpha = 0.4$) and black ($\alpha = 1.1$) with $V_{\max}$ set to either 20 or $35\,\mathrm{m\,s^{-1}}$ and m to either -3.8×10^{-5} or 0. The unperturbed HYCOM parameterization of C_D (Kara et al. 2002) is shown in magenta ($\alpha = 1$, $V_{\max} = 32.5\,\mathrm{m\,s^{-1}}$, and $m = 0$).

2. Drag coefficient

The ability to forecast tropical cyclone intensity changes depends crucially on understanding, and on modeling realistically, air–sea interactions at high wind speeds. A key advance in the recent past has been the recognition that, under the strong winds found in tropical cyclones, the drag coefficient C_D ceases to grow monotonically with increased wind speed; instead, it reaches a maximum and remains more or less constant beyond that. This can be qualitatively explained by the very high wind shearing off the ocean surface waves and thus limiting the atmosphere's influence on (and momentum transfer to) the ocean. The drag saturation is supported by numerous experimental and observational data as shown in Fig. 1. However, the different experiments vary in their estimates of the saturation value $C_D^{\max}$ and the corresponding wind speed $V_{\max}$ at which saturation occurs. Furthermore, the behavior of the drag coefficient beyond this maximum is not well known, and only dropsondes (Powell et al. 2003) in Fig. 1 suggest a decrease in C_D beyond $V_{\max}$. The aim of the present paper is to contribute to the discussion of the probable values of these parameters via an inverse modeling approach.

The impetus for the present study was a pilot application of PC expansions to study parametric uncertainty in ocean models (Alexanderian et al. 2012). Four

parameters were perturbed in Alexanderian et al. (2012): 3 K-profile parameterization (KPP) subgrid mixing parameters and the drag coefficient, and the impact of these perturbations on SST during the passage of Hurricane Ivan through the Gulf of Mexico in 2004 was investigated. The analysis showed that the uncertainty in SST is dominated by the uncertainty in the drag coefficient (particularly near the hurricane), and that the KPP subgrid mixing parameters contributed little despite the unrealistically large ranges explored. Hence, the present study focuses solely on drag coefficient parameters, and fixes the mixed layer parameters to their reference values.

Wind stresses are typically calculated using the drag law: $\boldsymbol{\tau} = \rho_a C_D \mathbf{V} V$ where ρ_a is the air density, C_D is the drag coefficient, $\mathbf{V}$ is the wind velocity vector at 10 m above sea level, and $V = \|\mathbf{V}\|$ is the wind speed. The drag coefficient in turn depends on the turbulent exchanges between the ocean and atmosphere and these effects are parameterized. In HYCOM this parameterization takes the following form:

$$C_D = C_{\mathrm{D0}} + C_{\mathrm{D1}}(T_s - T_a), \tag{1}$$

$$C_{\mathrm{D0}} = 10^{-3}[0.692 + 0.071\tilde{V} - 0.0007\tilde{V}^2], \tag{2}$$

$$C_{\mathrm{D1}} = 10^{-3}[0.083 - 0.0054\tilde{V} + 0.000\,093\tilde{V}^2], \tag{3}$$

$$\tilde{V} = \max[V_{\min}, \min(V_{\max}, V)], \tag{4}$$

where T_s and T_a are the sea surface and air temperatures, respectively. The wind speed is limited prior to applying it to compute the drag coefficient and the result is $\tilde{V}$; the HYCOM upper and lower bounds are $V_{\min} = 2.5\,\mathrm{m\,s^{-1}}$ and $V_{\max} = 32.5\,\mathrm{m\,s^{-1}}$. The upper limit reproduces the saturation behavior alluded to earlier. Equations (1)–(4) are from Kara et al. (2002) and are a computationally efficient quadratic fit to the Coupled Ocean–Atmosphere Response Experiment (COARE), version 2.5, parameterization. The dominant term in Eq. (1) is C_{D0} whereas C_{D1} represents the effect of the boundary layer stability on C_D.

To investigate the behavior of the drag coefficient at high wind speeds we introduce three control variables (Zedler et al. 2012) that will allow us to perturb the drag coefficient according to the following form:

$$C_D' = \alpha C_D \quad \text{for} \quad V < V_{\max}, \tag{5}$$

$$C_D' = \alpha[C_D + m(V - V_{\max})] \quad \text{for} \quad V > V_{\max}. \tag{6}$$

Equations (5)–(6) are simple modifications of Kara et al. (2002)'s parameterization: the parameter $V_{\max}$ is used to adjust the wind speed at which saturation occurs; the parameter α is used to adjust the size of the drag coefficient, while preserving the shape of the wind speed dependence; and for speeds greater than $V_{\max}$, the parameter m (slope of drag coefficient after saturation) allows for the possibility of decreasing drag with increasing wind.[2] Table 1 lists the parameter ranges explored in the present work. The formulation implies a saturation drag range between 1×10^{-3} and 2.5×10^{-3}. This range covers the low drag value reported in Zedler et al. (2009) and its upper limit is slightly higher than the drag value reported in Jarosz et al. (2007).[3] The $V_{\max}$ range is well within that suggested in Fig. 1 by the experimental and observational data.[4] The linear slope is varied between a minimum of $m^{(1)} = -3.8 \times 10^{-5}$, as suggested by the dropsonde data of Powell et al. (2003), and $m^{(2)} = 0$. The perturbed drag coefficient C_D' falls within the envelope of the experimental data as shown in Fig. 1.

TABLE 1. Table of parameter ranges for Eqs. (5) and (6).

Parameter	Min	Max
α	0.4	1.1
$V_{\max}$ ($\mathrm{m\,s^{-1}}$)	20	35
m	-3.8×10^{-5}	0.0

Finally, in anticipation of the PC expansions, we introduce the three independent, canonical stochastic variables ($|\xi_i| \le 1, i = \ldots, 3$) that will be used to characterize the three uncertain parameters; they are related linearly to the drag parameters according to

$$V_{\max} = \frac{V_{\max}^{(2)} - V_{\max}^{(1)}}{2}\xi_1 + \frac{V_{\max}^{(2)} + V_{\max}^{(1)}}{2}, \tag{7}$$

$$m = \frac{m^{(2)} - m^{(1)}}{2}\xi_2 + \frac{m^{(2)} + m^{(1)}}{2}, \tag{8}$$

$$\alpha = \frac{\alpha^{(2)} - \alpha^{(1)}}{2}\xi_3 + \frac{\alpha^{(2)} + \alpha^{(1)}}{2}. \tag{9}$$

The stochastic variables ($\xi_i, i = \ldots, 3$) will be assumed to be uniformly distributed over $[-1, 1]$, and we will use Bayesian inference to sharpen their distributions.

[2] The linear dependence on V is the simplest behavior one can assume to explore whether C_D decreases beyond saturation.

[3] It was not possible to extend the C_D range to include the high value reported in Shay and Jacob (2006) because our HYCOM simulations became unstable in this limit.

[4] The aircraft measurements suggest a $V_{\max}$ around $20\,\mathrm{m\,s^{-1}}$ whereas the dropsondes and laboratory experiments suggest a $V_{\max}$ of about $33\,\mathrm{m\,s^{-1}}$.

3. Fanapi observations and simulations

Before delving into the details of the surrogate construction and the inverse modeling methodologies, we visit some of the essential ingredients necessary to undertake these tasks, namely, the observational data and the simulated atmospheric and oceanic fields. We give a brief overview of the observational dataset available, follow it with a brief description of the atmospheric simulations necessary for the high-resolution atmospheric forcing fields, and finally present a first comparison of a reference HYCOM simulation with observational data.

a. Fanapi observations

Typhoon Fanapi formed as a tropical storm on 14 September 2010, southeast of Taiwan. It intensified into a typhoon on 16 September and started moving northwest. By 18 September, Fanapi strengthened into a category 3 typhoon heading toward Taiwan with maximum sustained wind speeds of 54 m s^{-1}. Fanapi made landfall early on 19 September at which time it changed direction and started moving southwest. Later that day, it weakened to a tropical storm that lasted 9 h before becoming a tropical depression on 20 September and totally dissipating on 21 September. Fanapi's track—as reported by the Joint Typhoon Warning Center (JTWC)—is shown in Fig. 2.

Fanapi developed and intensified within the observational array of a major field campaign ITOP. It is the best-observed typhoon over the northwest Pacific with unprecedented airborne observations of the atmosphere and the ocean from the U.S. Air Force C-130 aircraft. Our interest here is in the AXBT data that were collected along Fanapi's predicted track from 12 September to 20 September during different C-130 flights. The locations of AXBTs during selected flights are shown in Fig. 2. A total of 119 AXBTs were processed and used in the present work (unfortunately, data collected on the last day of operations, 19–20 September, were not available). Each AXBT measured the vertical temperature profile at the drop location; profiles on selected dates are shown in Fig. 3. The temperature profiles collected before Fanapi's intensification on 17 September reveal an SST of 29°C, and a mixed layer depth of about 50 m. Data collected on 17–18 September, after the typhoon intensified, show a cooler SST ranging between 24° and 28°C, and a mixed layer depth that varies between 50 and 100 m; it also shows an increase in the spread of temperatures between the different observation points.

b. Surface atmospheric forcing

A major hurdle in simulating the oceanic response to tropical cyclones is the availability of accurate and realistic atmospheric forcing datasets. In situ observations

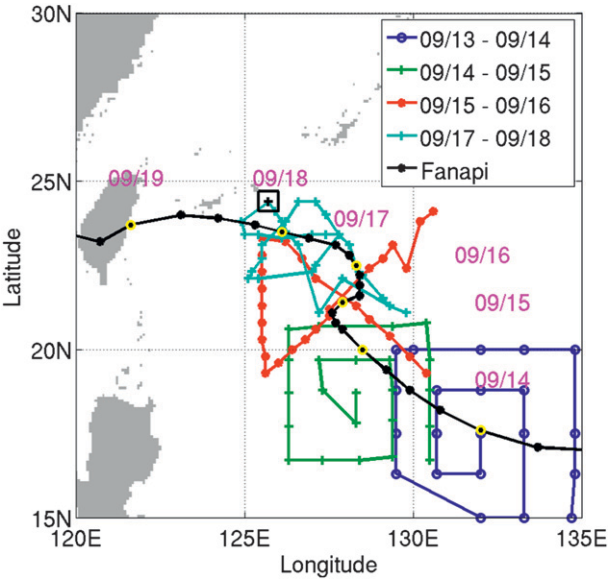

FIG. 2. Fanapi's track in 6-h intervals as reported by JTWC (black curve and circles) and the paths of selected C-130 flights during the period 13–18 Sep. The yellow circles on the track represent the typhoon center at 0000 UTC of the day indicated in the labels. The markers on the flight paths are the locations of the AXBT drops that totaled 119. The small box is 42 km × 42 km in area and is used for the analysis in the results section.

are limited and sparse in time and space, especially over the ocean. Global model analysis fields based on both in situ and satellite observations have relatively coarse spatial resolution (0.5°–1.0°), and cannot resolve the inner-core structure of tropical cyclones, where the strongest gradients and highest wind speeds are found. A useful alternative is to run a very high-resolution atmospheric model in hindcast mode and attempt to recreate, as realistically as possible, atmospheric conditions during Typhoon Fanapi. Here we resort to a triply nested (12, 4, and 1.3 km) high-resolution Weather Research and Forecasting Model (WRF) to simulate the atmospheric conditions during Fanapi's transit. Details of the model are presented in the appendix. Storm track and intensity (as measured by maximum 10-m wind speed and minimum sea level pressure) are useful and commonly used metrics to evaluate model fidelity against JTWC's best-track data.[5] These metrics are compared in Fig. 4, and confirm that the triply nested WRF has reproduced successfully Fanapi's intensity and track, except for an overintensification on 18–19 September.

[5] Note that the best-track data are also subject to measurement errors, and are best estimates based on available in situ and satellite observations. For example, the surface wind speed observations over the ocean are mostly from indirect satellite retrievals (Velden et al. 2006).

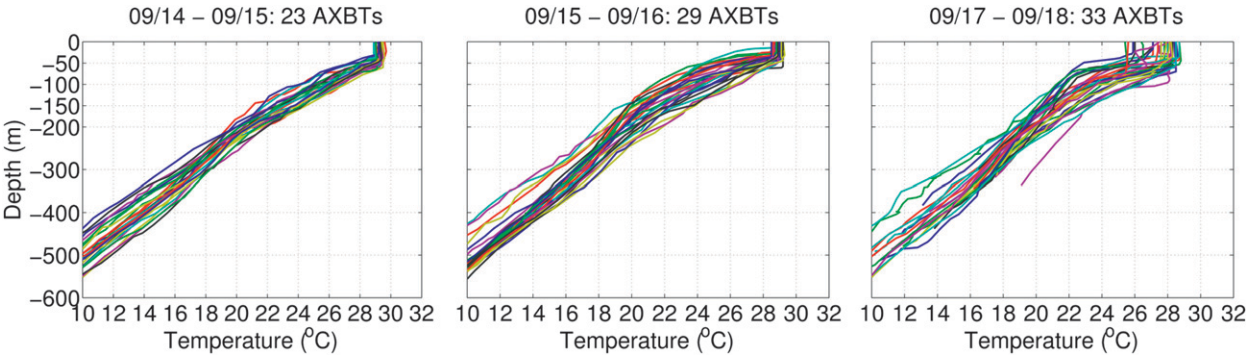

FIG. 3. (from left to right) Vertical temperature profiles from AXBTs for selected C-130 flights during Typhoon Fanapi as it intensifies during 14–18 Sep. The number of data points collected during each C-130 flight is indicated in the panel heading.

The high-resolution WRF simulations are adopted to compute HYCOM's atmospheric forcing. Most importantly, the WRF simulations provide the high-resolution 10-m wind speed vectors needed to calculate the wind stresses using Eqs. (1)–(4). The bottom-right panel of Fig. 4 shows a strong negative linear correlation between SST and wind speed at the AXBT locations. The simulated wind speeds range from 2 to 33.2 m s^{-1} and thus cover a substantial portion of the C_D range shown in Fig. 1. This functional relationship forms the basis for the inverse modeling approach.

c. *Comparing simulated and observed AXBTs*

As a check on HYCOM's ability to simulate the oceanic response realistically, we present here a first comparison between the AXBT data and their model

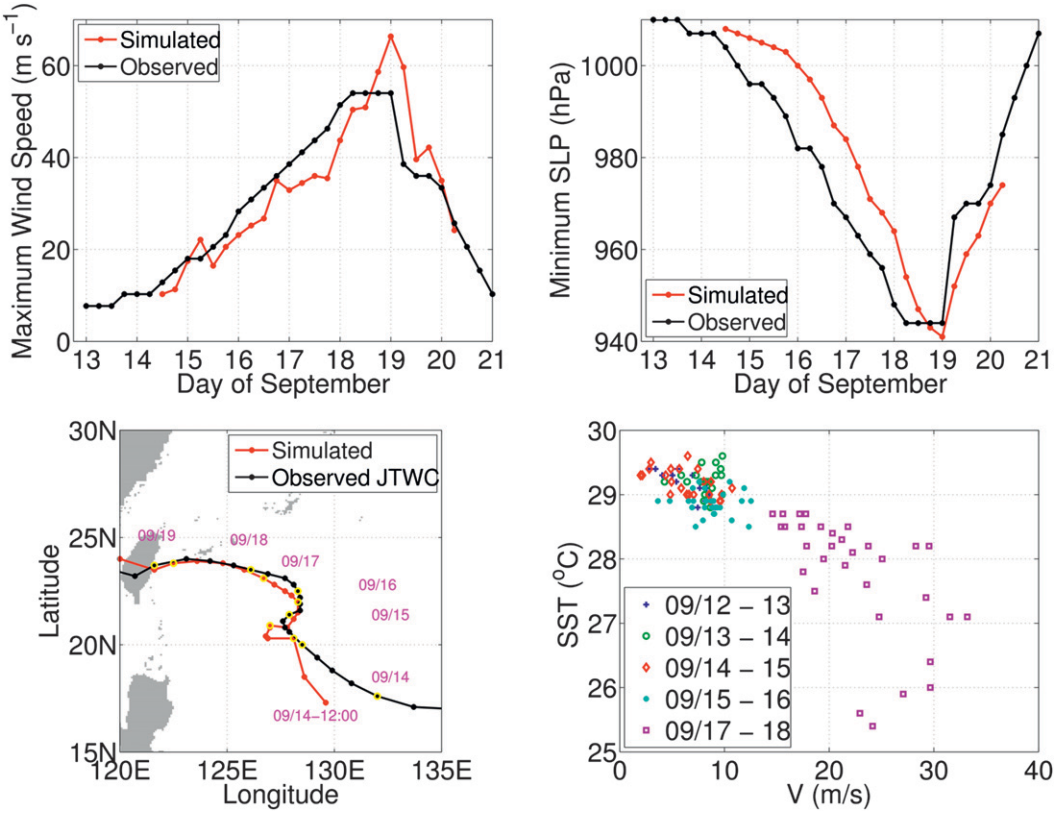

FIG. 4. Comparison of Fanapi's intensity [(top left) maximum wind speed and (top right) minimum sea level pressure] and (bottom left) track from high-resolution WRF simulation with JTWC's best-track data. (bottom right) SST vs the WRF wind speed at the AXBT locations for different dates.

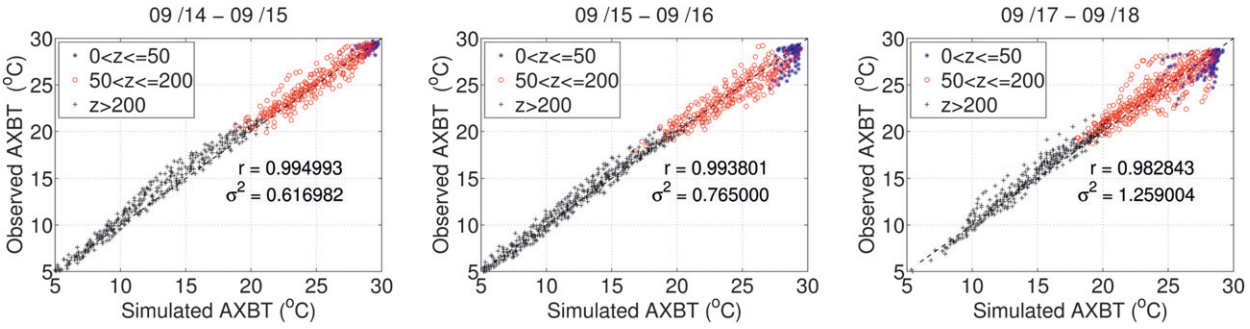

FIG. 5. (from left to right) Scatterplots of measured temperature (AXBTs) against their HYCOM counterparts during Typhoon Fanapi as it intensifies during 14–18 Sep. The data points are colored by the specified depth ranges (z in meters) as indicated in the legend. The correlation coefficient r and the error variance σ^2 is indicated on each plot.

counterparts for a reference simulation in which the drag parameters were set to their default values ($\alpha = 1$, $V_{\max} = 32.5\,\mathrm{m\,s^{-1}}$, and $m = 0$). The HYCOM initial conditions were obtained from the global 1/12° model, and the wind stresses were obtained from the triply nested WRF simulation; other HYCOM configuration settings are listed in the appendix. The scatterplots in Fig. 5 compare the observations and the HYCOM fields from 14 to 18 September as Fanapi intensifies. There is good agreement between the two datasets with no discernible trends until 17 September, a day after Fanapi becomes a typhoon. On the last day of flight operations (17–18 September, right plot in Fig. 5), the observations show lower temperature in the upper 50 m than HYCOM and a larger scatter in the depth range of 50–200 m, with cooler and warmer temperature differences. The difference between observed and simulated AXBTs has a maximum of 5°C. This difference is reflected in the correlation coefficients calculated for all data points of each day. The variance of the difference between observed and simulated AXBTs increases as Typhoon Fanapi intensifies. Bayesian inference, as detailed in the following section, will be used to estimate the ranges of drag parameters that minimize the scatter between observed and modeled ocean temperatures, and to estimate the variance of the noise in the measured data.

4. Bayesian inference

Let the temperature observations be arranged in a vector $\mathbf{T} = (T_1, T_2, \ldots, T_N)$; their model-predicted counterparts by $\mathbf{M} = (M_1, M_2, \ldots, M_N)$ and let $\mathbf{H} = \{\alpha, V_{\max}, m\}$ be the vector of control parameters. Bayes's theorem (Berger 1985; Epstein 1985; Berliner et al. 2003; Gelman et al. 2004; Bernardo and Smith 2007) can be invoked to reverse the conditional probabilities of the forward problem:

$$p(\mathbf{H}\,|\,\mathbf{T}) \propto p(\mathbf{T}\,|\,\mathbf{H})p(\mathbf{H}), \tag{10}$$

where $p(\mathbf{H})$ is the prior probability distribution function (pdf) of $\mathbf{H}$ and represents the a priori knowledge about $\mathbf{H}$ (before considering the observation data); $p(\mathbf{T}\,|\,\mathbf{H})$ is the likelihood and represents the probability of obtaining the data given the set of parameters $\mathbf{H}$; and finally $p(\mathbf{H}\,|\,\mathbf{T})$ is the posterior, which represents the probability that $\mathbf{H}$ is true given the data $\mathbf{T}$.

We adopt the following form for the likelihood function:

$$p(\mathbf{T}\,|\,\mathbf{H}) = \prod_{i=1}^{N} \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left[\frac{-(T_i - M_i)^2}{2\sigma^2}\right], \tag{11}$$

where we have assumed that the errors between model and observation, $\boldsymbol{\epsilon} = \mathbf{T} - \mathbf{M}$, are uncorrelated and normally distributed with mean zero and variance σ^2. The variance is unknown a priori; thus we treat it as a hyperparameter and augment the vector of control parameters to $\mathbf{H} = \{\alpha, V_{\max}, m, \sigma^2\}$. While in general σ^2 depends on the observations, in cases where the error amplitude is generally small and does not change throughout space and time, one may use a simplified model by assuming a spatially uniform σ^2. Furthermore, the typhoon wind forcing conditions vary significantly from one day to another and AXBT data collected on each day are exposed to different measurement errors; therefore, it is reasonable to use different variances for each day of data collection as indicated in section 6c.

The specification of the priors should reflect a priori knowledge about the parameters. In the absence of additional information regarding the parameters' distributions, a uniform prior is a reasonable choice:

$$p(\{\alpha, V_{\max}, m\}) = \begin{cases} \dfrac{1}{b_i - a_i} & \text{for} \quad a_i < \{\alpha, V_{\max}, m\} \le b_i, \\ 0 & \text{otherwise}, \end{cases} \tag{12}$$

where $[a_i, b_i]$ denote the parameter ranges defined in Table 1. The only information we know regarding the variance is that σ^2 is always positive; we thus assume a Jeffrey's prior (Sivia 2006) which can be expressed as

$$p(\sigma^2) = \frac{1}{\sigma^2} \quad \text{for} \quad \sigma^2 > 0, 0 \text{ otherwise}. \tag{13}$$

Consequently, Bayes's theorem leads to the following expression:

$$p(\mathbf{H} \mid \mathbf{T}) \propto \left\{ \prod_{i=1}^{N} \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left[\frac{-(T_i - M_i)^2}{2\sigma^2} \right] \right\} \times p(\sigma^2) p(\alpha) p(V_{\max}) p(m). \tag{14}$$

Note that our inference procedure accounts for uncertainty in the HYCOM drag parameters only, and does not include other sources of uncertainty such as HYCOM or WRF errors. As such, one could interpret the present formulation as a strongly constrained one. Accounting for model errors, or other sources of uncertainties such as initial conditions, is possible with the present methodology, but is beyond the scope of the present investigation.

Inferring the drag coefficient parameters requires sampling the posterior. In general, when the space of the unknown parameters is multidimensional, a suitable computational strategy is the MCMC method. We rely on an adaptive metropolis MCMC (Roberts and Rosenthal 2009; Haario et al. 2001) to sample the posterior distribution accurately and efficiently. This MCMC phase requires repeated (tens of thousands of) HYCOM simulations initialized with different values of the uncertain parameters; this step is prohibitively expensive. An alternative is to construct a surrogate model that can be used at a significantly reduced computational cost. Here we rely on PC expansions (briefly described in the next section) to build the surrogate, which, in addition, efficiently provides statistical properties, such as the mean, variance, and sensitivities.

5. Polynomial chaos expansion

Polynomial chaos expansions—more details of which can be found in Le Maître and Knio (2010)—are series-based methods that have been developed in the engineering community to represent uncertainty in the output of numerical simulations (Villegas et al. 2012; Lin and Karniadakis 2009; Xiu and Tartakovsky 2004). Let $R = R(\mathbf{x}, t, \boldsymbol{\xi})$ denote an output quantity of interest, such as temperature or velocity; this quantity depends on the usual space and time dimensions $(\mathbf{x}, t)$, but also on the uncertain variables $\boldsymbol{\xi}$ that went into specifying the model's input data (such as uncertain initial and boundary conditions, and parameters). PC expansions express the dependencies of R on the uncertain input variables $\boldsymbol{\xi}$ as a truncated expansion of the following form:

$$R(\mathbf{x}, t, \boldsymbol{\xi}) \doteq \sum_{k=0}^{P} \hat{R}_k(\mathbf{x}, t) \Psi_k(\boldsymbol{\xi}), \tag{15}$$

where $\hat{R}_k(\mathbf{x}, t)$ are unknown coefficients, P is finite and depends on the truncation strategy adopted, and the functions $\Psi_k(\boldsymbol{\xi})$ form an orthogonal basis of an underlying probability space. This series representation can be viewed as a spectral expansion of R along the stochastic dimensions. The Cameron–Martin theorem (Cameron and Martin 1947) asserts the existence and convergence of this series whenever R has finite variance. The series rate of convergence, and hence the number of terms to retain, depends on the smoothness of R with respect to $\boldsymbol{\xi}$. The series converges spectrally fast with P when R is infinitely smooth; the convergence rate becomes algebraic when R has finite smoothness (Canuto et al. 2006). In practice the series convergence is monitored via various error metrics; section 6a presents one possible error analysis for the present problem.

The choice of basis is dictated by the probability density function of the stochastic variable $\boldsymbol{\xi}$, which appears as a weight function in the stochastic space's inner product:

$$\langle \Psi_k, \Psi_m \rangle = \int \Psi_k(\boldsymbol{\xi}) \Psi_m(\boldsymbol{\xi}) \rho(\boldsymbol{\xi}) \, d\boldsymbol{\xi}. \tag{16}$$

When the basis functions are orthonormal $\langle \Psi_k, \Psi_m \rangle = \delta_{k,m}$, where $\delta_{k,m}$ is the Kronecker delta. For Gaussian distributions the basis functions are hence scaled Hermite polynomials whereas for uniform distributions, as is the case here, they are scaled Legendre polynomials. For multidimensional problems the basis functions are tensor products of 1D basis functions. The identification of the inner product weight function with the probability distribution of $\boldsymbol{\xi}$ simplifies the calculations of R's statistical moments. Noting that since $\Psi_0(\boldsymbol{\xi})$ is a constant that can be normalized to satisfy $\langle \Psi_0, \Psi_0 \rangle = 1$, the expectation and variance of R can be computed as

$$E[R] = \int R \rho(\boldsymbol{\xi}) \, d\boldsymbol{\xi} = \langle R, \Psi_0 \rangle = \hat{R}_0, \tag{17}$$

$$E[(R - E[R])^2] = \int (R - E[R])^2 \rho(\boldsymbol{\xi}) \, d\boldsymbol{\xi} = \sum_{k=1}^{P} \hat{R}_k^2 \langle \Psi_k, \Psi_k \rangle. \tag{18}$$

The series representation in Eq. (15) can thus be seen as combining approximation and probabilistic frameworks, a combination that has proven extremely useful in solving UQ problems.

A number of procedures have been devised to compute the expansion coefficients $\hat{R}_k$. Here we rely on the nonintrusive spectral projection (NISP) since it does not require any modification to HYCOM, and only ensemble runs at specified values of the uncertain parameters are needed. The NISP method is based on minimizing the $\mathcal{L}_2$ norm of the error between $R(\mathbf{x}, t, \boldsymbol{\xi})$ and its series representation, on exploiting the orthogonality of the basis, and on replacing the stochastic integrals by quadrature to obtain the following:

$$\hat{R}_k(\mathbf{x}, t) = \frac{\langle R, \Psi_k \rangle}{\langle \Psi_k, \Psi_k \rangle} \approx \frac{\langle R, \Psi_k \rangle_Q}{\langle \Psi_k, \Psi_k \rangle}. \tag{19}$$

The subscript Q refers to approximating the inner product integral with quadrature:

$$\begin{aligned} \langle R, \Psi_k \rangle &= \int R(\mathbf{x}, t, \boldsymbol{\xi}) \Psi_k(\boldsymbol{\xi}) \rho(\boldsymbol{\xi})\, d\boldsymbol{\xi} \approx \langle R, \Psi_k \rangle_Q \\ &= \sum_{q=1}^{Q} R(\boldsymbol{\xi}_q) \Psi_k(\boldsymbol{\xi}_q) \omega_q, \end{aligned} \tag{20}$$

where $\boldsymbol{\xi}_q$ and ω_q are multidimensional quadrature points and weights, respectively. The computation of the $\hat{R}_k$ can thus be expressed as a matrix-vector product of the following form:

$$\hat{R}_k(\mathbf{x}, t) = \sum_q \mathbf{P}_{kq} R(\mathbf{x}, t, \boldsymbol{\xi}_q), \quad \mathbf{P}_{kq} = \frac{\Psi_k(\boldsymbol{\xi}_q)\omega_q}{\langle \Psi_k, \Psi_k \rangle}, \tag{21}$$

where $\mathbf{P}_{kq}$ is the projection matrix and $R(\mathbf{x}, t, \boldsymbol{\xi}_q)$ is obtained from an ensemble of HYCOM *realizations* with the uncertain parameters set at the quadrature value $\boldsymbol{\xi}_q$. The NISP procedure can be carried out as a postprocessing operation using simple scripting languages like MATLAB, the only costs being those of the ensemble calculation and its storage.

The calculation of the $R(\mathbf{x}, t, \boldsymbol{\xi}_q)$ is the most expensive part of the inference procedure; thus, reducing the number of sampling points while maximizing their effectiveness is critical for the procedure's efficiency. The quadrature order should be commensurate with the truncation order, and should be high enough to avoid aliasing artifacts. The choice of quadrature rule is hence critical to the performance of the PC (in its NISP version at least). Tensorized Gaussian quadratures are one way to proceed, but their computational cost increases exponentially with the stochastic dimension of the problem (curse of dimensionality). An alternative is based on sparse Smolyak quadratures (Gerstner and Griebel 2003; Petras 2000, 2001, 2003) which require a smaller ensemble size, and that can be configured in a nested hierarchical structure so that additional sampling levels can be added incrementally without discarding previous calculations. Alexanderian et al. (2012) relied on Smolyak quadrature in their Hurricane Ivan UQ study. A level 5 Smolyak quadrature, requiring an ensemble with 385 members, was deemed necessary to obtain an accurate surrogate. One drawback in that calculation is that the space of random variables is sampled isotropically along all stochastic dimensions, whereas sensitivity analysis indicated that the drag coefficient uncertainty dominated the SST variance.

Given that the present HYCOM computational burdens are larger than the Gulf of Mexico simulations in Alexanderian et al. (2012) (larger domain and longer integration), we are motivated to adopt a recently developed methodology based on adaptive basis refinement and quadrature (Constantine et al. 2012; Winokur et al. 2013; Conrad and Marzouk 2012, manuscript submitted to *SIAM J. Sci. Comput.*). The new methodology uses a pseudospectral extension of the Smolyak algorithm that has the virtues of (i) enabling a general nonisotropic refinement of the sparse quadrature grid, (ii) steering the realizations toward the stochastic dimensions producing the most variance, and (iii) retention of all members of the PC basis that can be computed with no internal aliasing. The current implementation follows the construction outlined in Constantine et al. (2012), Winokur et al. (2013), and Conrad and Marzouk (2012, manuscript submitted to *SIAM J. Sci. Comput.*) using the nodes of the Gauss–Patterson quadrature rule. A variance-based criterion was used to adaptively refine the sparse grid, using the area-averaged SST in the control box (Fig. 6) on 18 September as the quantity of interest. The adaptive algorithm required 6 refinement steps, resulting in 67 independent realizations. The adapted grid had more nodes in the direction of the α parameter due to its larger impact on the SST variance (see Fig. 6).

6. Results

The results of the forward and inverse problems are presented and discussed in three sections. First, we perform an error analysis of the adaptive PC expansion construction, namely, to establish the validity of the corresponding surrogate. Second, we perform a sensitivity analysis and rank the impact of the different drag parameters on the predicted ocean temperature. Finally,

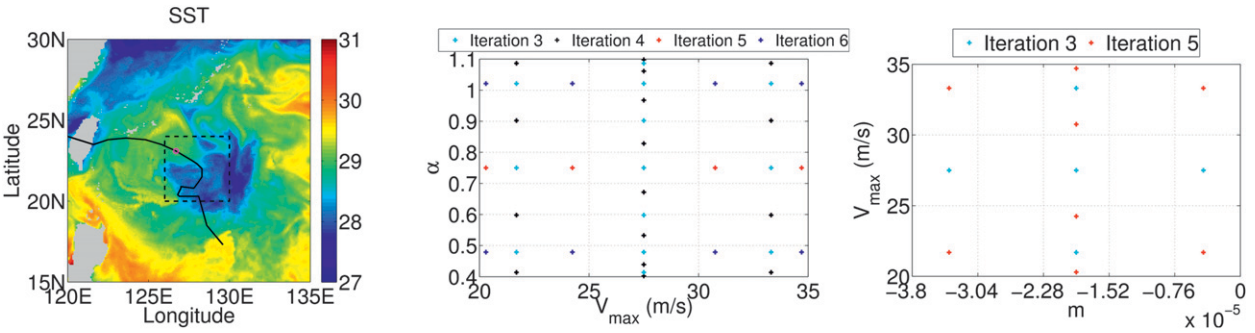

FIG. 6. (left) Contours of SST from a single realization $\boldsymbol{\xi} = (0, 0, 0)$ on 0000 UTC 18 Sep. A control box covering part of the typhoon track used for the construction of adaptive quadrature is also shown. (middle), (right) Two-dimensional projections of the nodes of the adaptive quadrature at different refinement levels.

posterior distributions for the drag parameters are determined and analyzed in light of the available AXBT data.

a. Error and convergence study

Unless otherwise noted, the results presented in the present error analysis are based on temperature fields averaged over the analysis box shown in Fig. 2. The center of this box (24.4°N, 125.7°E) coincides with one of the AXBT data points collected on 18 September that is close to Fanapi's track and experiences high wind speeds; the box size is chosen to include about seven HYCOM grid cells. Figure 7 shows the evolution of the area-averaged ocean temperature for the 67

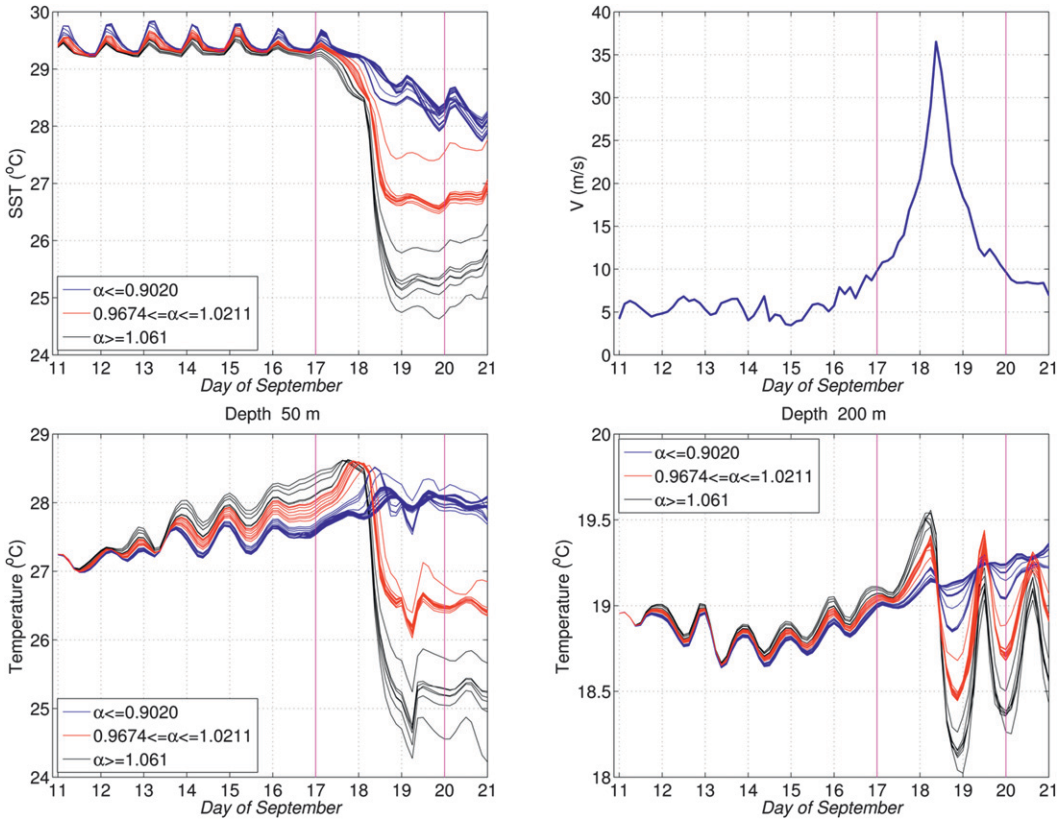

FIG. 7. Evolution of area-averaged ocean temperature for the different realizations at various depths: (top left) surface, (bottom left) 50 m, and (bottom right) 200 m. The temperature is color coded according to α values. (top right) The wind speed experienced at the area center, and the vertical lines in all panels indicate the times when Fanapi crosses the analysis area.

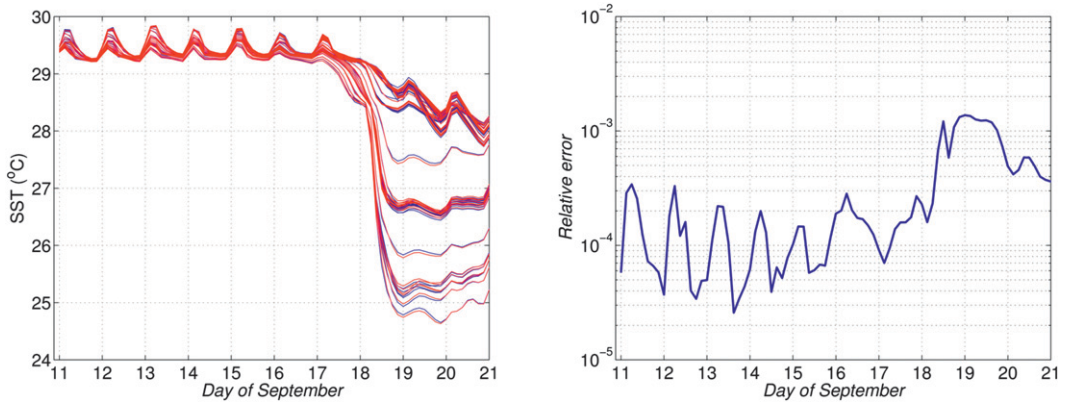

FIG. 8. (left) Evolution of the area-averaged SST realizations (blue) and of the corresponding PC estimates (red). (right) The relative normalized error calculated using Eq. (22).

realizations obtained using the adaptive PC expansion truncation. SST, which had been fluctuating around 29.5°C due to diurnal heating, cools by an amount that ranges from 1° to 4°C during Fanapi's passage, with the temperature drop increasing with α (top-left panel in Fig. 7). The maximum cooling occurs on 18 September when the wind speed peaks at 35 m s^{-1} as shown in the top-right panel of Fig. 7. The banded structure of the SST response is a first indication that the α parameter has a greater influence than the other two parameters. The temperature drop at 50-m depth ranges from 0°C for small α to 3°C for large α. The 50- and 200-m temperatures experience a warming trend prior to Fanapi's passage, with, surprisingly, warmer temperatures for the higher α values. At 200 m, the cooling associated with Fanapi's passage is totally absent for small α and reaches 1°C for large α. Note also that the 50- and 200-m panels display the same banded structure as SST.

To check the consistency of the approximation, in Fig. 8 we compare SST from the realizations with those obtained from the PC surrogate. The corresponding curves show an excellent agreement for all times including during high-wind periods. We define a global (in stochastic space) error metric to measure the relative normalized root-mean-square error between the left-hand side function in Eq. (15) and its PC representation at the sampling points:

$$E = \frac{\left[\sum_{\boldsymbol{\xi}\in\mathcal{S}} \left| T(\boldsymbol{\xi}) - \sum_{k=0}^{P} T_k \Psi_k(\boldsymbol{\xi}) \right|^2\right]^{1/2}}{\left[\sum_{\boldsymbol{\xi}\in\mathcal{S}} |T(\boldsymbol{\xi})|^2\right]^{1/2}}, \qquad (22)$$

where $\mathcal{S}$ is the 67-member ensemble. This error metric calculated over the analysis box is shown in the right panel of Fig. 8; the largest relative normalized error for SST is about 0.1% and occurs on 19 September when Fanapi reaches its maximum wind intensity. Contour maps of the relative normalized error for the entire simulation region are shown in Fig. 9 for various depths and dates to confirm the error trends of the analysis box. The error is largest after the typhoon intensifies (bottom row) for all depths. On both days, the maximum error is located at 50 m, which coincides with the depth of the original mixed layer measured by the AXBT. The maximum error magnitude recorded is about 1% and occurs on 18 September. The elevated error region is located to the right of the storm and extends from the surface down to 50 m, below which the impact of the input uncertainty decreases substantially. For further validation, Fig. 10 compares HYCOM vertical temperature profiles with those of selected AXBT observations on 14, 15, and 17 September. The plots show good agreement between the simulations and the observations at different locations and at different times.

Another type of consistency check is concerned with the convergence of the adaptive procedure used in the calculation of the PC expansion coefficients, and with its ability to improve their estimates with increased refinement levels. Since changes to the series coefficients impact the temperature statistical properties, in Fig. 11 we monitor the evolution of the probability density function of temperature[6] as the refinement level is increased. The different curves correspond to increased adaptive refinement levels (from 3 to 6). The plots indicate that the double-peaked distributions are not well resolved at level 3 but become weakly insensitive to

[6] The temperature pdf was calculated by randomly sampling the input uncertainties $\boldsymbol{\xi}$, summing the surrogate to obtain the corresponding temperatures, building a histogram of the ensuing distribution, and computing the final pdf.

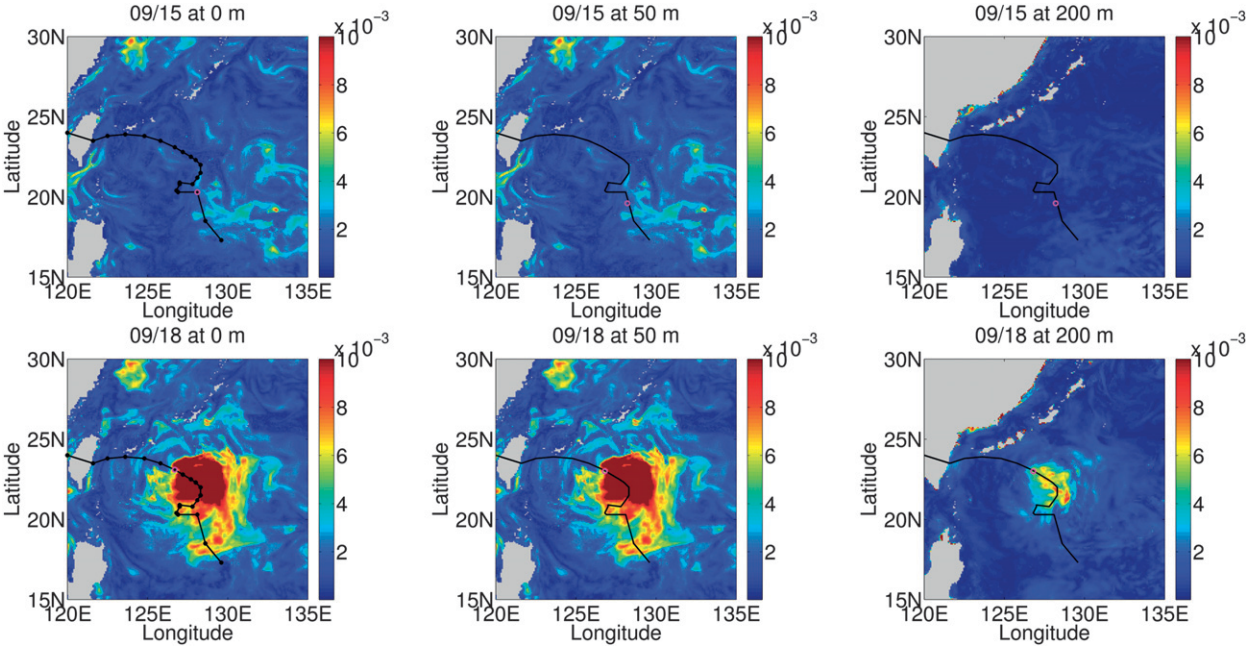

FIG. 9. Relative normalized error between realizations and the corresponding PC surrogates at different depths: (left) surface, (middle) 50 m, and (right) 200 m. (top) 0000 UTC 15 Sep and (bottom) 0000 UTC 18 Sep.

further refinement starting with level 4; note also that convergence is faster for the 200-m temperature than for SST. Figure 11 also includes the pdf obtained using an unadapted PC expansion series and calculated with an isotropic Smolyak quadrature with 4 levels (total of 87 realizations). This unadapted pdf is closer to the level 3 adaptive PC expansion but misses many of its details. Notice that all pdfs exhibit multiple peaks and extended tails either toward lower temperatures at the surface, or toward higher temperatures at 50- and 200-m depths.

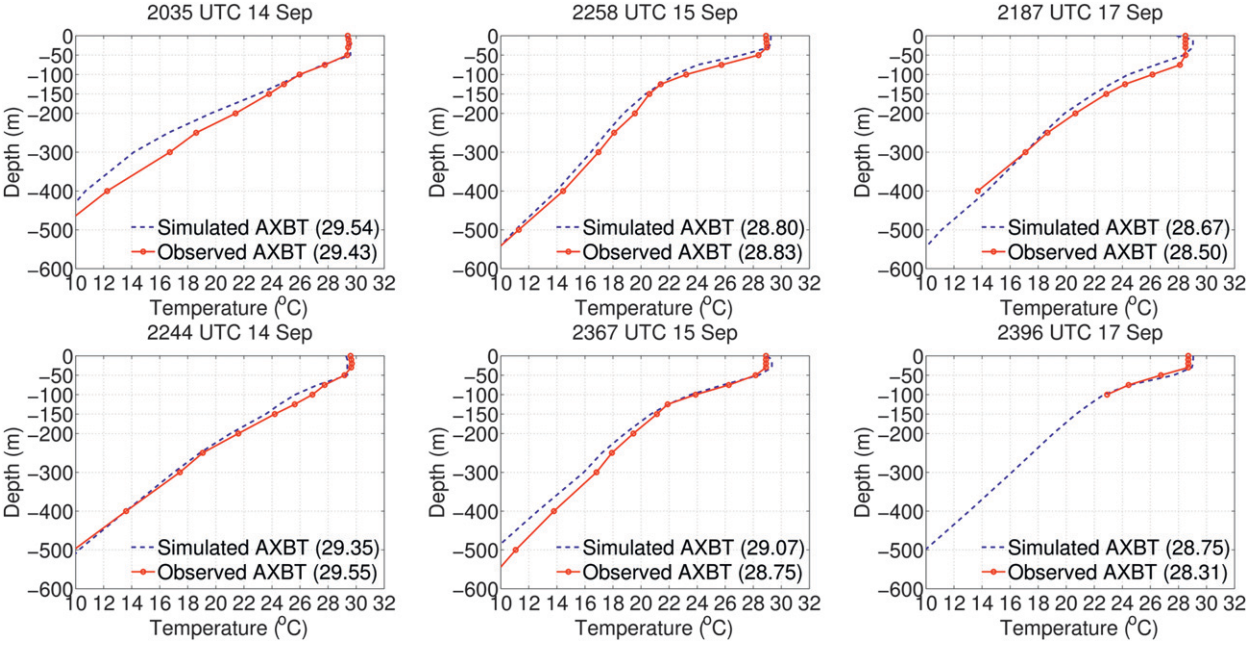

FIG. 10. Comparison of HYCOM vertical temperature profiles with AXBT observations on (left) 14, (middle) 15, and (right) 17 Sep. Temperature averages over the first 50 m are shown in the parentheses.

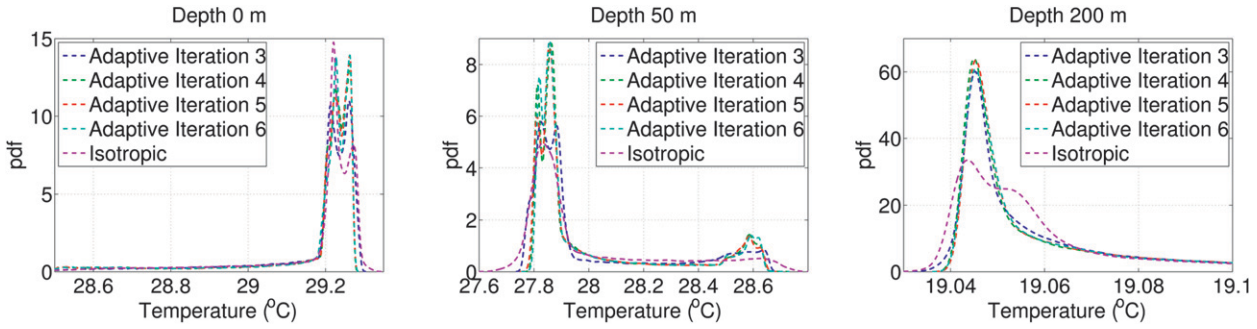

FIG. 11. Temperature pdfs using adaptive and isotropic PC refinement, as indicated. Plots are generated for 0000 UTC 18 Sep, and different depths: (left) surface, (middle) 50 m, and (right) 200 m.

The various error metrics and model–data comparisons presented above provide confidence that with six refinement levels the adaptive PC expansion is a faithful model surrogate. They also illustrate the efficiency of the adaptive approach, and its advantages over isotropic refinement.

b. *Statistical and sensitivity analysis*

We now exploit the PC surrogate to study the statistics of the temperature, and to quantify its sensitivity to the uncertain input parameters and thus to anticipate their impact on the inverse problem. In doing so, it is emphasized that no additional HYCOM simulations were needed to obtain the information presented below; rather, this information is obtained either directly from the coefficients of the PC surrogate or by sampling the corresponding representations for different values of $\boldsymbol{\xi}$.

The left panel of Fig. 12 shows the evolution of the analysis box-averaged SST and its two standard deviations bounds. The standard deviation is insignificant during the first few days of the simulation but increases as the intensifying typhoon approaches the analysis area. The temperature drop is clearly observed after Fanapi's passage. To quantify the contribution of each uncertain parameter to the variance in SST, we calculate the total sensitivity index using the PC coefficients (Alexanderian et al. 2012; Sudret 2008; Crestaux et al. 2009). The total sensitivity index of each of the uncertain parameters is shown in the right panel of Fig. 12. Throughout the simulation, the multiplicative drag factor α is dominant and contributes most to the SST variance compared with the two other drag parameters. The saturation wind speed $V_{\max}$ exhibits a small sensitivity index before the typhoon intensifies and the drag coefficient slope (after saturation) m appears to be an insignificant contributor to the variance.

The same picture emerges when one examines the temperature statistics below the surface. The left panel of Fig. 13 shows the time evolution of the PC-mean vertical temperature profile, and reveals a mixed layer depth of about 50 m and an average cooling of about 2°C once Fanapi arrives. The right panel shows the associated standard deviation and the confinement of the uncertainties to the upper 50 m. The total sensitivity indices

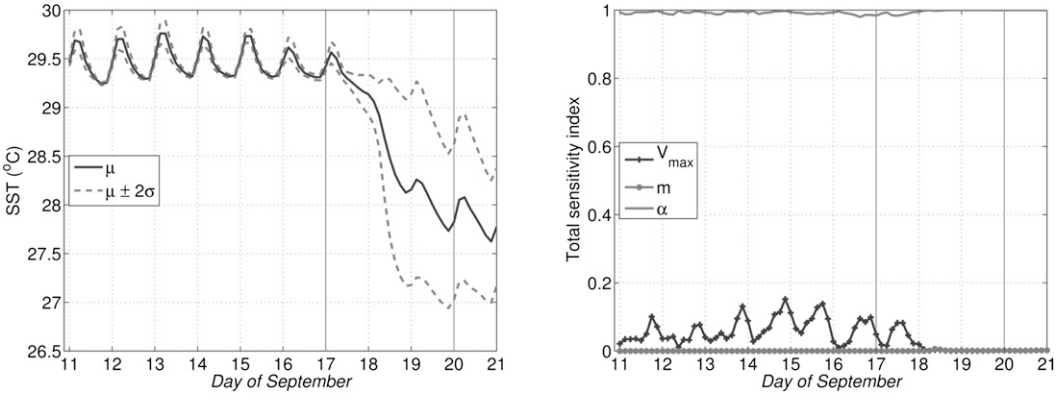

FIG. 12. (left) Evolution of area-averaged SST bounded by two standard deviations. (right) Evolution of the total sensitivity indices of the three uncertain parameters. Fanapi's passage over the analysis area is indicated by the two vertical lines.

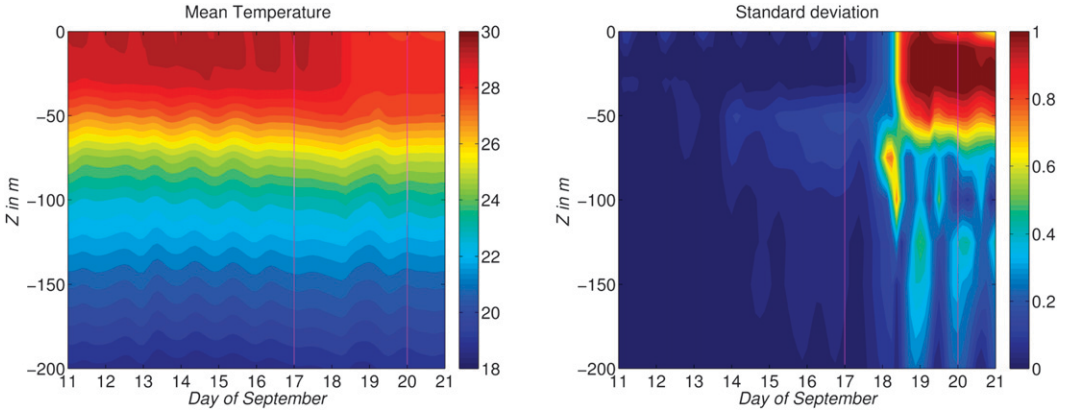

FIG. 13. Contours of (left) mean temperature and (right) its standard deviation.

(not shown) indicate that α contributes most to the uncertainty in temperature throughout the upper 50 m.

The area-averaged SST dependence on α and $V_{\max}$ (for fixed $m = 0$) on 17, 18, and 19 September is shown in Fig. 14, which displays the SST response surfaces as contour plots. The most striking features are the relatively flat horizontal contours on 17 and 18 September, indicating that SST depends only mildly on $V_{\max}$ even during peak winds; unsurprisingly, these contours turn completely flat on 19 September (and afterward) when the winds dip below 10 m s^{-1}. On the other hand, α exerts a strong influence on SST even under mild wind conditions. This influence is relatively weak for the low α range and increases substantially for higher α as evidenced by the packed contours. The temperature response at 50- and 200-m depths (not shown) exhibits roughly the same structure as at the surface, with milder dependence on $V_{\max}$ even during the peak winds of 18 September.

c. *Inferring drag coefficient parameters*

We now exploit the surrogate model in the Bayesian inference of the drag coefficient parameters. As mentioned in section 4, an adaptive MCMC method is used to sample the posterior distributions (Roberts and Rosenthal 2009; Haario et al. 2001) and consequently update the drag parameters' distributions in light of the AXBT data. This sampling, demanding tens of thousands of forward simulations, would have been prohibitively expensive in the absence of the surrogate because the generation of each sample would have required an independent HYCOM realization. The surrogate provides a computationally efficient alternative, and requires only summing the PC expansion for different values of the seed $\boldsymbol{\xi}$. The setup of the Bayesian inference problem, MCMC chains for the input drag parameters and the corresponding posterior distributions are presented and discussed in this section.

The AXBT data provided on different days (12–17 September) and at different depths were used in the likelihood function [Eq. (11)] to update the input parameters. Because the data collected on each day are exposed to different measurement errors, the variance between the data and the model were assumed to be different for each set of AXBT data. MCMCs of 100 000 iterations are obtained for the drag parameters: α, $V_{\max}$, and m as well as for the variance σ^2 on different days. Figure 15 (top row) shows the drag parameters' sample

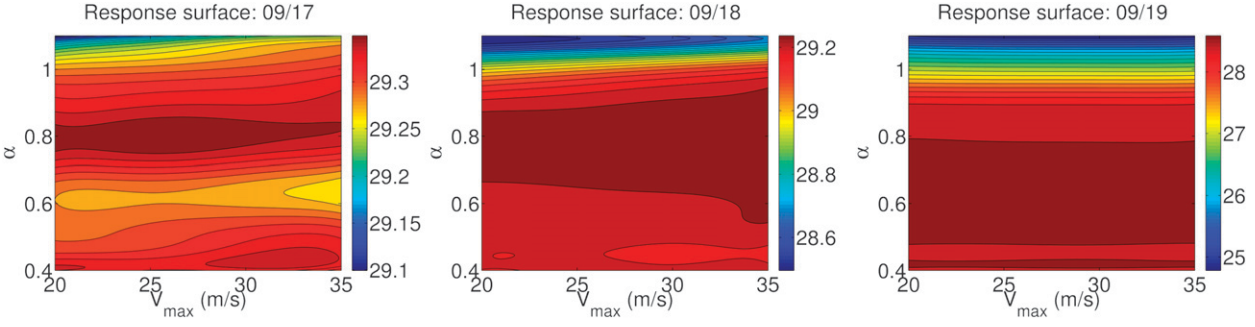

FIG. 14. SST response surface as function of α and $V_{\max}$, with fixed $m = 0$. Plots are generated on different days, as indicated.

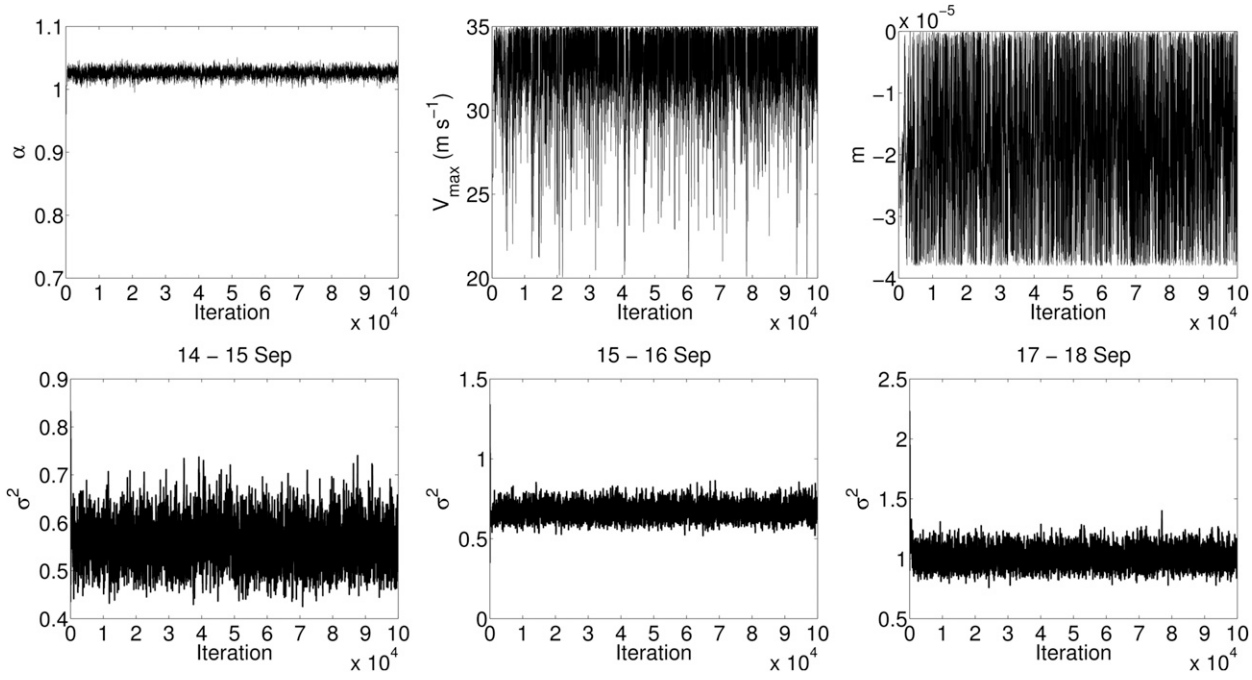

FIG. 15. (top) Chain samples for (left) α, (middle) $V_{\max}$, and (right) m. (bottom) Chain samples for σ^2 generated on different days, as indicated.

chains for different iterations of the MCMC algorithm. The top-left panel shows a well-mixed chain for α. The $V_{\max}$ chain appears to be concentrated in the higher end of the parameter range, with values between 25 and 35 m s^{-1}. Finally, the chain for m spans the entire range of the prior, and so it appears that the observations are not informative concerning this uncertain input. The chains for the variances at selected days are also shown in Fig. 15 (bottom row) and appear to be well mixed.

The computed MCMC chains can be readily used to determine the posterior distribution; kernel density estimation (KDE) is used for this purpose (Parzen 1962; Silverman 1986). (The first 1500 iterates, associated with the burn-in period, are discarded.) The resulting posterior pdfs of the drag parameters are shown in Fig. 16 (top row). As expected from the chains shown in Fig. 15, the posterior pdf of $V_{\max}$ exhibits a well-defined peak, with a maximum a posteriori (MAP) estimate around 34 m s^{-1}; an extended tail toward the lower velocity values is also observed. Note that the MAP occurs near the edge of the assumed prior of $V_{\max}$, which as previously noted represented our best estimate of the range of variability of the saturation wind speed before conducting the analysis. However, as shown in Fig. 6 (middle), the MAP coordinates are effectively bracketed by members of the sparse grid. Also note that the MAP value is slightly higher than the 32.5 m s^{-1} used operationally in HYCOM. In contrast, the posterior pdf of m appears to be fairly flat and similar to the prior. This is an indication that the observational data were not useful to refine our prior knowledge of m. This could be explained by the fact that winds at the AXBT locations peaked at maximum speeds of 33.2 m s^{-1} (Fig. 4, bottom right), lower than the inferred $V_{\max}$ whereas m affects C_D for wind speeds higher than $V_{\max}$. It thus appears that in the present setting, nonlocal effects of the high winds experienced within the domain do not seem to affect the temperature profiles at the AXBT locations, and that, consequently, direct measurements at higher wind speeds would be desirable. Finally, the α posterior shows a clear localization compared with the prior, and the distribution appears to be normal and centered at $\alpha = 1.026$. This stretching of the drag curve relative to the operational HYCOM one may be related to the fact that, over the period in which AXBT data are available, the simulated wind speeds are slightly below the observed wind speeds (Fig. 4). The posterior distribution of the variance is also shown in Fig. 16 (bottom row) for selected dates. The variances corresponding to the MAP estimates of the parameters are lower than those obtained with the HYCOM operational values (listed in Fig. 5): 0.6741 versus 0.765 on 15–16 September and 1 versus 1.26 on 17–18 September. Thus, the MAP values have reduced the discrepancies between simulated and observed temperature. The square root of the MAP variances—the temperature standard deviation—is a

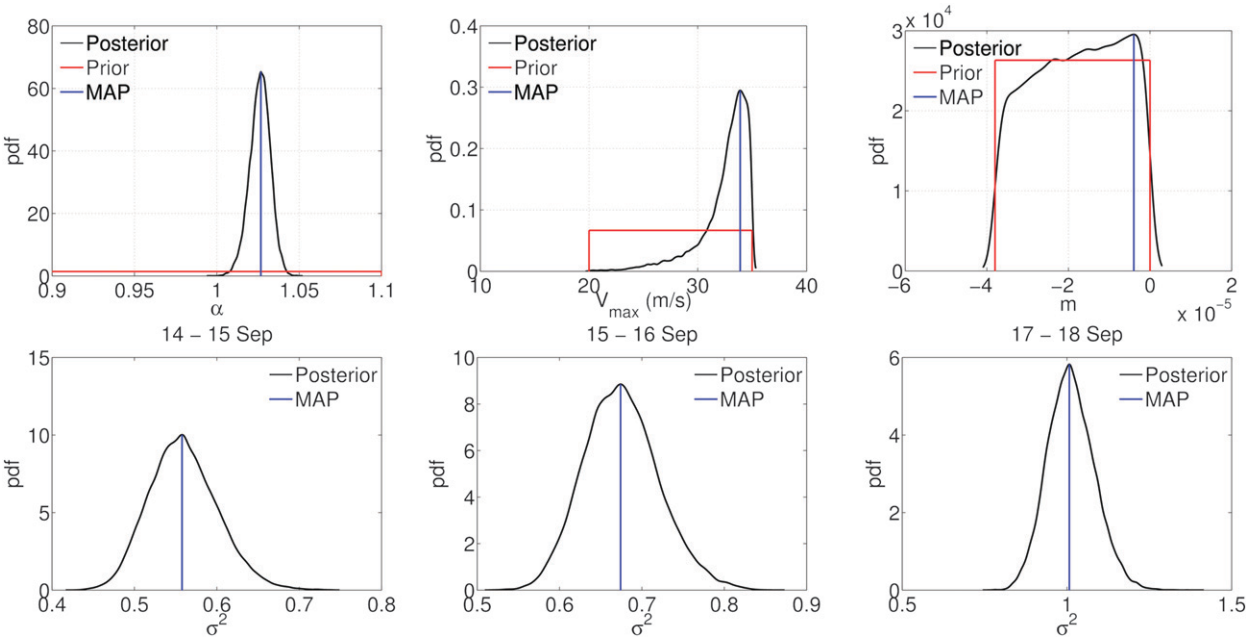

FIG. 16. (top) Prior and posterior distributions for the drag parameters and (bottom) posterior distributions of the variance between simulations and observations on different dates. MAP values of the posterior distributions are also indicated.

reflection of the mismatch between model and AXBT data. This mismatch grows slowly from 0.75°C on 14 September to about 1°C on 17 September.

The Kullback–Leibler (KL) divergence is a statistical measure that quantifies the distance between two probability distributions p and q (Kullback 1959); it is defined as

$$D_{\mathrm{KL}}(p\|q) = \int_{-\infty}^{\infty} p(x) \ln \frac{p(x)}{q(x)} dx. \qquad (23)$$

The variable D_{KL} satisfies most of the intuitive notions of a distance (i.e., it is nonnegative, it is zero if p and q are identical, and increases as p and q grow further apart), except that it is not symmetric in p and q and does not satisfy the triangle inequality. When p and q are identified with the posterior and prior distributions of a Bayesian inference problem, D_{KL} becomes a useful measure of the information gain contributed by the new data to the prior knowledge. The KL information gains due to the AXBT observations for α, $V_{\max}$, and m are 3.5, 1, and 0.087, respectively. The largest gain is thus associated with α, followed by saturation wind speed; the present AXBT data appears to be largely uninformative concerning the drag coefficient slope after saturation, m. These values reflect the intuitive notions suggested by Fig. 16 whereby the posterior and prior of m are very similar whereas those of α and $V_{\max}$ are substantially different.

With the MCMC chains computed, we can now construct the joint posterior distributions of the parameters using KDE. Since the posterior of m is close to the prior, we show in Fig. 17 the joint posterior distributions of α and $V_{\max}$ (left panel), and of α and α^2 on 17–18 September (right panel). These pdfs clearly show a single peak located at $\alpha = 1.026$ and $V_{\max} = 34\,\mathrm{m\,s^{-1}}$, and that can be identified with the MAP estimates from the marginalized pdfs. The right panel also clearly shows a tight estimate for α with little spread around it. Setting $m = 0$ (default HYCOM value), and substituting the MAP estimates for $V_{\max}$ and α in Eq. (5), we obtain the "optimal" (most likely) curve describing the drag coefficient C_D's dependence on wind speed; this curve is shown in Fig. 18. In addition to decreasing the variance between simulations and observations, the impact of the MAP values on temperature can be assessed with the help of the surrogate. The results for the area-averaged SST and mixed layer depth (other quantities of interest can be treated similarly) are displayed in Fig. 19. The default drag parameters underestimate the cooling by 0.5°C and the mixed layer depth by as much as 15 m; note that the mixed layer depth feels the impact of the increased drag from 18 to 20 September, otherwise, the two curves overlap.

7. Discussion and conclusions

The present study relied on AXBT data and on Bayesian inference to improve our understanding of drag parameterization at high wind speeds. It concentrated on three uncertain parameters: a multiplicative factor

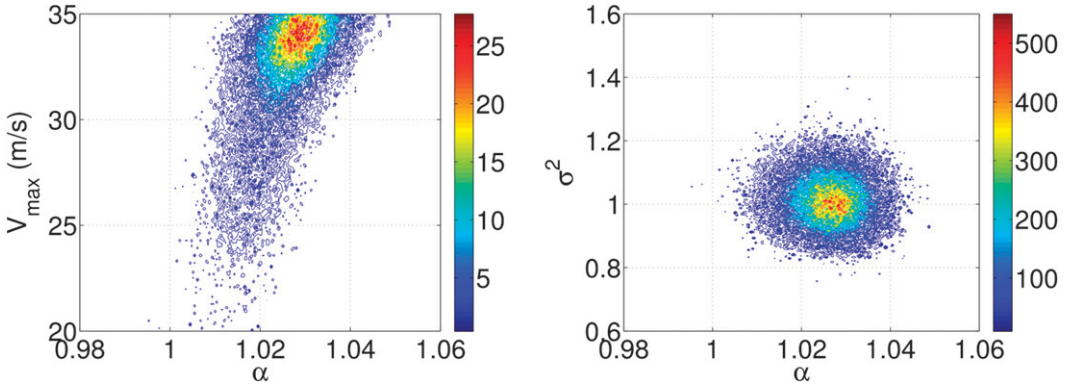

FIG. 17. (left) Joint posterior distribution of α and V_{max}. (right) Joint posterior of α and σ^2, generated for 17–18 Sep.

modulating the drag coefficient's magnitude α, the maximum wind speed at which saturation occurs (V_{max}), and the slope of the drag coefficient after saturation (m). The implementation of our research program hinged on a number of critical ingredients including (i) the ITOP that yielded the AXBT data; (ii) the WRF hindcasts to provide the high-resolution realistic atmospheric forcings for the ocean model; (iii) a robust and realistic ocean model, HYCOM, with a reliable mixed layer parameterization and realistic initial conditions for the forward model; and (iv) the adaptive PC expansions that produced a faithful and efficient surrogate at the modest cost of a 67-member HYCOM ensemble, and that made the Bayesian inference feasible. The main outcomes of this effort are the identification of the following MAP estimates for the drag parameters: that the drag coefficient saturates at about 2.3×10^{-3} with a saturation wind speed of at least $34\,m\,s^{-1}$, and that the available AXBT data was not informative with regard to the potential drop in the drag coefficient beyond saturation. Furthermore, the posterior probability density displayed a sharp localization for α with a Gaussian-like profile that decayed very fast away from its maximum, whereas the V_{max} posterior was skewed, decaying very fast for values greater than $34\,m\,s^{-1}$ but slowly for smaller values.

A number of recommendations can be made concerning the collection of AXBT data, and their influence on the drag parameters' posterior distributions. First, the inferred MAP value for V_{max} abuts the edge of the interval explored, and the situation is exasperated by the lack of AXBT data with wind speeds higher than $33.2\,m\,s^{-1}$; this value is substantially smaller than the maximum speeds recorded in the best-track data ($54\,m\,s^{-1}$) and in the WRF simulation ($60\,m\,s^{-1}$). It would thus be extremely useful if future observational programs are able to drop AXBTs at locations where the wind speed exceeds $35\,m\,s^{-1}$; future inverse modeling would also do well to consider a wider prior provided that data at higher speeds is available. Second, temperature measurements near the surface are more valuable than those at depth since the wind stress influence on temperature is weak at depths below 100 m (Fig. 13). The PC expansion and Bayesian inference methodologies can be configured to investigate the optimal deployment of AXBT data (i.e., targeted observations) so as to maximize the information gain with respect to specific questions (Huan and Marzouk 2013).

The present study highlights a number of advantages to relying on Bayesian inference to solve the inverse modeling problem. These include no a priori assumptions on the model's differentiability as would be required by adjoint and optimization methods, and no restriction on the statistics of the input and output as would be required in a Kalman filter approach. The Bayesian approach yields posterior probability distributions that encapsulate all statistical information about

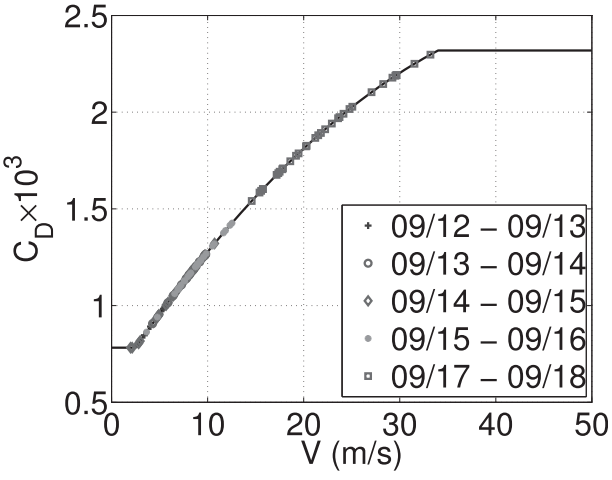

FIG. 18. Optimal drag coefficient C_D using MAP estimate of the three drag parameters. The symbols refer to AXBT data used in the Bayesian inference.

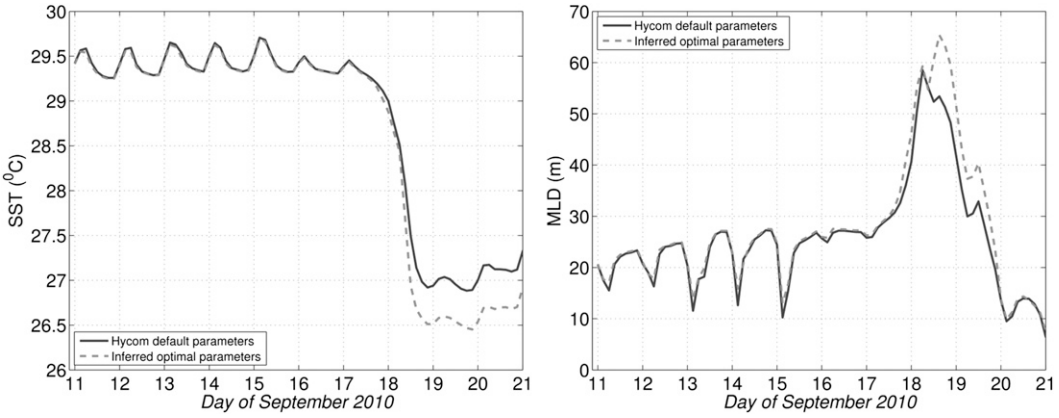

FIG. 19. Comparison of the area-averaged (left) SST and (right) MLD evolution during Fanapi using HYCOM default drag parameters and inferred optimal parameters.

the parameters at hand. The availability of prior and posterior probabilities permits us to quantify the information gain, as measured by the KL divergence, brought in by the AXBT data. Finally, it is important to point out that the PC expansion approach allows us to probe the dependency of output uncertainties on input uncertainties and to construct explicitly their response surfaces, thus yielding important insights into the inverse problem.

The present study does not address a number of issues relevant to tropical cyclones and ocean modeling. Chief among those is the reliance on an ocean-only model to carry out the forward modeling even though the air–sea interaction problem is intrinsically a coupled one requiring the use of, at least, coupled ocean–atmosphere models, if not ocean–wave–atmosphere models. It was deemed best, however, for the purposes of the present initial exploration, to defer running expensive coupled models to a later stage once the approach's potential was demonstrated in a relatively simpler context. The present UQ and inverse analysis considered only momentum coupling between the ocean and atmosphere and did not include the heat flux exchange between the two systems as mediated by the enthalpy coefficient. The heat flux exchange (and the behavior of the enthalpy coefficient) at high wind speeds, and in the presence of sea spray, is the subject of heated debate among tropical cyclone scientists, particularly concerning its role in modulating intensity changes and how to represent it realistically in hurricane models (Emanuel 1995; Andreas and Emanuel 2001; Emanuel 2003; Drennan et al. 2007; Haus et al. 2010; Andreas 2011; Jeong et al. 2012; Bell et al. 2012). This type of inquiry is best carried out using a coupled atmosphere–ocean model where the impact of the heat flux uncertainty on tropical cyclone intensity can be addressed directly.

Acknowledgments. This research was supported by the Office of Naval Research Award N00014-101-0498 and Department Research Initiative (DRI) ITOP Grant N00014-08-1-0576, by the U.S. Department of Energy, Office of Advanced Scientific Computing Research, Award DE-SC0007020, and by the Gulf of Mexico Research Initiative Contracts SA1207GOMRI005 (CARTHE) and SA12GOMRI008 (DEEP-C). We thank Jeff Kerlings, Robert Lee, Peter Black, and the 53rd Weather Reconnaissance Squadron of the Air Force Reserve Command for the AXBT deployment during ITOP. We would also like to thank the anonymous reviewers for their helpful comments on improving this manuscript.

APPENDIX A

Atmospheric Model Configuration

Realistic atmospheric forcing fields are required to accurately simulate the oceanic response to tropical cyclones. Present global atmospheric or coupled ocean–atmosphere models run by major operational centers have coarse space–time resolutions, and as such, cannot capture the dynamically relevant small-scale wind stress gradients in tropical cyclones; an alternative is to use high-resolution hindcasts to improve the representation of relevant dynamical features. WRF, version 3.2.1 (Skamarock et al. 2008), was used in a triply nested domain configuration to simulate the atmospheric conditions during Typhoon Fanapi.

The horizontal grid spacings were 12, 4, and 1.3 km and had 600 × 445, 301 × 301, and 301 × 301 points, respectively. The outer domain spanned the region 3°S–47°N, 97°E–180°, and was centered at 23.5°N, 140°E; the two inner nests were storm-following moving domains,

similar to Davis et al. (2008). The model had 36 vertical levels with the lowest level at 17 m from the surface, and 10 levels within the lowest 1 km. The moist convection is explicitly represented using the WRF single-moment 5-class (WSM5) microphysical scheme with rain, snow, cloud water, cloud ice, and water vapor as prognostic variables (Hong and Chen 2004) on the two inner domains, while the Kain–Fritsch cumulus parameterization scheme (Kain and Fritsch 1993) is used on the outer domain. The Yonsei University (YSU) boundary layer scheme (Hong et al. 2006) was used with wind-dependent surface roughness based on Donelan et al. (2004) and surface heat exchange coefficients based on Garratt (1992).

WRF was initialized with National Centers for Environmental Prediction (NCEP) analysis fields. SST was updated daily using the NCEP SST analysis field throughout the WRF simulation of Typhoon Fanapi. Also, 6-hourly nudging of zonal and meridional winds was performed using the NCEP analyses on the outer domain to ensure that the large-scale steering flow for Fanapi is as close to reality as possible, and that, therefore, the track of the simulated storm is close to that of the best-track data. The simulation period was from 0000 UTC 11 September to 0000 UTC 21 September. The outer domain was used for the first four days to spin up the precursor of Fanapi. The two inner-moving domains were used from 0000 UTC 15 September to 1800 UTC 20 September at which point the storm made landfall in China.

APPENDIX B

HYCOM Configuration

The HYCOM (Bleck 2002; Chassignet et al. 2003, 2007; Halliwell 2004) computational domain is essentially a nested subregion of the operational Global HYCOM (Chassignet et al. 2009), spanning the western North Pacific Ocean, 1°21′3″–43°8′51″N, 109°16′35″–170°4′35″E. Its horizontal grid resolution is $^1/_{12}$° at the equator with a total of 761 × 583 grid points (~6.5-km spacing on average). The model is configured with 27 layers with σ_2 target densities of 28.10, 28.90, 29.70, 30.50, 30.95, 31.50, 32.05, 32.60, 33.15, 33.70, 34.25, 34.75, 35.15, 35.50, 35.80, 36.04, 36.20, 36.38, 36.52, 36.62, 36.70, 36.77, 36.83, 36.89, 36.97, 37.02, and 37.06. The top model layer thickness is fixed to 3 m. The bottom topography is based on a quality controlled 2.5-minute gridded elevations/bathymetry for the world (ETOPO2.5) dataset with the western and northern edges land masked and islands clipped. The subgrid-scale parameterization adopted for the ocean mixed layer is *K*-profile parameterization (KPP; Large et al. 1994; Halliwell 2004). The three KPP parameters that were perturbed in Alexanderian et al. (2012) were held fixed at their reference values of 0.25 for the critical Richardson number, and at 10^{-4} and 10^{-5} $m^2 s^{-1}$ for the background viscosity and diffusivity, respectively.

HYCOM was initialized from a 10 September 2010 global HYCOM analysis state and run in a one-way nested (but free running) mode. Baroclinic velocities, salinity, potential temperature, and layer interfaces were relaxed to the global HYCOM solution in a buffer zone while the barotropic components were allowed to enter and exit the computational domain according to their wave propagation characteristics. The surface fluxes of momentum and heat were calculated from bulk formula that relied on various atmospheric fields, including 10-m winds, 2-m air temperature, 2-m water vapor mixing ratio, precipitation, and incoming solar radiation. These atmospheric fields were obtained from WRF, which dynamically downscaled the NCEP analysis fields.

REFERENCES

Aksoy, A., F. Zhang, and J. W. Nielsen-Gammon, 2006a: Ensemble-based simultaneous state and parameter estimation in a two-dimensional sea-breeze model. *Mon. Wea. Rev.,* **134,** 2951–2970.

——, ——, and ——, 2006b: Ensemble-based simultaneous state and parameter estimation with MM5. *Geophys. Res. Lett.,* **33,** L12801, doi:10.1029/2006GL026186.

Alexanderian, A., J. Winokur, I. Sraj, A. Srinivasan, M. Iskandarani, W. Thacker, and O. Knio, 2012: Global sensitivity analysis in an ocean general circulation model: A sparse spectral projection approach. *Comput. Geosci.,* **16,** 757–778.

Andreas, E. L, 2011: Fallacies of the enthalpy transfer coefficient over the ocean in high winds. *J. Atmos. Sci.,* **68,** 1435–1445.

——, and K. A. Emanuel, 2001: Effects of sea spray on tropical cyclone intensity. *J. Atmos. Sci.,* **58,** 3741–3751.

Annan, J., J. Hargreaves, N. Edwards, and R. Marsh, 2005: Parameter estimation in an intermediate complexity earth system model using an ensemble Kalman filter. *Ocean Modell.,* **8,** 135–154, doi:10.1016/j.ocemod.2003.12.004.

Bell, M. M., M. T. Montgomery, and K. A. Emanuel, 2012: Air–sea enthalpy and momentum exchange at major hurricane wind speeds observed during CBLAST. *J. Atmos. Sci.,* **69,** 3197–3222.

Berger, J., 1985: *Statistical Decision Theory and Bayesian Analysis*. Springer-Verlag, 634 pp.

Berliner, L. M., R. F. Milliff, and C. K. Wikle, 2003: Bayesian hierarchical modeling of air–sea interaction. *J. Geophys. Res.,* **108,** 3104, doi:10.1029/2002JC001413.

Bernardo, J., and A. F. M. Smith, 2007: *Bayesian Theory*. 2nd ed. Wiley Series in Probability and Statistics, John Wiley & Sons, 640 pp.

Bleck, R., 2002: An oceanic general circulation model framed in hybrid isopycnic cartesian coordinates. *Ocean Modell.,* **4,** 55–88.

Cameron, R. H., and W. T. Martin, 1947: The orthogonal development of non-linear functionals in series of Fourier-Hermite functionals. *Ann. Math.,* **48,** 385–392.

Canuto, C. G., M. Y. Hussaini, A. Quarteroni, and T. A. Zang, 2006: *Spectral Methods: Fundamentals in Single Domains*. Scientific Computation Series, Vol. 22, Springer-Verlag, 581 pp.

Chassignet, E. P., L. Smith, G. Halliwell, and R. Bleck, 2003: North Atlantic simulation with the Hybrid Coordinate Ocean Model (HYCOM): Impact of the vertical coordinate choice, reference density, and themobaricity. *J. Phys. Oceanogr.,* **33,** 2504–2526.

——, H. E. Hurlburt, O. M. Smedstad, G. R. Halliwell, P. J. Hogan, A. J. Wallcraft, R. Baraille, and R. Bleck, 2007: The HYCOM (HYbrid Coordinate Ocean Model) data assimilative system. *J. Mar. Syst.,* **65,** 60–83.

——, and Coauthors, 2009: U.S. GODAE: Global ocean prediction with the HYbrid Coordinate Ocean Model (HYCOM). *Oceanography,* **22** (2), 64–75.

Chen, S. S., J. F. Price, W. Zhao, M. A. Donelan, and E. J. Walsh, 2007: The CBLAST-Hurricane program and the next-generation fully coupled atmosphere–wave–ocean models for hurricane research and prediction. *Bull. Amer. Meteor. Soc.,* **88,** 311–317.

——, W. Zhao, M. A. Donelan, and H. L. Tolman, 2013: Directional wind-wave coupling in fully coupled atmosphere–wave–ocean models: Results from CBLAST-Hurricane. *J. Atmos. Sci.,* in press.

Constantine, P., M. Eldred, and E. Phipps, 2012: Sparse pseudospectral approximation method. *Comput. Methods Appl. Mech. Eng.,* **229–232,** 1–12.

Crestaux, T., O. L. Maitre, and J.-M. Martinez, 2009: Polynomial chaos expansion for sensitivity analysis. *Reliab. Eng. Syst. Saf.,* **94,** 1161–1172.

Davis, C., and Coauthors, 2008: Prediction of landfalling hurricanes with the advanced hurricane WRF model. *Mon. Wea. Rev.,* **136,** 1990–2005.

Donelan, M. A., B. K. Haus, N. Reul, W. J. Plant, M. Stiassnie, H. C. Graber, O. B. Brown, and E. S. Saltzman, 2004: On the limiting aerodynamic roughness of the ocean in very strong winds. *Geophys. Res. Lett.,* **31,** L18306, doi:10.1029/2004GL019460.

Dowd, M., 2011: Estimating parameters for a stochastic dynamic marine ecological problem. *Environmetrics,* **22,** 501–515.

Drennan, W. M., J. A. Zhang, J. R. French, C. McCormick, and P. G. Black, 2007: Turbulent fluxes in the hurricane boundary layer. Part II: Latent heat flux. *J. Atmos. Sci.,* **64,** 1103–1115.

Emanuel, K. A., 1995: Sensitivity of tropical cyclones to surface exchange coefficients and a revised steady-state model incorporating eye dynamics. *J. Atmos. Sci.,* **52,** 3969–3976.

——, 2003: A similarity hypothesis for air–sea exchange at extreme wind speeds. *J. Atmos. Sci.,* **60,** 1420–1428.

Epstein, E. S., 1985: *Statistical Inference and Prediction in Climatology: A Bayesian Approach*. Amer. Meteor. Soc., 199 pp.

Fasham, M. J. R., and G. T. Evans, 1995: The use of optimization techniques to model marine ecosystem dynamics at the JGOFS station at 47°N 20°W. *Philos. Trans. Roy. Soc. London,* **348,** 203–209, doi:10.1098/rstb.1995.0062.

Fennel, K., M. Losch, J. Schöter, and M. Wenzel, 2001: Testing a marine ecosystem model: Sensitivity analysis and parameter optimization. *J. Mar. Syst.,* **28** (1–2), 45–63, doi:10.1016/S0924-7963(00)00083-X.

French, J. R., W. M. Drennan, J. A. Zhang, and P. G. Black, 2007: Turbulent fluxes in the hurricane boundary layer. Part I: Momentum flux. *J. Atmos. Sci.,* **64,** 1089–1102.

Garratt, J. R., 1992: *The Atmosphere Boundary Layer.* Cambridge Atmospheric and Space Science Series, Cambridge University Press, 316 pp.

Gelman, A., J. B. Carlin, H. L. Stern, and D. B. Rubin, 2004: *Bayesian Data Analysis*. 2nd ed. Chapman and Hall/CRC, 668 pp.

Gerstner, T., and M. Griebel, 2003: Dimension–adaptive tensor–product quadrature. *Computing,* **71,** 65–87, doi:10.1007/s00607-003-0015-5.

Haario, H., E. Saksman, and J. Tamminen, 2001: An adaptive metropolis algorithm. *Bernoulli,* **7,** 223–242. [Available online at http://www.jstor.org/stable/3318737.]

Halliwell, G. R., 2004: Evaluation of vertical coordinate and vertical mixing algorithms in the HYbrid-Coordinate Ocean Model (HYCOM). *Ocean Modell.,* **7,** 285–322, doi:10.1016/j.ocemod.2003.10.002.

Haus, B. K., D. Jeong, M. A. Donelan, J. A. Zhang, and I. Savelyev, 2010: Relative rates of sea–air heat transfer and frictional drag in very high winds. *Geophys. Res. Lett.,* **37,** L07802, doi:10.1029/2009GL042206.

Hong, S.-Y., J. Dudhia, and S.-H. Chen, 2004: A revised approach to ice microphysical processes for the bulk parameterization of clouds and precipitation. *Mon. Wea. Rev.,* **132,** 103–120.

——, Y. Noh, and J. Dudhia, 2006: A new vertical diffusion package with an explicit treatment of entrainment processes. *Mon. Wea. Rev.,* **134,** 2318–2341.

Hu, X.-M., F. Zhang, and J. W. Nielsen-Gammon, 2010: Ensemble-based simultaneous state and parameter estimation for treatment of mesoscale model error: A real-data study. *Geophys. Res. Lett.,* **37,** L08802, doi:10.1029/2010GL043017.

Huan, X., and Y. M. Marzouk, 2013: Simulation-based optimal Bayesian experimental design for nonlinear systems. *J. Comput. Phys.,* **232,** 288–317, doi:10.1016/j.jcp.2012.08.013.

Jarosz, E., D. A. Mitchell, D. W. Wang, and W. J. Teague, 2007: Bottom-up determination of air-sea momentum exchange under a major tropical cyclone. *Science,* **315,** 1707–1709, doi:10.1126/science.1136466.

Jeong, D., B. K. Haus, and M. A. Donelan, 2012: Enthalpy transfer across the air–water interface in high winds including spray. *J. Atmos. Sci.,* **69,** 2733–2748.

Jones, E., J. Parslow, and L. Murry, 2010: A Bayesian approach to state and parameter estimation in a phytoplankton-zooplankton model. *Aust. Meteor. J.,* **59,** 7–16.

Kain, J. S., and J. M. Fritsch, 1993: Convective parameterization for mesoscale models: The Kain-Fritsch scheme. *The Representation of Cumulus Convection in Numerical Models, Meteor. Monogr.,* No. 24, Amer. Meteor. Soc., 165–170.

Kara, A. B., P. A. Rochford, and H. E. Hurlburt, 2002: Air–sea flux estimates and the 1997–1998 ENSO event. *Bound.-Layer Meteor.,* **103,** 439–458.

Kondrashov, D., C. Sun, and M. Ghil, 2008: Data assimilation for a coupled ocean–atmosphere model. Part II: Parameter estimation. *Mon. Wea. Rev.,* **136,** 5056–5076.

Kullback, S., 1959: *Information Theory and Statistics*. John Wiley and Sons, 395 pp.

Large, W. G., J. C. McWilliams, and S. C. Doney, 1994: Oceanic vertical mixing: A review and a model with a non-local boundary layer parameterization. *Rev. Geophys.,* **32,** 363–403.

Le Maître, O. P., and O. M. Knio, 2010: *Spectral Methods for Uncertainty Quantification*. Applications to Computational Fluid Dynamics Series, Vol. 16, Springer-Verlag, 536 pp.

Lin, G., and G. E. Karniadakis, 2009: Sensitivity analysis and stochastic simulations of non-equilibrium plasma flow. *Int. J. Numer. Methods Eng.,* **80,** 738–766, doi:10.1002/nme.2582.

Losa, S. N., G. A. Kivman, and V. A. Ryabchenko, 2004: Weak constraint parameter estimation for a simple ocean ecosystem model: What can we learn about the model and data? *J. Mar. Syst.,* **45,** 1–20, doi:10.1016/j.jmarsys.2003.08.005.

Marzouk, Y. M., and H. N. Najm, 2009: Dimensionality reduction and polynomial chaos acceleration of Bayesian inference in inverse problems. *J. Comput. Phys.,* **228,** 1862–1902, doi:10.1016/j.jcp.2008.11.024.

——, ——, and L. A. Rahn, 2007: Stochastic spectral methods for efficient Bayesian solution of inverse problems. *J. Comput. Phys.,* **224,** 560–586, doi:10.1016/j.jcp.2006.10.010.

Mattern, P., K. Fennel, and M. Dowd, 2012: Estimating time-dependent parameters for a biological ocean model using an emulator approach. *J. Mar. Syst.,* **96–97,** 32–47, doi:10.1016/j.jmarsys.2012.01.015.

Olson, R., R. Sriver, M. Goes, N. M. Urban, H. D. Matthews, M. Haran, and K. Keller, 2012: A climate sensitivity estimate using Bayesian fusion of instrumental observations and an earth system model. *J. Geophys. Res.,* **117,** D04103, doi:10.1029/2011JD016620.

Parzen, E., 1962: On estimation of a probability density function and mode. *Ann. Math. Stat.,* **33,** 1065–1076. [Available online at http://www.jstor.org/stable/2237880.]

Petras, K., 2000: On the Smolyak cubature error for analytic functions. *Adv. Comput. Math.,* **12,** 71–93.

——, 2001: Fast calculation of coefficients in the Smolyak algorithm. *Numer. Algorithms,* **26,** 93–109.

——, 2003: Smolyak cubature of given polynomial degree with few nodes for increasing dimension. *Numer. Math.,* **93,** 729–753, doi:10.1007/s002110200401.

Powell, M. D., P. J. Vickery, and T. A. Reinhold, 2003: Reduced drag coefficient for high wind speeds in tropical cyclones. *Nature,* **422,** 279–283, doi:10.1038/nature01481.

Roberts, G. O., and J. S. Rosenthal, 2009: Examples of adaptive MCMC. *J. Comput. Gr. Stat.,* **18,** 349–367, doi:10.1198/jcgs.2009.06134.

Sanford, T. B., J. F. Price, J. B. Girton, and D. C. Webb, 2007: Highly resolved observations and simulations of the ocean response to a hurricane. *Geophys. Res. Lett.,* **34,** L13604, doi:10.1029/2007GL029679.

Shay, L., and S. Jacob, 2006: Relationship between oceanic energy fluxes and surface winds during tropical cyclone passage. *Atmosphere-Ocean Interactions,* Vol. 2, W. Perrie, Ed., *Advances in Fluid Mechanics,* WIT Press, 115–142.

Silverman, B. W., 1986: *Density Estimation for Statistics and Data Analysis*. Chapman and Hall, 176 pp.

Sivia, D. S., 2006: *Data Analysis: A Bayesian Tutorial*. Oxford Science Publications, 264 pp.

Skamarock, W. C., and Coauthors, 2008: A description of the advanced research WRF version 3. NCAR Tech. Note NCAR/TN-475+STR, 113 pp. [Available online at http://www.mmm.ucar.edu/wrf/users/docs/arw_v3_bw.pdf.]

Smedstad, O. M., and J. J. O'Brien, 1991: Variational data assimilation and parameter estimation in an equatorial Pacific Ocean model. *Prog. Oceanogr.,* **26,** 179–241, doi:10.1016/0079-6611(91)90002-4.

Srinivasan, A., and Coauthors, 2010: Many task computing for modeling the fate of oil discharged from the deep water horizon well blowout. Preprints, *2010 IEEE Workshop on Many-Task Computing on Grids and Supercomputers (MTAGS),* New Orleans, LA, IEEE, 1–7. [Available online at http://ieeexplore.ieee.org/xpl/articleDetails.jsp?arnumber=5699424.]

Sudret, B., 2008: Global sensitivity analysis using polynomial chaos expansions. *Reliab. Eng. Syst. Saf.,* **93,** 964–979, doi:10.1016/j.ress.2007.04.002.

Tarantola, A., 2005: *Inverse Problem Theory and Methods for Model Parameter Estimation*. SIAM, 342 pp.

Thacker, W. C., A. Srinivasan, M. Iskandarani, O. M. Knio, and M. L. Henaff, 2012: Propagating boundary uncertainties using polynomial expansions. *Ocean Modell.,* **43–44,** 52–63, doi:10.1016/j.ocemod.2011.11.011.

Velden, C., and Coauthors, 2006: The Dvorak tropical cyclone intensity estimation technique. *Bull. Amer. Meteor. Soc.,* **87,** 1195–1210.

Villegas, M., F. Augustin, A. Gilg, A. Hmaidi, and U. Wever, 2012: Application of the polynomial chaos expansion to the simulation of chemical reactors with uncertainties. *Math. Comput. Simul.,* **82,** 805–817, doi:10.1016/j.matcom.2011.12.001.

Vogel, C., 2002: *Computational Methods for Inverse Problems*. SIAM, 183 pp.

Walsh, K. J., P. Sandery, G. B. Brassington, M. Entel, C. Siegenthaler-Ledrian, J. D. Kepert, and R. Darbyshire, 2010: Constraints on drag and exchange coefficients at extreme wind speeds. *J. Geophys. Res.,* **115,** C09007, doi:10.1029/2009JC005876.

Winokur, J., P. P. Conrad, I. Sraj, O. M. Knio, A. Srinivasan, W. C. Thacker, Y. Marzouk, and M. Iskandarani, 2013: A priori testing of sparse adaptive polynomial chaos expansions using an ocean general circulation model database. *Comput. Geosci.,* doi:10.1007/s10596-013-9361-3, in press.

Xiu, D., and D. Tartakovsky, 2004: Uncertainty quantification for flow in highly heterogeneous porous media. *Computational Methods in Water Resources,* C. T. Miller et al., Eds., Developments in Water Science Series, Vol. 55, Elsevier, 695–703.

Zedler, S., P. Niiler, D. Stammer, E. Terrill, and J. Morzel, 2009: Ocean's response to Hurricane Frances and its implications for drag coefficient parameterization at high wind speeds. *J. Geophys. Res.,* **114,** C04016, doi:10.1029/2008JC005205.

——, G. Kanschat, R. Korty, and I. Hoteit, 2012: A new approach for the determination of the drag coefficient from the upper ocean response to a tropical cyclone: A feasibility study. *J. Oceanogr.,* **68,** 227–241.