

Implications of Multivariate Non-Gaussian Data Assimilation for Multiscale Weather Prediction

JONATHAN POTERJOY^{a,b}

^a *Department of Atmospheric and Oceanic Science, University of Maryland, College Park, College Park, Maryland*

^b *NOAA/Atlantic Oceanographic and Meteorological Laboratory, Miami, Florida*

(Manuscript received 31 August 2021, in final form 28 February 2022)

ABSTRACT: Weather prediction models currently operate within a probabilistic framework for generating forecasts conditioned on recent measurements of Earth's atmosphere. This framework can be conceptualized as one that approximates parts of a Bayesian posterior density estimated under assumptions of Gaussian errors. Gaussian error approximations are appropriate for synoptic-scale atmospheric flow, which experiences quasi-linear error evolution over time scales depicted by measurements, but are often hypothesized to be inappropriate for highly nonlinear, sparsely observed mesoscale processes. The current study adopts an experimental regional modeling system to examine the impact of Gaussian prior error approximations, which are adopted by ensemble Kalman filters (EnKFs) to generate probabilistic predictions. The analysis is aided by results obtained using recently introduced particle filter (PF) methodology that relies on an implicit nonparametric representation of prior probability densities—but with added computational expense. The investigation focuses on EnKF and PF comparisons over monthlong experiments performed using an extensive domain, which features the development and passage of numerous extratropical and tropical cyclones. The experiments reveal spurious small-scale corrections in EnKF members, which come about from inappropriate Gaussian approximations for priors dominated by alignment uncertainty in mesoscale weather systems. Similar behavior is found in PF members, owing to the use of a localization operator, but to a much lesser extent. This result is reproduced and studied using a low-dimensional model, which permits the use of large sample estimates of the Bayesian posterior distribution. Findings from this study motivate the use of data assimilation techniques that provide a more appropriate specification of multivariate non-Gaussian prior densities or a multiscale treatment of alignment errors during data assimilation.

SIGNIFICANCE STATEMENT: Numerical predictions of Earth's atmosphere require computer models, which represent known physical processes governing the evolution of atmospheric flow, and a clever use of statistical methods to construct a complete model representation of the true atmosphere from incomplete measurements. The second requirement is built on assumptions for the shape of error distributions for variables that are input into the model for generating predictions. The present study explores the fidelity of these error assumptions for regional weather forecasting using a novel technique that avoids common approximations that go into operational weather prediction systems.

KEYWORDS: Bayesian methods; Kalman filters; Data assimilation; Ensembles

1. Introduction

Numerical predictions of Earth's atmosphere are imperfect, owing to computational and scientific limitations and the observability of the dynamical system at hand. Even as these limitations diminish, atmospheric flow exhibits intrinsic predictability limits that place a firm barrier for what can be achieved with deterministic predictions (e.g., Lorenz 1969; Rotunno and Snyder 2008). Therefore, a probabilistic framework is needed. Within this framework, various data assimilation techniques now exist for solving user-specified parts of a Bayesian posterior density, which informs how predictions account for environmental measurements of the true dynamical system (Reich and Cotter 2015).

For global weather prediction, there exist additional challenges in how data assimilation operates for multiple scales of motion. Midlatitude synoptic-scale weather systems are dominated by horizontal motions that are in approximate geostrophic balance. Weather phenomena at these scales tend to

be characterized well by satellite and in situ observing systems, which leads to small forecast errors over typical 6–12-h windows used by operational global data assimilation methods. Therefore, Gaussian approximations used by ensemble Kalman filters (EnKFs; Evensen 1994; Bishop et al. 2000; Anderson 2001; Whitaker and Hamill 2002) and three- and four-dimensional variational methods (3DVar and 4DVar; Thepaut and Courtie 1991; Courtier et al. 1994) provide reasonably accurate results. For this reason, all national weather forecast centers currently use various forms of EnKFs, 3DVar, 4DVar, or hybrids of these methods for operational prediction (Bannister 2017). Among the above methods, techniques that combine four-dimensional data assimilation with ensemble statistics have become standard practice. This choice likely follows from practical benefits of smoothers; i.e., assimilating observations over a window can be less costly than assimilating observations at multiple times, and induce a smaller amount of sampling error (Kurosawa and Poterjoy 2021). Smoothers also apply a Gaussian assumption on the smoothing density—rather than the filtering density—which is likely more appropriate for weakly nonlinear data

Corresponding author: Jonathan Poterjoy, poterjoy@umd.edu

DOI: 10.1175/MWR-D-21-0228.1

© 2022 American Meteorological Society. For information regarding reuse of this content and general copyright information, consult the [AMS Copyright Policy](#) ([www.ametsoc.org/PUBSReuseLicenses](#)).

assimilation problems (Morzfeld and Hodyss 2019). For applications such as global weather prediction, computational limitations typically restrict smoothers to a high-resolution deterministic analysis, leading to Gaussian filters or smoothers being adopted for adjusting a reduced-resolution ensemble between observation times (e.g., Zhang et al. 2009; Buehner et al. 2010a,b; Poterjoy and Zhang 2014).

Sub-synoptic-scale flow evolves on a faster time scale and exhibit a higher dependence on parameterized physical processes in models, such as radiation, unresolved turbulent motions, and cloud microphysics (Stensrud 2007). Physical parameterizations provide energy sources and sinks that are not represented explicitly in models, thus adding to the uncertainty in short-term predictions. For resolved processes, satellite and radar winds can provide important dynamical information regarding the time evolution of mesoscale features, but thermodynamic variables are often poorly observed. This information can be inferred from satellites and radars; however, the mapping between observations and prognostic model variables is nonlinear. Therefore, short-range forecasts for mesoscale weather features are associated with elevated levels of uncertainty, even if they are embedded in well-constrained synoptic-scale systems. "Alignment errors" (Hoffman et al. 1995) at these scales pose a major source of uncertainty and are best represented by multivariate non-Gaussian additive errors for model variables (Lawson and Hansen 2005; Hodyss and Reinecke 2013).

For multiscale weather prediction, which is characterized by well-resolved synoptic-scale flow and large mesoscale uncertainty, the potential benefits of data assimilation algorithms that move beyond long-standing Gaussian assumptions remain unknown. A part of this unknown stems from limitations in the computational cost available for running modern weather models. Current operational data assimilation systems rely on ensemble forecasts of size $O(10^2)$ to estimate prior statistics for data assimilation, which is likely too small for quantifying higher-order statistics beyond the first two moments, even for marginal quantities (Miyoshi et al. 2014). Because of this shortcoming, a natural starting point for solving challenges posed by non-Gaussian errors is to focus more narrowly on observation-space prior statistics and observation likelihoods that are known to be non-Gaussian. In this context, nonlinearity in measurement operators may provide clear parametric choices for non-Gaussian observations and observation-space priors (e.g., Fletcher and Zupanski 2006; Bishop 2016). For prior statistics that exhibit irregular deviations from Gaussianity, some filters leverage nonparametric representations of marginal densities (e.g., Anderson 2010, 2020), which act as an intermediate step toward a fully non-Gaussian methodology. Furthermore, Gaussian filters can be modified to include higher-order multivariate statistics, such as skewness (e.g., Hodyss 2012), or treat alignment errors that yield non-Gaussian prior uncertainty (e.g., Ravela et al. 2007; Nehr Korn et al. 2015; Stratman et al. 2018; Ying 2019).

Not until recently have more general multivariate non-Gaussian data assimilation methods based on particle filters (Doucet et al. 2001) been applied for real weather app-

lications (Robert et al. 2018; Potthast et al. 2019; Poterjoy et al. 2019, 2021). Particle filters (PFs) do not make strict assumptions for the shape of prior and posterior error distributions, but require ensemble sizes that are unobtainable in modern weather prediction systems (Snyder et al. 2015; Slivinski and Snyder 2016). Therefore, assumptions are still needed to cope with the large dimension of the data assimilation problem at hand. Borrowing from EnKFs, spatial localization presents a common strategy that has yielded numerous types of high-dimensional PF methods; examples include Bengtsson et al. (2003), Poterjoy (2016), Poterjoy and Anderson (2016), Penny and Miyoshi (2016), Lee and Majda (2016), and Morzfeld et al. (2018). We refer readers to Farchi and Bocquet (2018) and Leeuwen et al. (2019) for reviews on localized PFs as well as alternative strategies for adapting PFs for high-dimensional problems.

The current study uses an experimental regional modeling system introduced by Poterjoy et al. (2021, hereafter PAW21), to reveal deficiencies in Gaussian-based data assimilation methodology that may pose long-term challenges for regional and global weather prediction models. For this purpose, a recently proposed localized PF (Poterjoy et al. 2019; Poterjoy 2022, hereafter P22) is compared with an EnKF during monthlong data assimilation experiments that are dominated by the passage of extratropical and tropical weather systems. These experiments use the Hurricane Weather Research and Forecasting (HWRF) Model (Gopalakrishnan et al. 2010; Atlas et al. 2015), with a domain that covers extensive portions of the Atlantic Ocean, Caribbean, and southeastern United States. The experiments target a period that features several major tropical cyclones from the 2017 eastern North Atlantic hurricane season. While forecast results from these experiments are briefly discussed in PAW21, the current study provides a more detailed examination of the EnKF and local PF members produced over the experiment period. The real-data experiments reveal major differences between the EnKF and local PF, which are hypothesized to come from Gaussian assumptions in the EnKF. This hypothesis is examined using a simplified two-dimensional (2D) data assimilation application, which replicates some of the behavior observed in HWRF experiments. The 2D problem, which consists of an axisymmetric vortex in zero mean flow, exhibits qualitatively similar behavior to displaced weather features depicted by the HWRF ensemble at scales less than 150 km, but with the added advantage of having full control over the source of prior uncertainty.

The manuscript is organized in the following manner. Section 2 introduces the two data assimilation methods used for this study. Section 3 briefly describes the regional modeling system and period of interest for numerical experiments. Section 4 summarizes findings from monthlong comparisons of the Gaussian and non-Gaussian data assimilation methods using the HWRF Model; details of these findings are further explored using a low-dimensional vortex application in section 5. The last section discusses major findings from this study and implications for atmospheric prediction.

2. Data assimilation methodology

The data assimilation methods explored in this study includes the square root EnKF of Whitaker and Hamill (2002) and an iterative version of the Poterjoy et al. (2019) local PF. The iterative PF is a simplified version of the adaptive tempering method discussed in P22. It is briefly described in this section alongside the EnKF.

a. The ensemble Kalman filter

Ensemble Kalman filters apply a Monte Carlo approximation of the model state uncertainty, assuming Gaussian errors (Evensen 1994). This study uses the EnKF within the community Grid point Statistical Interpolation (GSI) data assimilation package. GSI is used by operational NOAA weather prediction models including the operational HWRf Model used for the current study (Biswas et al. 2018).

The GSI EnKF operates by assimilating a vector of observations \mathbf{y} serially to update an ensemble of model state vectors $\{\mathbf{x}_n\}_{n=1}^{N_e}$. The ensemble mean follows the same update as the Kalman filter, while perturbations are transformed to match the posterior error covariance estimated by the Kalman filter equations (e.g., Jazwinski 1970). Instead of using tangent linear approximations of measurement operators, the full operator is applied to each member to yield a set of observation-space prior members $\{\mathbf{h}\mathbf{x}_n\}_{n=1}^{N_e}$, which are updated alongside the original state vectors. Using superscripts $-$ and $+$ to denote quantities before and after assimilating the i th observation in \mathbf{y} , respectively, the mean after assimilating observations up to y_i is given by

$$\bar{\mathbf{x}}^+ = \bar{\mathbf{x}}^- + \mathbf{K}(y_i - \overline{hx_i^-}), \tag{1}$$

$$\mathbf{K} = \frac{\hat{\mathbf{P}}^-}{\hat{\sigma}_i^- + r_i}, \tag{2}$$

$$\hat{\mathbf{P}}^- = \frac{1}{N_e - 1} \sum_{n=1}^{N_e} \left[(\mathbf{x}_n^- - \bar{\mathbf{x}}^-)(hx_i^- - \overline{hx_i^-}) \right], \tag{3}$$

$$\hat{\sigma}_i^- = \frac{1}{N_e - 1} \sum_{n=1}^{N_e} \left(hx_i^- - \overline{hx_i^-} \right)^2, \tag{4}$$

where $\hat{\mathbf{P}}^-$ and $\hat{\sigma}_i^-$ indicate ensemble-estimated error covariance and variance, overbars indicate ensemble averages, \mathbf{K} is the Kalman gain matrix, and r_i indicates the error variance for the i th observation. Perturbations about the mean are updated separately using

$$\mathbf{x}'_n{}^+ = \mathbf{x}'_n{}^- - \alpha \mathbf{K} \mathbf{h} \mathbf{x}'_n{}^-, \tag{5}$$

$$\alpha = \left(1 + \sqrt{\frac{\hat{\sigma}_i^-}{\hat{\sigma}_i^- + r_i}} \right)^{-1}. \tag{6}$$

The observation space ensemble, $\{\mathbf{h}\mathbf{x}_n\}_{n=1}^{N_e}$, is updated in parallel with the state space ensemble to avoid repeated

calculations of measurement operators. Model states updated by (1)–(6) are then used as input for the next observation in the sequence and the steps are repeated until all observations are assimilated. We also note that the above equations can be modified to assimilate batches of observations with correlated errors, which is not done for the current study.

The mean, variance, and covariance calculations in (1)–(6) come from sample estimates over the ensemble and are prone to sampling error. Therefore, heuristic modifications of these quantities are needed when N_e is small and when statistics are biased because of errors not accounted for in the ensemble; e.g., unknown sources of model or observation error. The current study adopts the strategies commonly used within the GSI framework, namely, covariance localization via an element-wise product of \mathbf{K} and Eq. (4.10) of Gaspari and Cohn (1999) and relaxation to prior spread (RTPS) posterior inflation (Whitaker and Hamill 2012).

b. The local particle filter

A serial implementation of the local PF was first introduced by Poterjoy (2016) and modified in Poterjoy et al. (2019) to more closely resemble the EnKF described above. This section briefly introduces the local PF using notation that is similar to the EnKF update described by (1)–(6).

As in the EnKF, the local PF first applies the full measurement operator on each ensemble member—also called “particles”—to obtain $\{\mathbf{h}\mathbf{x}_n\}_{n=1}^{N_e}$. Also like the EnKF, $\{\mathbf{h}\mathbf{x}_n\}_{n=1}^{N_e}$ is updated alongside the state-space ensemble $\{\mathbf{x}_n\}_{n=1}^{N_e}$. For the i th observation in \mathbf{y} the state mean is given by

$$\bar{\mathbf{x}}^+ = \sum_{n=1}^{N_e} \omega_{i,n} \circ \mathbf{x}_n^-, \tag{7}$$

where $\omega_{i,n}$ is a vector containing weighting coefficients for the n th particle given all observations up to y_i and \circ indicates a Schur product. These weights determine marginal posterior quantities from a weighted sum of prior particles; see P22 for a full description of the multivariate posterior characterized by the local PF, which describes the origin of these weights. In general, each $\omega_{i,n}$ is a function of normalized likelihoods for marginal variables in each particle. Elements of $\omega_{i,n}$ take the same form as the standard bootstrap PF weights (Gordon et al. 1993), except they taper to $1/N_e$ at large distances from each y_i . The function adopted for the tapering is the same used for joint observation-model space localization in the EnKF. Refinements introduced in Poterjoy et al. (2019) ensure that localization impacts posterior mean estimates in a manner that is identical to the EnKF localization.

Because of localization, perturbations about $\bar{\mathbf{x}}^+$ for each variable become linear combinations of prior particles. To maintain the Bayesian convergence property of PFs in the vicinity of measurements, the local PF first draws particles with replacement from the univariate posterior for hx_i , which is represented by a weighted sum of delta functions (Gordon et al. 1993). The sampling step results in a set of indices, $\{k_n\}_{n=1}^{N_e}$, which inform how to resample for a given measurement.

The sampling indices are then used in the perturbation update equation:

$$\mathbf{x}'_n{}^+ = \mathbf{r}_{1,i} \circ (\mathbf{x}_{k_n}^- - \bar{\mathbf{x}}^+) + \mathbf{r}_{2,i} \circ (\mathbf{x}_n^- - \bar{\mathbf{x}}^+), \quad (8)$$

where $\mathbf{r}_{1,i}$ and $\mathbf{r}_{2,i}$ ensure that the update satisfies the first two moments dictated by importance weights everywhere in state space (Poterjoy 2016). Note that because of the sampling step, this equation still maintains the original bootstrap PF solution in the vicinity of measurements.

To maintain particle diversity, $\mathbf{r}_{1,i}$ and $\mathbf{r}_{2,i}$ can be modified to increase the mixing of particles that are duplicated during resampling (Poterjoy et al. 2019; Feng et al. 2020). The amount of mixing is controlled by the parameter γ discussed in P22 and used for all experiments performed in this study. This strategy helps stabilize the filter when N_e is small and ensures sampled particles yield unique solutions in the absence of localization. The local PF also uses the weight regularization introduced in P22, which heuristically decreases the impact of observations on particle weights if they fall below a target effective ensemble size. Without regularization, the $\omega_{i,n}$ vectors can easily collapse to zero for all but a single particle when presented with very accurate or dense measurements. Last, P22 explores the impact of fitting the first two moments alone for unobserved variables in (8), and introduces strategies for sampling from non-Gaussian posteriors despite this assumption. The current study adopts the likelihood factorization or “tempering” method discussed in section 4 of P22, which samples from the local PF posterior density via a set of intermittent transitions. For the current study, we factor the likelihood into three parts, which is sufficient for achieving the desired improvements in the current applications. This process involves assimilating measurements three times after inflating observation error variance by a factor of 3 (Emerick and Reynolds 2012). We demonstrate the effectiveness of the iterative approach on a two-dimensional geophysical problem in section 5.

The local PF shares a common algorithm structure with the serial EnKF, thus allowing for the new filter to be included in the GSI software package for direct comparisons with current data assimilation methods used for operational weather prediction and research. As discussed in P22, the regularization and tempering strategies adopted for this study provide significant benefits over previous implementations of the local PF. These steps, however, bring additional computational cost, which amounts to a PF update step that is a factor of 10 greater than the EnKF for the real application presented in sections 3 and 4.¹ Improving the computational efficiency of the iterative local PF is an ongoing topic of research and will be discussed in a future study.

¹ This estimate comes from assimilating $O(10^6)$ observations for the domain discussed in section 3, using 240 cores on the NOAA Jet high-performance computing system.

3. The AOML-UMD ensemble prediction system

The current study adopts an experimental regional HWRf modeling system called the “AOML-UMD ensemble system,” which is introduced in PAW21. This modeling system is built within the same software package used by the operational HWRf Model but with added workflow options and scripts for performing uninterrupted sequential data assimilation over long periods. Unlike the operational HWRf, it is not periodically reset from the operational NOAA Global Forecast System (GFS) Data Assimilation System (GDAS) and instead relies on the GFS for boundary conditions alone. The data assimilation step yields updates to all prognostic model variables, namely, horizontal winds, hydrostatic pressure, total pressure, specific humidity, vertical acceleration and “total condensate” from the Ferrier–Aligo microphysics scheme used by HWRf (Aligo et al. 2014). We note that this strategy differs from the operational HWRf, which initiates the model with zero vertical acceleration and total condensate each forecast time. To facilitate a simplified analysis of data assimilation strategies, the AOML-UMD ensemble system uses a single static grid with no moving nests. The suite of assimilated observations include all operationally available conventional and clear-air radiance measurements—the only exception being aircraft reconnaissance measurements routinely collected from storms targeted during hurricane seasons. We exclude these measurements for the present study, owing to the coarse model resolution used in numerical experiments. As discussed in PAW21, radiance measurements are assimilated using time-dependent bias correction coefficients that were estimated using large-domain experiments.

For this study, the HWRf domain has 248×298 horizontal grid points, with an approximate grid spacing of 18 km. Figure 1 illustrates the domain coverage and the location of conventional (non-radiance) measurements assimilated for a single cycle. As noted in PAW21, we adopt the same physical parameterizations used for the 2017 operational HWRf, but choose the same model level configuration as the 2016 version. This choice results in 61 vertical levels that extend up to 2 hPa, which is higher than the 2017 HWRf and presents theoretical benefits for radiance measurement operator calculations.

Data assimilation experiments performed within the AOML-UMD ensemble system use 60 members, which are updated every 6 h. After weeklong tuning experiments examining EnKF and local PF sensitivity to data assimilation parameters, we arrived at coefficients for localization cutoff length scales and inflation that are appropriate for each method; these parameters are summarized in Table 1. In general, the optimal cutoff lengths for the EnKF are found to be longer than those identified for the local PF. Similar to past experiments (e.g., Poterjoy et al. 2019), this discrepancy reflects the higher sensitivity of the local PF to sampling error, which is one drawback of applying a delta function approximation for prior distributions.

We also note that experiments performed in this study do not use supplementary strategies for maintaining balance following data assimilation steps, such as incremental analysis

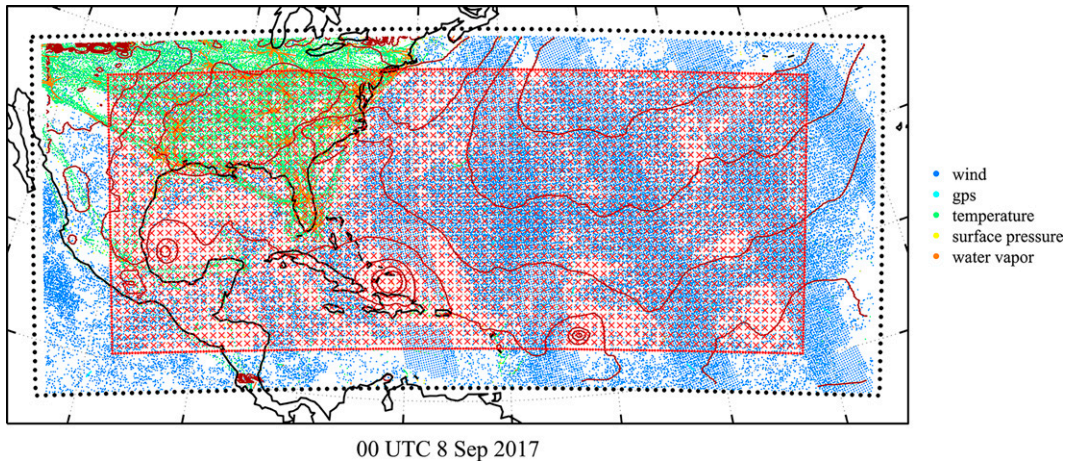


FIG. 1. The dotted line indicates HWRF domain boundaries used for all modeling experiments, which are plotted alongside (solid black lines) land boundaries. For reference, red contours of MSLP are plotted every 5 hPa for a single PF member along with locations of conventional measurements at a single time (0000 UTC 8 Sep 2017). The red hatched region indicates the verification domain.

update (Bloom et al. 1996) or digital filter initialization (Lynch and Huang 1992). While these methods are known to improve forecast skill, they bring additional complexity for interpreting results from the perspective of data assimilation research. Given the scope of this study, the adopted methodology emphasizes transparency over best operational practices.

4. Results from monthlong experiments

In this section, we briefly summarize the findings reported in PAW21 regarding ensemble forecast skill obtained from EnKF and local PF experiments. We then provide a close examination of how the two filters generate posterior solutions during the experiments and draw conclusions based on the underlying assumptions behind each data assimilation system.

As discussed previously, the AOML-UMD ensemble system provides the framework for analyzing data assimilation systems for multiscale weather applications. Our experiments encompass a monthlong period spanning all of September 2017, and cover a domain that featured several major hurricanes. The formation, intensification, and interaction of these storms with extratropical waves introduce numerous forecast challenges—and presents an ideal test problem for non-Gaussian data assimilation methodology. To formulate experiments, each ensemble starts from GDAS ensemble members on 0000 UTC 1 September, and uses boundary conditions from the GDAS ensemble thereafter. For verification purposes, we ignore the first 5 days of sequential data

assimilation, which reduces any memory of the initial GDAS analysis. We arrived at this configuration by examining domain-average root mean squared differences (RMSDs) between the prior EnKF mean and GDAS deterministic analysis over the full monthlong period. While most variables quickly approach quasi-steady RMSDs, specific humidity (q_u) requires about 4 days to adjust (Fig. 2), thus motivating the 5-day spinup period.

a. Forecast skill

Following spinup, we quantify the performance of the EnKF and local PF by running 120-h ensemble forecasts from the first 20 members generated by each data assimilation method. These forecasts run twice daily (every 12 h), yielding a set of 52 ensemble forecasts to verify. As in PAW21, we quantify errors in ensemble-mean forecasts using volume-averaged RMSDs to the GDAS analysis. The selected verification volume is 500 km from lateral domain boundaries and stretches from the surface to 16 km using horizontal and vertical grid spacings of 54 km and 500 m, respectively (hatched region in Fig. 1). The verification domain is selected to remove most large correlations between verifying grid points, and is far enough from the domain edges to not be heavily impacted by the GDAS boundary conditions. Errors calculated over this domain tend to be dominated by synoptic-scale flow, but are also sensitive to storm track, structure, and intensity when major hurricanes exist in forecasts. While forecasts should ideally be verified using raw measurements, the lack of upper-air in situ measurements over oceans presents challenges. Note that mature global data assimilation systems,

TABLE 1. Choices of parameters for HWRF experiments.

Method	Horizontal localization cutoff	Vertical localization cutoff	Relaxation coefficient	N_{eff}	γ
EnKF	0.75 scale height	1.5 scale height	0.95	—	—
Local PF	0.50 scale height	1.0 scale height	—	24	0.3




FIG. 2. Root-mean-square fit of EnKF mean q_v to GDAS q_v over the monthlong HWRf experiments.

like GDAS, rely heavily on radiances measurements to cope with this shortcoming. The GDAS analysis is generated by the Environmental Modeling Center (EMC) of the NOAA National Centers for Environmental Prediction (NCEP) using a four-dimensional ensemble-variational analysis over 6-h windows (Kleist and Ide 2015). Compared to the AOML-UMD ensemble system, GDAS uses a higher model resolution and more observations. GDAS also uses quality control and data assimilation procedures that are more thoroughly tested than HWRf—owing to its long history as a global data assimilation system. It is also not expected to be biased in favor of the EnKF or local PF. Therefore, this product provides an appropriate dataset for verifying forecast results in the current study, especially those at later lead times, when forecast error is much larger than error in GDAS analyses. Furthermore, we omit a detailed comparison of ensemble spread for these experiments as it is difficult to verify this quantity without error estimates of the GDAS analysis—but discuss the time-evolution of variance in choice variables in the latter part of this section.

Figure 3 (reproduced from Fig. 11 of PAW21) shows volume- and forecast-average RMSDs calculated for zonal wind (u), meridional wind (v), temperature (T), and q_v . Despite providing a closer fit to the GDAS analysis, the

verification shows that EnKF ensemble forecasts tend to lose skill at a faster rate than local PF forecasts. In general, the local PF produces more accurate mean forecasts than the EnKF at lead times greater than 24 h for all variables. Poterjoy et al. (2017) noted similar forecast behavior when comparing the local PF to an EnKF in idealized simulations of a mesoscale convective system. They found the EnKF to produce smaller posterior mean RMSEs than the local PF, but degraded forecast accuracy, owing to physically inconsistent multivariate adjustments to thermodynamic and hydro-meteor variables within individual ensemble members. To supplement this verification, PAW21 also examine the frequency at which the local PF provides lower forecast RMSDs than the EnKF, and find that results are not dominated by a subset of cycles; i.e., the local PF produces more accurate forecasts than the EnKF for a majority of the forecasts.

To identify precursors that explain the observed difference in error growth between experiments (i.e., Fig. 3), we further analyze results using potential vorticity (PV). PV is adopted for this purpose, as it allows for a compact quantification of errors in wind, temperature, and geopotential for the balanced part of synoptic and mesoscale flow over the domain (e.g., Dirren et al. 2003). PV is also sensitive to sharp gradients in wind and potential temperature, which are expected to be large in the presence of phase errors at the mesoscale, even at early forecast lead times. From visual inspection of GDAS analyses, PV also tends to evolve smoothly with time on isentropic surfaces, suggesting that adjustments induced by data assimilation are rather small between data assimilation cycles, compared to the raw variables themselves. Baroclinic waves that dip southward into the verification domain exhibit clear signatures in upper-tropospheric PV. Likewise, tropical cyclones produce coherent structures of elevated PV in the lower to middle troposphere, with corresponding negative PV above. The interaction of tropical cyclones with upper-tropospheric PV anomalies also presents an insightful

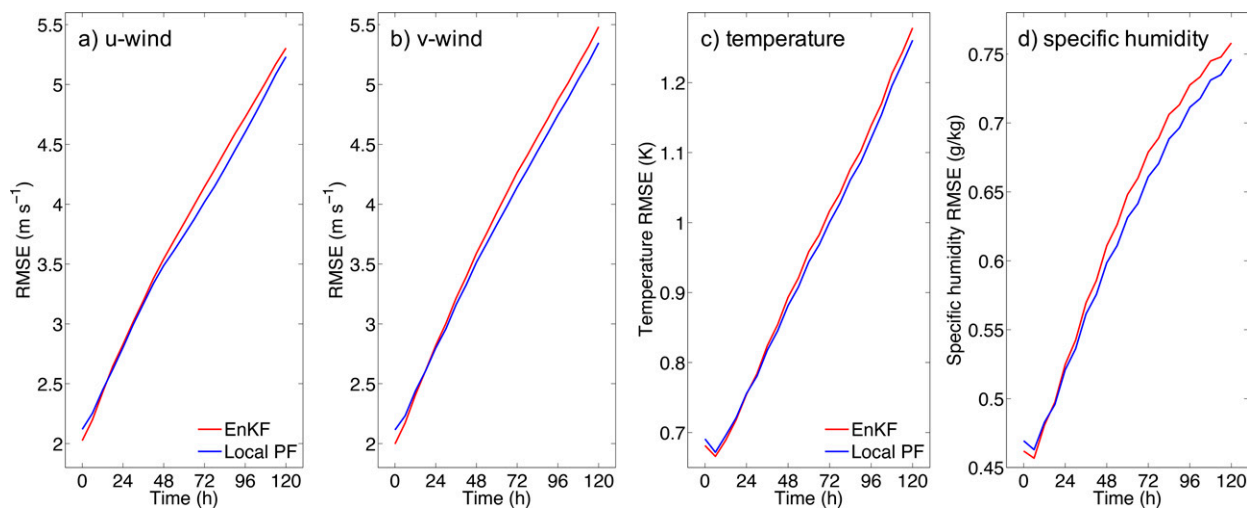


FIG. 3. Ensemble-mean forecast RMSDs as a function of forecast lead time for (a) u , (b) v , (c) T , and (d) q_v for the EnKF (red lines) and local PF (blue lines), averaged over all forecasts generated during HWRf experiments.

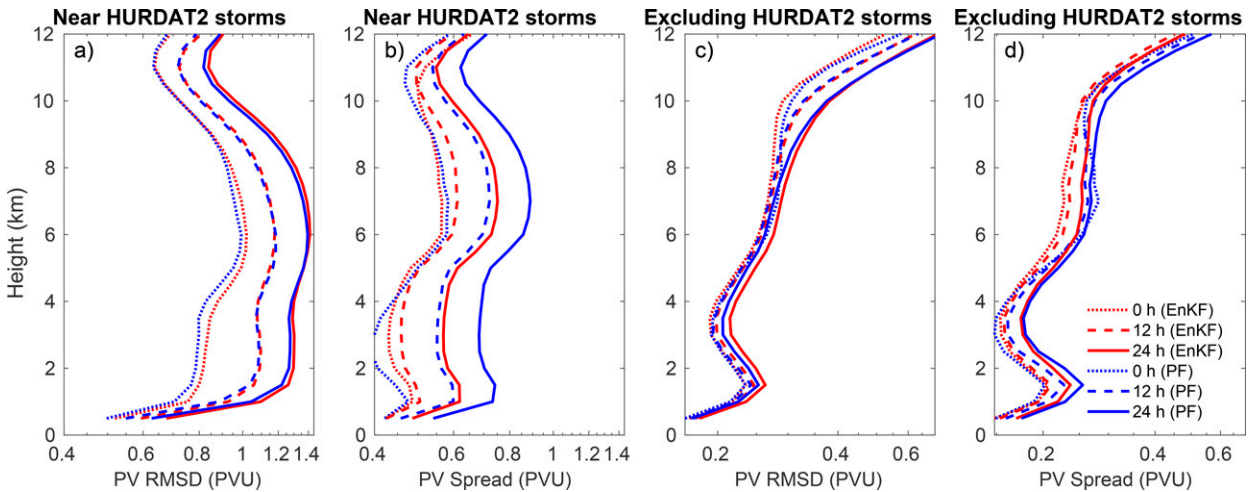


FIG. 4. Domain-average ensemble-mean forecast RMSDs and spread for PV at 0-, 12-, and 24-h forecast lead times. (a) RMSDs and (b) spread calculated from points within 500 km of storms indicated in the HURDAT2 database. (c),(d) As in (a) and (b), but for the remaining grid points.

conceptual framework for analyzing the evolution of storms embedded in complex environments (Molinari et al. 1998). While the present study does not provide a detailed analysis of PV over the course of the experiment, the aggregate effects of PV errors are quantified during the first 48 h of ensemble forecasts. For this verification, PV is first calculated on isentropic surfaces before calculating RMSDs and spread and interpolating results to height surfaces for visualization. We further stratify the PV verification to isolate error contributions from tropical cyclones. This step involves separating verification points that land within 500 km of tropical and posttropical cyclone locations, as depicted in the Hurricane Database (HURDAT2; Landsea and Franklin 2013), from those beyond 500 km of storm locations. Figure 4 shows profiles of domain-wide ensemble-mean forecast errors and spread for PV at 0, 12, and 24 h, again using GDAS for verification. From this analysis, local PF forecasts tend to produce

comparable or slightly improved ensemble-mean forecasts for tropospheric PV by 12 h—with clear benefits by 24 h. This finding occurs for grid points in the vicinity of storms as well as those outside this region, thus suggesting similar benefits in both regimes. Though not shown, further inspection of error profiles in raw variables show a similar trend in tropospheric temperature, which suggests that some of the benefits in local PF forecasts may stem from a more accurate depiction of static stability in lower levels of the verification domain. Furthermore, EnKF ensemble forecasts tend to exhibit a smaller amount of 0–12 h growth in spread through the depth of the troposphere, suggesting that these forecasts sample a smaller subspace of possible model solutions. This behavior is more clearly seen in the verification performed near storms (Fig. 4b), which we will analyze for the remaining portions of this manuscript. All else being equal, this property of the EnKF alone would lead to forecast degradations for the ensemble mean at later lead times.

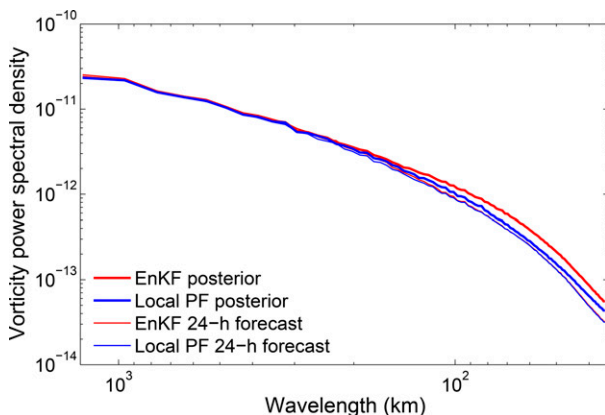


FIG. 5. Power spectral density for ζ averaged from 24-h forecast members (fine lines) and posterior members (heavy lines) for the EnKF (red lines) and local PF (blue lines).

b. Multivariate posterior updates

Findings from Poterjoy et al. (2017) suggest that PF-based data assimilation methodology provides more appropriate multivariate updates than EnKFs—which is one property of geophysical applications that is challenging to verify, owing to the large state dimension. While numerous variables were analyzed for this study, we will focus more narrowly on the horizontal wind field produced by EnKF and local PF members, as the behavior of the two methods can be easily replicated using low-dimensional models. From qualitative inspections of posterior members, major differences exist in the rotational part of horizontal wind increments, which can be easily quantified via spectral analysis.

To perform an objective comparison of how each method represents posterior winds, we first calculate two-dimensional power spectral densities for the vertical vorticity (ζ) fields in EnKF and local PF members, and sum the magnitudes within

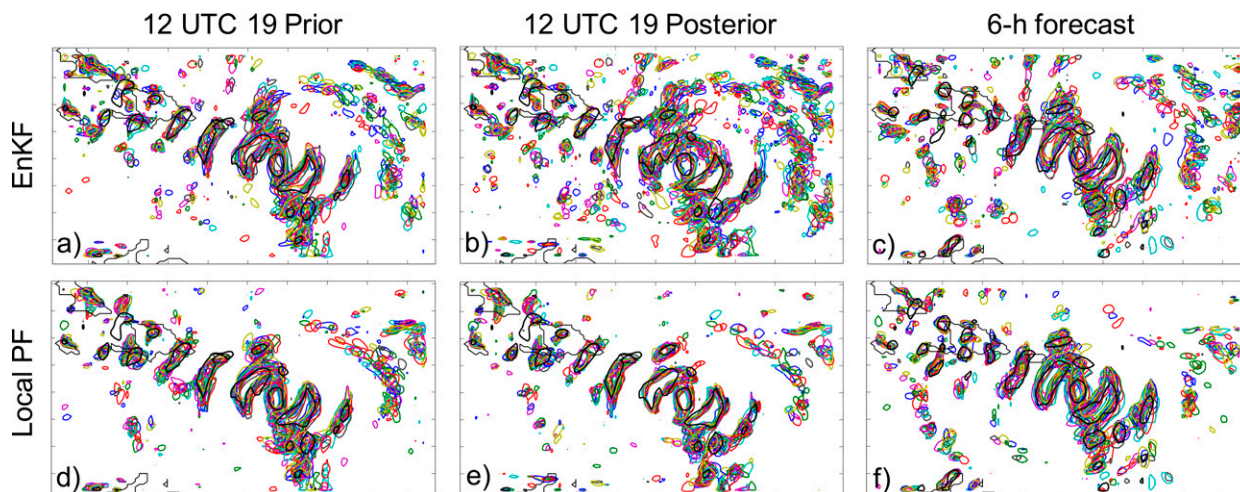


FIG. 6. Vertical vorticity at 850 hPa filtered to remove wavelengths > 150 km. (top) EnKF and (bottom) local PF members for the (left) prior, (center) posterior, and (right) 6-h forecast from the 1200 UTC 19 Sep cycle. For reference, black contours show the GDAS analysis at (left),(center) 1200 UTC 19 Sep and (right) 1800 UTC 19 Sep.

bins. We then average each power spectra over the same 20 members used to generate forecasts. The posterior power spectra (heavy lines in Fig. 5) are then compared with average power spectra calculated from 24-h EnKF and local PF forecast members produced over the same period (thin lines Fig. 5). From this analysis, it is clear that both the local PF and EnKF produce similar wind spectra for 24-h forecasts, which is a lead time when transient adjustments following forecast initialization are negligible. The two power spectra are indistinguishable from each other in Fig. 5, owing to the time-averaged mean spectra being mostly a reflection of the HWRF Model climatology for winds over the monthlong period. Nevertheless, notable differences emerge in the posterior ζ power density spectra, which will be discussed for the remainder of this section.

The spectral analysis suggests that both data assimilation methods, to varying degree, systematically modify the distribution of kinetic energy at scales smaller than 150 km. The spurious increase in energy introduced by the EnKF, however, is only partially matched by the local PF. To explore the implications of this result, we examine the ζ field for a single data assimilation step and pressure level. For illustrative purposes, we choose a single storm (Hurricane Maria at 1200 UTC 19 September) and interpolate ζ to the 850-hPa pressure level after applying a high-pass filter to remove wavelengths greater than 150 km, which is the minimum length scale at which EnKF and local PF members contain similar spectral energy (Fig. 5). Figure 6 shows the resulting magnitude of ζ “anomalies” for the first 20 prior, posterior, and 6-h forecast members colored contours—with top and bottom panels corresponding to the EnKF and local PF, respectively, and the GDAS analysis (black) overlaid in each panel for reference. The resulting member fields contain regions of elevated ζ anomalies in and around Maria. The largest values coincide with similar mesoscale features found in the GDAS analysis, but with clear variability in size,

magnitude, and location. At this end of the spectrum, differences emerge between the EnKF and local PF experiments due to the observability of flow at these scales and discrepancies between how the two methods infer unobserved state variables.² While the two data assimilation methods produce qualitatively similar prior members at this time (Figs. 6a,d) the EnKF posterior members contain numerous small-scale ζ anomalies that do not exist in the prior (or local PF) members (Figs. 6b,e). The noted increase in small-scale (<150 km) ζ anomalies between the EnKF prior and posterior members presents an example of how the EnKF induces a notable increase in spectral energy at these scales, as quantified in Fig. 5. We also note that the variance in EnKF posterior members is not substantially different from the prior, owing to the use of RTPS when generating the posterior ensemble. For most members, the small-scale features generated by the EnKF update step vanish 6 h into the forecast, which is visible in comparisons between Figs. 6b and 6c. While the Maria example is purely anecdotal, the adjustment of winds during prediction steps are clearly visible in short- to medium-range ensemble forecasts generated over the full monthlong experiment. Figure 7 quantifies this behavior using domain-mean ensemble ζ spread, averaged over all 52 ensemble forecasts. The EnKF forecasts show a drop in spread in the first six hours of the forecasts, which precedes the expected increase at later lead times. This result is also consistent with the analysis of ensemble spread for PV, which shows a slower increase in spread for EnKF members versus PF members over the first 12 h—particularly in the lower troposphere.

Following the spectral analysis, we suspect forecast improvements obtained from the local PF come largely from more appropriate multivariate updates to ensemble members,

² The advective time scale for length scales near 100 km falls near or below the 6-h observation schedule used for data assimilation.




FIG. 7. Ensemble spread in EnKF (red line) and local PF (blue line) ζ forecasts averaged over all ensemble HWRf forecasts.

which may not be adequately captured in verifications of the ensemble mean at short lead times. Though not shown, similar behavior is found across thermodynamics variables in the model, but to a lesser degree. Localization is one factor that will inevitably lead to kinematic inconsistencies between posterior members and solutions permitted by the dynamical model (Kepert 2009; Greybush et al. 2011). While interpreting these findings, it is important to emphasize that localization serves a slightly different role in reducing filter degeneracy in the PF versus the EnKF (P22). In particular, the influence of localization on perturbation updates can be vastly different between the pair of filters. Unlike the EnKF, perturbation updates for the PF are a function of innovations, which means the influence of localization on higher-order posterior moments is also a function of innovations. Nevertheless, we do not expect the strategy adopted for localization in the PF to be better suited for reducing the observed noise at the mesoscale. We also note that an improper specification of representativeness error—combined with localization—could lead to an overfitting of observations during data assimilation that would induce differences between posterior and forecast members at smaller scales. These theories, however, are less likely given that the local PF uses shorter localization cutoff lengths yet produces a smaller increase in spectral density for shorter wavelengths. Furthermore, we note that the choice of posterior inflation for the EnKF also influence the makeup of posterior members and the resulting growth in error variance. Alternative approaches, which relax a fraction of the perturbation updates (Zhang et al. 2004), have been shown to improve balance in initial conditions and increase the rate at which ensemble spread grows during integration (Whitaker and Hamill 2012).

The next section adopts a simple two-dimensional model to reproduce and analyze findings identified in the HWRf experiments. We will demonstrate that a large portion of the spurious small-scale vorticity updates for the EnKF can be produced by inappropriate multivariate Gaussian approximations alone.

5. The non-Gaussian data assimilation problem posed by vortex alignment errors

The current section provides a more detailed comparison of the EnKF and local PF using a low-dimensional application that mimics the behavior uncovered in section 4. In general, we hypothesize that findings from section 4 come from non-Gaussianity in the multivariate prior densities used to characterize flow at scales < 150 km in the HWRf experiments. Numerical experiments performed in this section serve the purpose of reproducing the data assimilation challenges found from real-data experiments, while illustrating how the iterative local PF solves the underlying problem. The latter objective builds off of methodology discussed in P22, namely, weight regularization and tempering (see section 2), which are known to provide benefits for sparsely observed dynamical systems. While these strategies have demonstrated value for idealized low-dimensional models (see P22), this study is the first to explore their value for geophysical applications.

The idealized experiments use simulated observations from a kinematic, axisymmetric vortex, which introduces an application that can be solved easily using a Bayesian filter with little approximation. Hodyss and Reinecke (2013) show that large displacement errors in hurricane-like vortices tend to be characterized by complex multivariate structure in prior errors. They find that EnKF corrections to the mean can be greatly improved by adopting a quadratic nonlinear update—a finding that stems from large skewness in the presence of position uncertainty. We follow a similar methodology to explore challenges associated with displacement errors in the HWRf experiments, but focus primarily on updates made to individual ensemble members rather than ensemble statistics.

The low-dimensional application provides full control over the parameters used to define the data assimilation problem. We select choices of observation network and prior density that are motivated by real weather applications to compare ensemble updates performed with the filters used in section 4, namely, the ensemble square root EnKF and iterative local PF with three iterations (denoted local IPF). We supplement this comparison by also showing posterior members generated using a regularized local PF (denoted local RPF) and the standard bootstrap PF with no localization. The local RPF is identical to a single step of the local IPF—thus demonstrating the role of iterations—and the bootstrap PF provides estimates of the Bayesian posterior with a large sample approximation. The EnKF, local RPF, and local IPF all use the same ensemble size as section 4 ($N_e = 60$) and the bootstrap PF adopts a much larger ensemble of $N_e = 10^4$ so it can provide an accurate depiction of the true Bayesian solution. The choice of N_e for the bootstrap PF is the approximate

minimum N_e needed for this method to yield an effective ensemble size of 60 from posterior weights for the data assimilation problems constructed in this section.

a. Idealized vortex model

To reproduce data assimilation challenges posed by meso-scale weather systems with various degrees of prior uncertainty, we consider the case of a vortex in zero mean flow. We ignore the many complexities listed in the introduction for modeling subsynoptic flow and assume that the cumulative effects of model error, suboptimal data assimilation, and intrinsic predictability limits contribute to uncertainty in the location, size, and intensity of a discrete vortex. In reality, the same uncertainty sources exist over a spectrum of scales in the atmosphere, which complicates the interpretation of results for multiscale geophysical models.

For simplicity, the vortex wind field is modeled using a Rankine vortex profile (Acheson 1990). In cylindrical coordinates with the vortex center at the origin, the tangential winds (u_θ) are a function of radius (r) alone:

$$u_\theta = \begin{cases} U_\theta \frac{r}{R}, & r < R, \\ U_\theta \frac{R}{r}, & r \geq R, \end{cases} \quad (9)$$

The model has two parameters for controlling vortex size and intensity, U_θ and R , which specify maximum wind speed and maximum wind radii, respectively. Both the radial wind component (u_r) and vertical wind component (u_z) are assumed to be zero. This model provides a forced vortex in solid-body rotation for the region located R units from the domain center, and a free vortex outside R . To reproduce the data assimilation problem of interest, we consider a vortex on a 2D Cartesian domain comprised of 91×91 equally spaced grid points. At each grid point, u_θ is first calculated using (9) with r being the distance of each grid point from the domain center, defined by coordinates ($i_c = 46, j_c = 46$). We choose parameters for a “truth” state to be fixed at $U_\theta^t = 30 \text{ m s}^{-1}$ and $R^t = 12$, and adopt i_c and j_c as the true center location. Figure 8a shows values of wind speed for a cross section through the domain center for the designated true solution, and Fig. 8b shows a single 15 m s^{-1} wind speed contour for this vortex on the 2D domain.

The designated truth remains near the mode of each prior used for numerical experiments. The chosen priors and observation networks allow the EnKF and local PF to both produce relatively accurate estimates of the posterior mean—despite large deviations from Gaussianity. By construction, most of the observed differences in behavior can then be attributed to differences in perturbations about the mean solution. This design choice provides a closer match to the data assimilation experiments performed with HWRP—where the EnKF produces accurate posterior mean solutions, but shows evidence of improperly specified cross-variable error covariance in posterior perturbations.

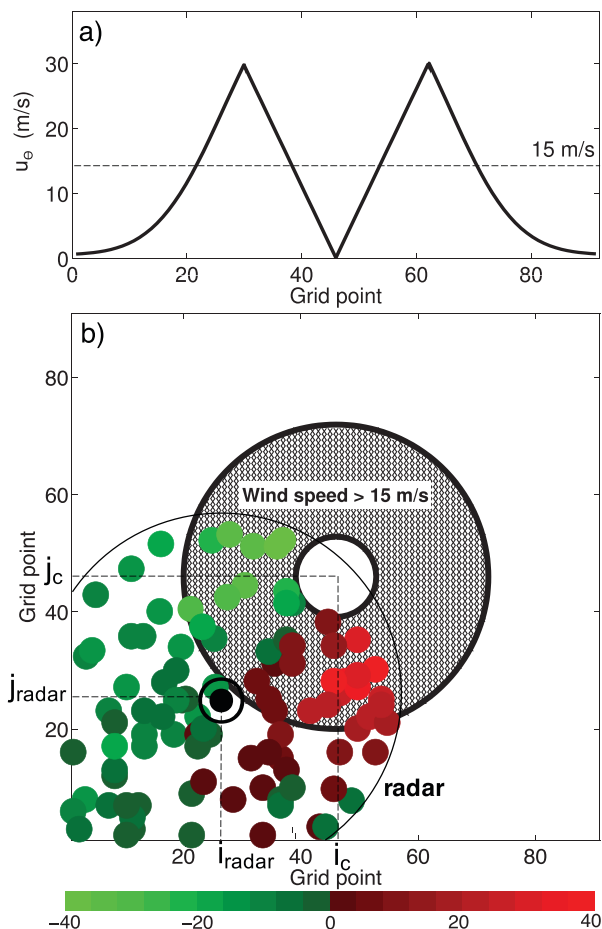


FIG. 8. (a) Tangential wind speed as a function of grid points, calculated using the Rankine vortex model with point 46 as center location. (b) The 15 m s^{-1} wind speed contours for the vortex placed on a 2D grid; values greater than 15 m s^{-1} are indicated by hatched region. The green and red markers indicate the location and magnitude of radial wind observations created for a synthetic radar located at coordinate $(i_{\text{radar}}, j_{\text{radar}})$.

b. Observation network

We generate two observation networks from the designated truth solution. Both networks resemble uniformly sampled synthetic observations from a “radar” placed at coordinate $(i_{\text{radar}} = 25, j_{\text{radar}} = 25)$ in the bottom left of the domain (black marker in Fig. 8b). At each location, winds from the truth state are projected onto the direction of a hypothetical radar beam looking outward from the radar. A random error is then drawn from $N(0, \sigma_o^2)$ for $\sigma_o = 3$ and added to each wind value to form y_i for $i = 1, 2, \dots, N_y$. The first network covers the entire domain with a uniform density of $N_y = 300$ radial winds. The second network restricts the radar to have a scan radius of 30 grid points, reducing N_y to 100, but retaining the same data density. The green and red dots in Fig. 8b indicate the location and value of measurements from the second observation network, which covers only one quadrant of the vortex.

The first observation network presents a scenario where the vortex is well observed at the current time, even if prior uncertainty is large (to be discussed in section 5c). The second observation network presents an application where the vortex is only partially observed. The data assimilation problem posed by the second network is similar to the problem of assimilating airborne Doppler radar observations collected from tropical cyclones (Zhang et al. 2011; Weng and Zhang 2012). In this case, radar measurements collected during aircraft reconnaissance flights can only observe portions of the tropical cyclone vortex at a given time, thus requiring data assimilation to infer winds on unobserved quadrants of the storm. A similar challenge exists for assimilating measurements for continental mesoscale convective systems. Land-based radars provide observations at high spatial and temporal frequencies, including winds at multiple scan angles—but not without gaps in vertical and horizontal coverage. Therefore, it is common for a radar to observe only portions of mesovortices found within bow echoes and supercell thunderstorms.

The second observation network is likely more representative of the HWRF experiments presented in section 4, where a majority of measurements assimilated in the vicinity of storms come from radiances or radiance-derived atmospheric motion vectors, which must infer the full 3D wind, mass, and moisture variables from prior error statistics.

c. Choices of prior

For this application, the objective is to generate samples from the posterior density $p(\mathbf{x}|\mathbf{y})$ where \mathbf{x} contains winds on the 2D grid described in section 5a and \mathbf{y} is populated with synthetic radar observations. To mimic a real data assimilation application, this exercise requires first drawing samples from a prior density $p(\mathbf{x})$, which can be transformed into samples from $p(\mathbf{x}|\mathbf{y})$. Each experiment uses one of four different choices of prior, which we sample from by randomly perturbing vortex parameters before calculating the wind components at each grid point. This approach bypasses the need to define each $p(\mathbf{x})$ explicitly.

The first three priors use identical axisymmetric vortices, but induce position uncertainty by perturbing the center position of each vortex. In doing so, the resulting priors maintain the dominant source of error variance (cf. Figs. 3 and 4 of Poterjoy and Zhang 2011), while yielding a challenging data assimilation problem for Gaussian filters. We draw the position of each member randomly from a Gaussian centered on the true vortex center; i.e., each member uses coordinates i_n and j_n sampled from $N(i_c, \sigma_p^2)$ and $N(j_c, \sigma_p^2)$, respectively, for $n = 1, \dots, N_e$, where σ_p is a prescribed position error standard deviation that changes for each prior. All vortices generated in these experiments have zero radial winds when viewed in a cylindrical coordinate system with the vortex center at the origin. This coordinate system, however, is unobtainable for real applications, which are typically characterized by multiple scales of motion with ambiguous center positions. To be consistent with weather models, we

convert all winds into zonal (\mathbf{u}) and meridional (\mathbf{v}) components so that

$$\mathbf{x} = \begin{pmatrix} \mathbf{u} \\ \mathbf{v} \end{pmatrix}.$$

For the priors specified here, $p(\mathbf{x})$ is close to a Gaussian when position error is sampled from a Gaussian with $\sigma_p \ll R$, where R controls the size of the vortex (Chen and Snyder 2007). Therefore, σ_p/R provides a controllable means of changing the shape of $p(\mathbf{x})$ to improve or reduce the fidelity of Gaussian assumptions. We explore this property by performing data assimilation experiments with three different priors, each using $\sigma_p/R = (0.25, 0.50, 1.0)$ for fixed R and U_θ . To place in the context of a real-world application, the average 6-h error from official National Hurricane Center (NHC) model-guided predictions was about 30 km during the 2010–18 Atlantic hurricane seasons (NWS 2018), which is similar to the radius of maximum winds expected for a mature hurricane (Knaff et al. 2015). The ratio σ_p/R is presumably smaller for well-observed and larger storms, which motivates the range of values chosen in this demonstration.

We also generate samples from a fourth prior that uses $\sigma_p/R = 1.0$, but allows R and U_θ to vary across members. This prior contains both size and intensity uncertainty and will be denoted by experiments entitled “structure uncertainty.” Despite the naming convention, this prior has the same amount of position uncertainty as the third prior described above, but with R drawn from $N(R', 9)$ and U_θ drawn from $N(U_\theta', 9)$.

d. Experiment results

This subsection discusses results obtained by performing data assimilation using the described sets of observations and priors. All experiments use localization coefficients modeled from an exponentially decaying function: $\rho_{ij} = \exp[-d_{ij}^2/(2r_{loc}^2)]$, where d_{ij} is the physical distance between observation y_i and state variable x_j . Table 2 summarizes the data assimilation and model parameters used for each prior, which are kept constant for the pair of observation networks. Parameter choices follow from a tuning of each method based on posterior mean RMSEs from the first observation network, which is typical for real data assimilation applications. The localization length scales arrived at for this experiment are about a factor of 10 greater than the domain width used in the experiments. This data assimilation configuration roughly matches the setup for the vorticity anomalies in Fig. 6, except no observations exist far from the idealized vortex. It is also

TABLE 2. Choices of parameters for Rankine vortex experiments.

Method	r_{loc}	N_{eff}	γ
EnKF	800	—	—
Local RPF	800	6	0.3
Local IPF	800	6	0.3
Bootstrap PF	—	—	—

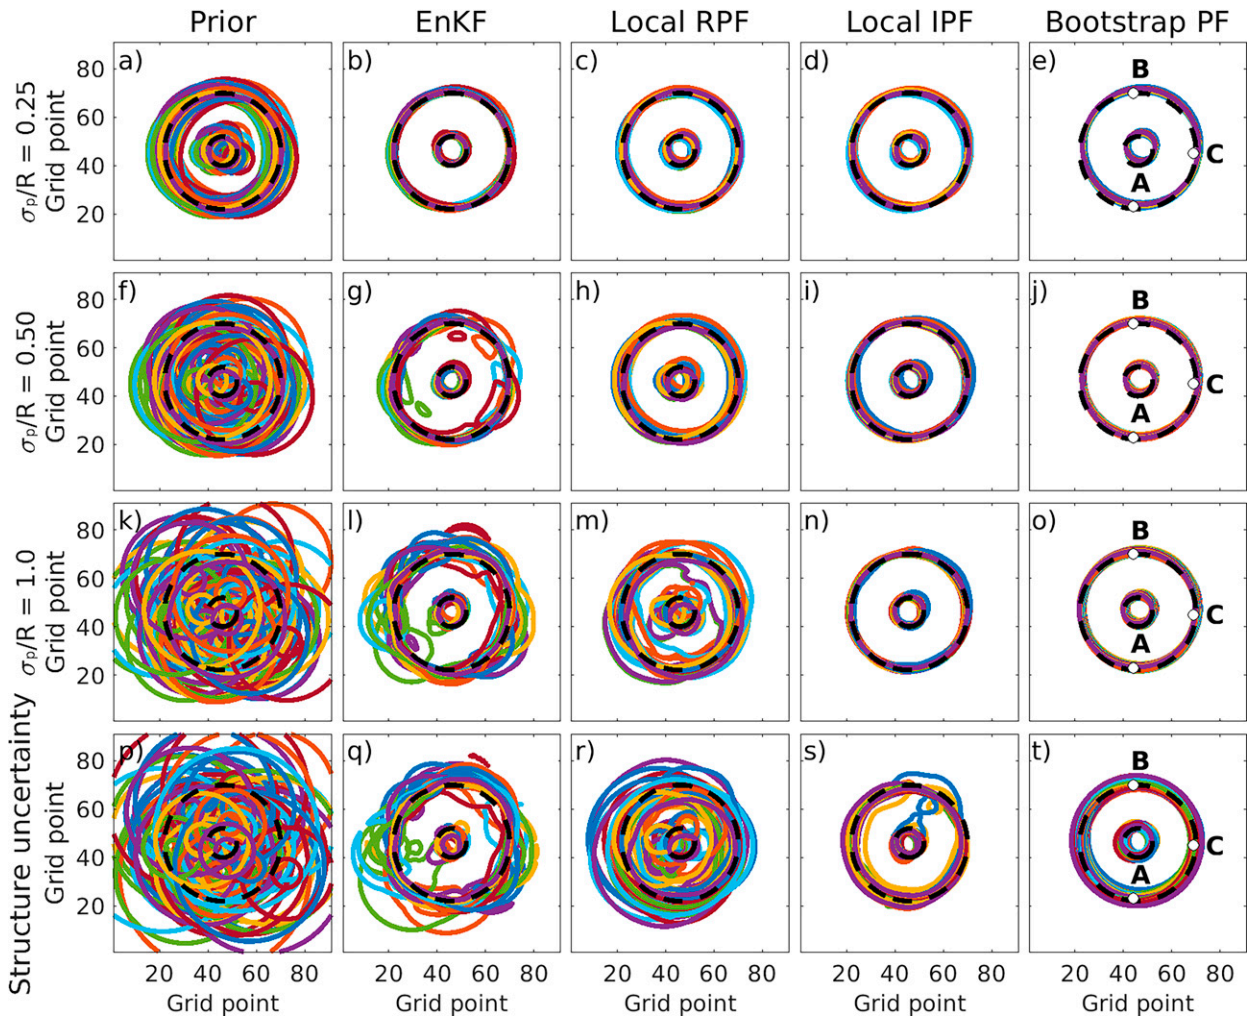


FIG. 9. Colored contours show 15 m s^{-1} wind for (from left to right) prior members, EnKF posterior members, local RPF posterior members, local IPF posterior members, and bootstrap PF posterior members. Members are plotted for the (from top to bottom) $\sigma_p/R = 0.25$ experiment, $\sigma_p/R = 0.50$ experiment, $\sigma_p/R = 1.0$ experiment, and “structure uncertainty” experiment. A black dashed contour shows the true winds and white markers indicate the location of points A, B, and C, chosen for scatterplots in Figs. 10 and 12.

noteworthy to mention that we omit the posterior relaxation step from the EnKF, which was adopted in HWRP experiments to maintain spread through successive cycles. While this step is not needed for the interpretation of findings from experiments we note that the additional inflation step would only amplify any detrimental implications uncovered for non-Gaussian priors.

Figure 9 shows 15 m s^{-1} wind speed contours for posterior members obtained from the first observation network ($N_y = 300$) using each of the four choices of prior. For reference, the first column shows the prior members adopted for each experiment. For experiments that use $\sigma_p/R = 0.25$ for the prior (Figs. 9a,b) each data assimilation method produces posterior members that are visibly similar (first row of Fig. 9). This finding is expected, owing to the Gaussian approximations becoming increasingly more valid for small σ_p/R (Chen and Snyder 2007). Increasing position error by a factor of 2

(second row of Fig. 9) and 4 (third row of Fig. 9), however, yields vast differences across experiments. The EnKF produces large asymmetries in members that no longer resemble the Bayesian solution indicated in the fifth column. Likewise, the local RPF is constructed to maintain low-likelihood particles in order to retain a user-specified effective ensemble size. Therefore, the regularized version of the PF exhibits much larger posterior uncertainty in 15 m s^{-1} wind contours than the Bayesian estimate.

By including an additional recursion to the PF update steps, (7) and (8) shift particles incrementally into the posterior distribution assumed by the local PF (see section 2.1 of P22)—with each iteration yielding a diverse set of particles that match posterior moments. In doing so, the local IPF relies less on regularization to obtain a set of posterior particles with a specified effective ensemble size. As described in P22, this step also allows the PF to sample more accurately

from non-Gaussian posterior distributions using knowledge of the first two moments alone. The resulting iterations produce members that more closely resemble samples from the Bayesian posterior—comparing the fourth and fifth columns of Fig. 9. The advantages of the iterations are more substantial when structure uncertainty is introduced to the prior. The 60-member ensemble produces very few particles that are close to the observed vortex position, size, and intensity for this problem, which forces the local RPF to sample several posterior members that are far from the observations.

We further examine findings from the first set of experiments by visualizing portions of the prior and posterior error distributions depicted by members. For this purpose, we focus more narrowly on zonal winds for a pair of points A and B [denoted $u(A)$ and $u(B)$, respectively], which are indicated in the fifth column of Fig. 9. These points are chosen near the outer 15 m s^{-1} contours shown in Fig. 9. In the context of data assimilation, the marginal dependence between the chosen points controls how information is spread spatially over the domain, thus determining how measurements on one side of a vortex update unobserved variables on the opposite side. At these points, we plot samples from the joint prior error distribution of $u(A)$ and $u(B)$ using the $N_e = 10^4$ ensemble adopted for the bootstrap PF experiments (blue markers in Fig. 10). Similar to the findings of Hodyss and Reinecke (2013) for vortex pressure fields, it is clear that the prior has a complex multivariate structure that is non-Gaussian for a wide range of parameters. For low σ_p/R , the prior exhibits a triangular shape that is bounded on all three sides (first row of Fig. 10). Increasing σ_p/R beyond a critical value extends the region of nonzero probability to include a second basin, which exhibits two modes that extend outward from the origin and become increasingly narrower for large magnitudes of $u(A)$ and $u(B)$ (third row of Fig. 10). By inducing structure uncertainty into the displaced vortices, the resulting prior retains a similar shape, but with a blurring of the boundaries that define the region of nonzero probability (fourth row of Fig. 10). Furthermore, we note that the marginal probability densities estimated for point winds within the vortex tend to be close to Gaussian, but shift toward skewed distributions for large σ_p/R , similar to the behavior described by Hodyss and Reinecke (2013). For example, the gray curves plotted along x and y axes in Fig. 10 show kernel-estimated marginal densities for $u(A)$ and $u(B)$, respectively, for each prior. Therefore, the non-Gaussian data assimilation problem constructed in this section is one characterized by nonlinear relationships across random variables, but near-Gaussian marginal distributions. Such behavior is not easily detected using ensemble sizes readily available for weather prediction and is difficult to identify in the HWRf experiments, despite similar displacement errors. We will demonstrate that the nonlinear filters have clear advantages for prior distributions of this type, even when deviations from Gaussianity are not clearly identifiable.

As illustrated in Fig. 10, the data assimilation challenge for a displaced vortex then becomes one of how to shift members from various parts of a rather complicated prior error distribution into the subset of this domain determined by

measurements to contain nonzero posterior probability. The red markers in Fig. 10 depict how each filter performs this task—with varying degrees of success. For this particular application, the region of highest prior probability exists along the lower-right boundary, which contains the truth in our simulations (indicated by yellow markers in the figure). When provided with measurements from the radar, the Bayesian-estimated posterior is restricted to solutions near this boundary. The local IPF provides the closest match to the Bayesian solution, followed by the local RPF—which again provides a gross overestimate of the posterior uncertainty. Because the EnKF assumes the prior and posterior distributions are Gaussian, the mapping of prior members into the posterior can be conceptualized as a multivariate linear regression step with regression coefficients determined by (2). This assumption causes the EnKF to retain the shape of the prior in the posterior estimate. Prior members sampled along the periphery of the distribution—such as those already along the lower-right bounding region—become shifted into portions of the domain that have zero prior probability. This factor leads to various unphysical artifacts to appear in wind contours such as the large vortex asymmetries depicted in Fig. 9.

As demonstrated for the HWRf experiments, both the local PF and EnKF exhibit spuriously large power spectral energy for the posterior wind field at small scales. We apply a similar methodology outlined in section 4 to show that asymmetries induced by data assimilation assumptions manifest themselves in a similar way for the axisymmetric vortex application. To reproduce this analysis, we calculate ζ for EnKF and local IPF posterior members and perform a 2D spectral decomposition of the asymmetric part of ζ (denoted ζ'); i.e., we locate the center of each vortex and remove an azimuthal mean calculated with respect to each center. Removing the azimuthal mean helps isolate parts of the solution that are not physically allowed by the Rankine vortex model, since any nonzero value for ζ' is spurious. Figure 11 shows the mean of these calculations, averaged over the EnKF and local IPF ensembles for the four choices of prior.³ Consistent with Fig. 5 both data assimilation methods produce spuriously large spectral density, but with higher values induced by the EnKF for the first 10 wavenumbers owing to Gaussian assumptions. The spectral energy grows expectedly with σ_p/R for the EnKF, and is larger than the local IPF for all priors, even when $\sigma_p/R = 0.25$.

The finding that the EnKF yields measurably higher spectral energy than the local IPF for the first prior in Fig. 5 is unexpected, given the visual comparisons of wind contours in Fig. 9. We perform a closer inspection of this result by enlarging the portion of the domain containing the posterior densities in Figs. 10a and 10c, and plotting EnKF and local IPF members on the same graphic (indicated by red and green markers in Fig. 12a, respectively). This exercise illustrates that

³ These results are plotted as a function of wavenumber rather than wavelength to avoid any over-interpretation of length-scale dependence, which is ultimately determined arbitrarily by choices of domain and vortex size.

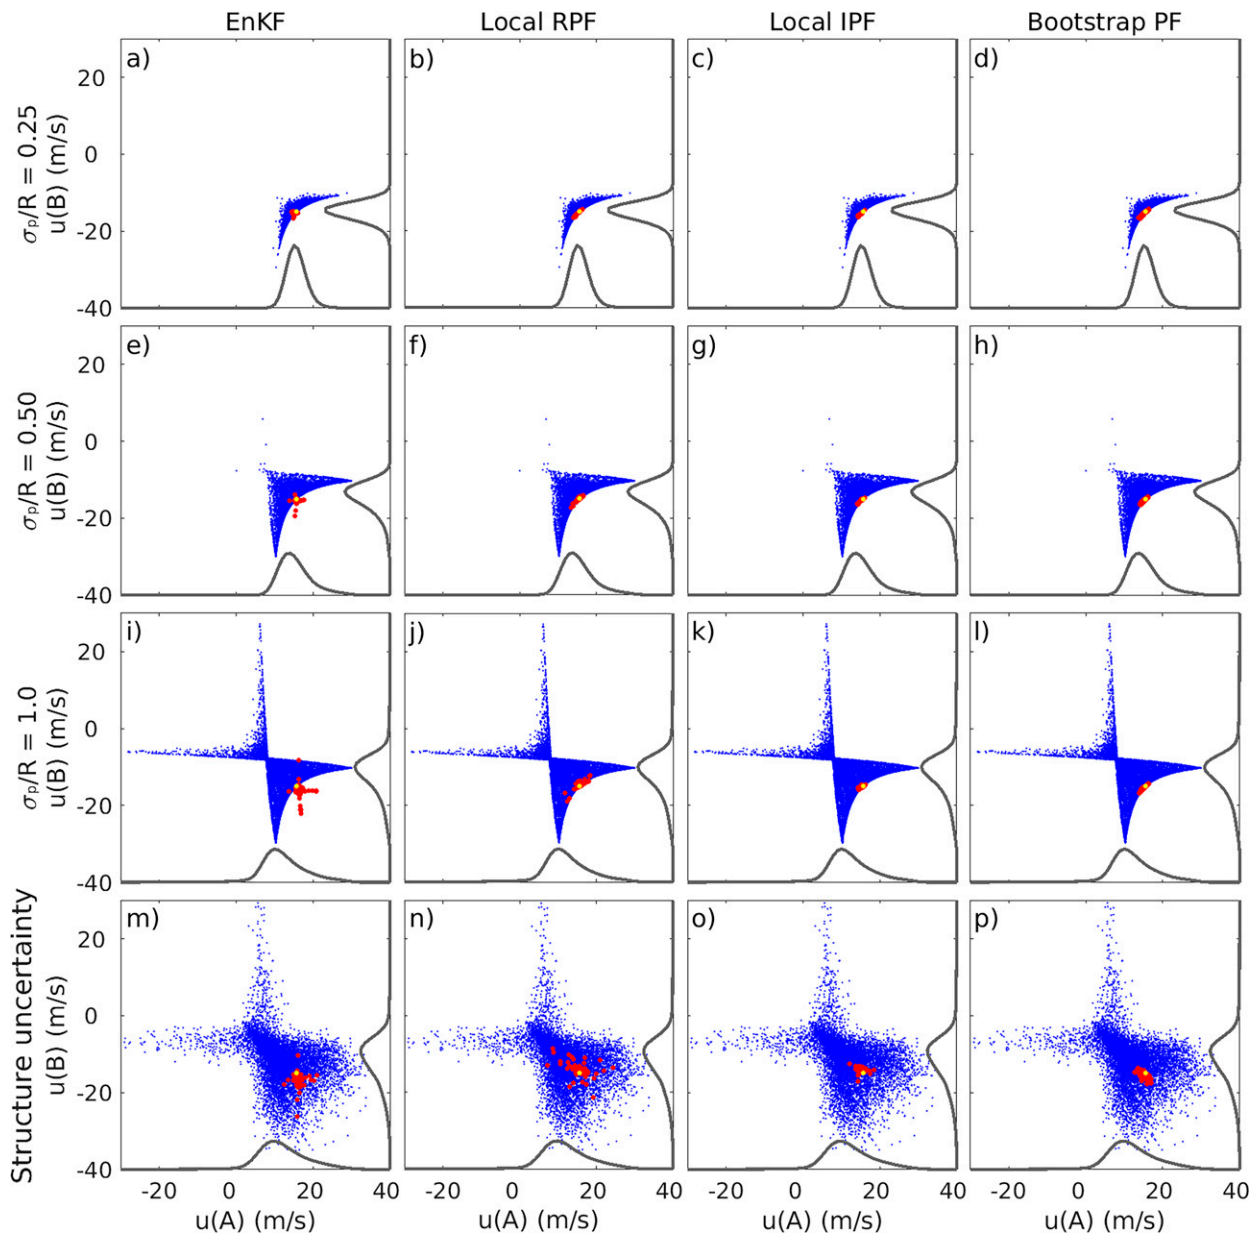


FIG. 10. Scatterplots of prior u values for points A and B (blue dots; indicated by white markers in the fifth column of Fig. 9). Marginal pdfs for each variable (gray contours) are scaled and plotted on each axis for reference. Red markers in each column correspond to (from left to right) EnKF posterior members, local RPF posterior members, local IPF posterior members, and bootstrap PF posterior members. The yellow marker indicates the true solution.

the first prior distribution is sufficiently far from a Gaussian to cause the EnKF to shift members outside the bounded region, which occurs to a much smaller degree with the local IPF. Differences between the two filters become even more apparent when a third variable is included in the joint density comparisons. For this purpose, we choose zonal winds at a point C located half-way between A and B , in the right hemisphere of the vortex. For the three-variable prior density, solutions are restricted to follow the hyperplane illustrated in Fig. 12b, which is estimated from the $N_e = 10^4$ sample and plotted in its

place. The heuristic decoupling of variables brought about by localization causes the local IPF to produce members that land off of the hyperplane; these solutions represent members with unphysical behavior, such as asymmetries. Nevertheless, the local IPF members are much closer to physically allowable solutions than those produced by the EnKF. Both methods yield asymmetries that translate into an increase in spectral density, but the contribution from Gaussian assumptions is far more egregious than the localization strategy adopted for this application.

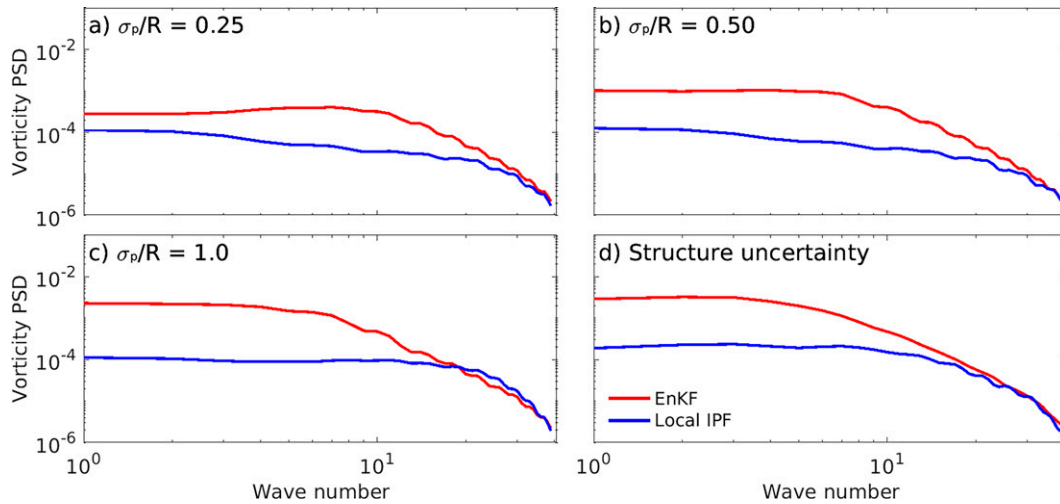


FIG. 11. Power spectral density for ζ averaged over EnKF members (red lines) and local IPF members (blue lines) for all four choices of prior used in idealized vortex experiments.

Results obtained from assimilating the first observation network already provide insight into the HWRF experiments in section 4. In summary, deviations from a Gaussian can occur for even small alignment errors, which cause filters that make this assumption to induce a spurious increase in spectral energy. Gaussian approximations tend to break down for smaller scales, which likely leads to this result for ζ characterized by length scales < 150 km in the HWRF experiments. An additional set of idealized experiments performed using the second observation network present an identical set of challenges, except a larger number of state variables are completely unobserved and must be inferred through data assimilation. In this regard, the second network has more in common with real geophysical observing systems than the first one.

Figure 13 shows prior and posterior wind contours for all experiments performed using the second, sparser observation network. Asymmetries induced by Gaussian assumptions in

the EnKF become more visible when the radar measures only a portion of the vortex, as demonstrated for all priors including when $\sigma_p/R = 0.25$. Likewise, the local IPF continues to produce posterior members that more closely resemble the Bayesian solution for all priors. The difficulty posed by the sparse network for the EnKF follows the same problems highlighted by Fig. 12; i.e., multivariate updates determined by the mean and error covariance alone can easily shift members into unphysical directions. This deficiency becomes less problematic when the vortex approaches full observability—as in the first observation network tested in this section. With a spatially dense network of observations, the EnKF can still provide an accurate estimate of the posterior mean (e.g., Posselt and Bishop 2012). In fact, posterior mean RMSEs are comparable for the EnKF and local IPF for experiments performed with the first observation network, despite the nonphysical behavior exhibited by posterior members (not shown). Under these circumstances, posterior

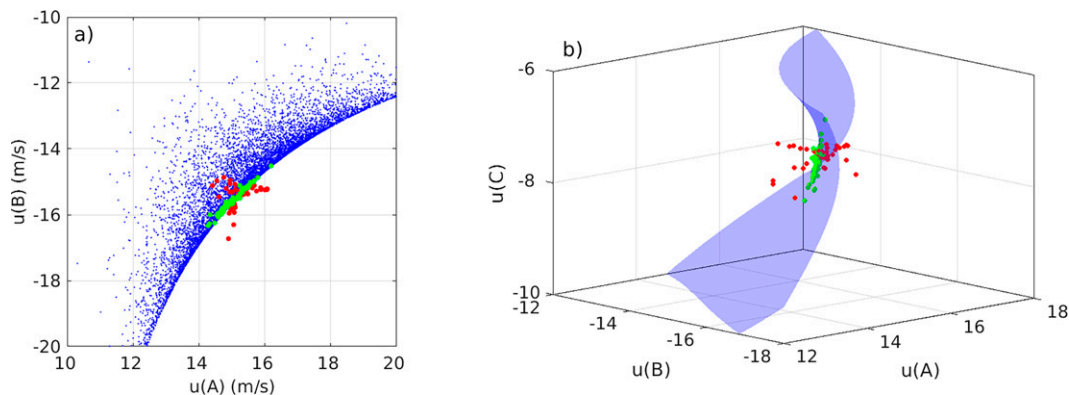


FIG. 12. (left) Enlarged version of Fig. 10a, but with local IPF members (green markers) plotted alongside EnKF members (red markers). (right) Probability density visualization is extended to the left to include a third random variable $u(C)$; it also replaces prior samples (blue markers) with a 3D depiction of the hyperplane that all prior samples land on.

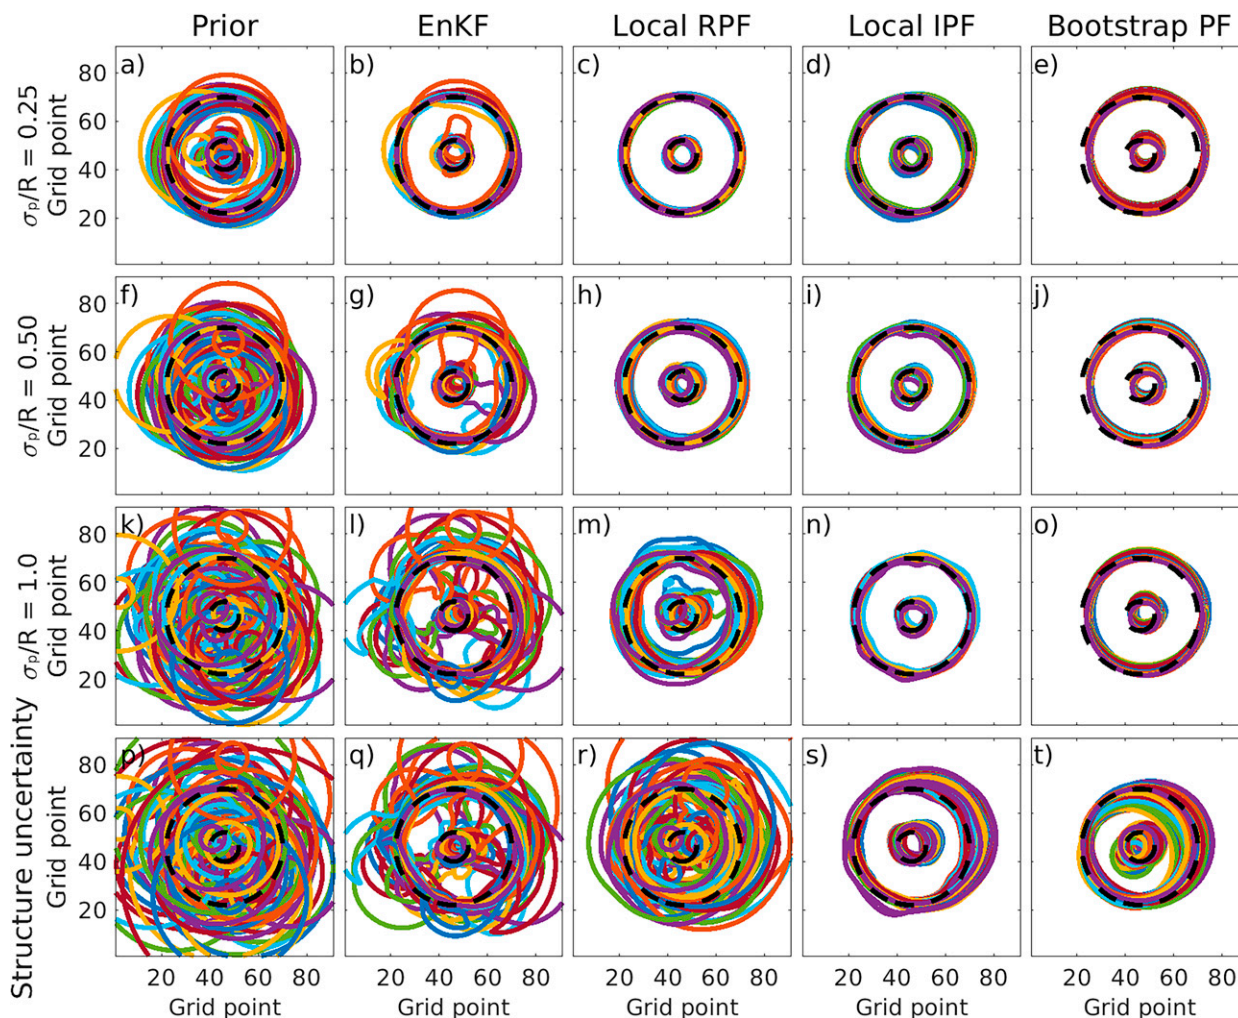


FIG. 13. As in Fig. 9, but for a partially observed vortex (second observation network).

members can still draw closer to the true solution as the variance approaches zero. Real weather applications, however, still suffer from the practical limitations discussed in section 1, which means a near-zero posterior variance is currently not obtainable.

6. Summary and conclusions

The current study presents findings from monthlong regional modeling experiments performed using an experimental HWRF weather prediction system (Poterjoy et al. 2021). The prediction system was designed to explore science questions related to multiscale atmospheric modeling within future implementations of the NOAA Unified Forecasting System (UFS)—including a prototype NOAA Hurricane Analysis and Forecast System (HAFS; Marks et al. 2019). This research focuses more narrowly on data assimilation challenges posed by multivariate non-Gaussian prior uncertainty, which often exists for loosely constrained storm-scale weather phenomena. For this purpose, an iterative localized

PF (Poterjoy 2022, hereafter P22) is added to the NOAA GSI data assimilation package and compared with a square-root EnKF (Whitaker and Hamill 2002). The EnKF—in various forms—is the current method of choice for initializing probabilistic forecasts within operational weather prediction systems (Bannister 2017). While EnKFs use Gaussian assumptions for prior errors and likelihoods, PFs adopt a non-parametric representation of priors and provide much greater flexibility for choosing likelihood functions that match error characteristics of observations. Both methods rely on a sample representation of probability densities used during data assimilation, and implement common practices for reducing the sample size needed to estimate these densities for geophysical models defined on spatial grids (i.e., localization).

Posterior and forecast ensembles produced by the pair of data assimilation methods are scrutinized using monthlong sequential data assimilation experiments performed over an active month of the 2017 Atlantic hurricane season. Local PF forecasts produce slightly more skillful ensemble-mean predictions at later lead times and exhibit a more linear increase

in ensemble spread than EnKF forecasts. The most striking differences between posterior solutions exist at scales where large alignment errors in mesoscale weather features have a significant role in determining prior error statistics. At these scales, EnKF updates induce a large spurious increase in spectral density that exists to a much lesser extent in local PF members. This result occurs despite the EnKF using less-restrictive localization length scales than the local PF, which is often suggested to be a major source of imbalance during data assimilation.

Idealized low-dimensional experiments are constructed to provide a more rigorous investigation into results obtained from the real-data application. In addition to allowing full control over prior error sources, these experiments permit the use of the standard (nonlocalized) bootstrap PF with large samples, thus providing a Bayesian posterior for verification. Findings from these experiments illustrate the complex, non-Gaussian shape of multivariate prior wind errors in the presence of position uncertainty, which is known to pose large challenges for Gaussian filters (Lawson and Hansen 2005; Hodys and Reinecke 2013). Linear updates performed by the EnKF draw members into unphysical directions that manifest as a spurious increase in spectral density for posterior members. For this specific application, Gaussian approximations on the prior density are more egregious than modulating dependence across variables using localization. The EnKF produces a similar spurious increase in spectral density for scales that exhibit large displacement errors, thus leading us to speculate that similar behavior exists in HWRP experiments. We also acknowledge that improperly specified representativeness error and differences in localization strategy may also contribute to this behavior. The idealized demonstration also does not explain the dynamic implications of such behavior, which is a topic for a future study. It does, however, provide evidence to the hypothesis that imbalances induced by Gaussian approximations can significantly degrade the performance of ensemble forecasts at later lead times. Furthermore, marginal probability densities estimated for winds at discrete grid points in the idealized application remain close to Gaussian, even if their joint distributions are far from Gaussian. For applications of this type, univariate hypothesis testing may be insufficient for detecting deviations from Gaussianity, which complicates the formation of adaptive strategies for combining PFs with EnKFs (e.g., Kurosawa and Poterjoy 2021; P22).

Findings from real and idealized experiments suggest that benefits can be obtained by adopting data assimilation methodology that either allows for nonlinear dependence in random variables or transform members to resemble samples from a multivariate Gaussian. The local PF adopted for this study is a clear example of a method that directly handles nonlinear dependence, but similar benefits may be realizable by addressing alignment errors separately during data assimilation (e.g., Hoffman et al. 1995; Lawson and Hansen 2005; Ravela et al. 2007; Nehrkorn et al. 2015; Stratman et al. 2018; Ying 2019).

Last, the current study does not explore nonlinear observation processes and nonlinear measurement operators, which

provide additional sources of non-Gaussianity that are common for real geophysical applications. A natural application of the local PF beyond the current study is for the assimilation of all-sky radiance measurements, which encompass the above challenges. While the current implementation of the local PF is likely too computationally demanding for operational use, major cost savings can be obtained by adopting alternative tempering strategies discussed in P22 or performing particle updates only when a user-specified effective ensemble size is reached. Both strategies will be discussed alongside results in future studies that apply the local PF for high-resolution regional weather applications.

Acknowledgments. This study was supported by NOAA Grant NA19NES4320002 [Cooperative Institute for Satellite Earth System Studies (CISESS)] at the University of Maryland/ESSIC, NOAA Grant NA20OAR4600281, and a U.S. National Science Foundation CAREER Award AGS1848363. Computing for this project was provided by the NOAA Jet high performance computing system. The author thanks Chris Snyder and two anonymous reviewers for insightful comments that improved the quality of the manuscript.

REFERENCES

- Acheson, D. J., 1990: *Elementary Fluid Dynamics*. Oxford University Press, 397 pp.
- Aligo, E., B. S. Ferrier, J. Carley, E. Rogers, M. Pyle, S. J. Weiss, and I. L. Jirak, 2014: Modified microphysics for use in high resolution NAM forecasts. *27th Conf. on Severe Local Storms*, Madison, WI, Amer. Meteor. Soc., 16A.1, <https://ams.confex.com/ams/27SLS/webprogram/Paper255732.html>.
- Anderson, J. L., 2001: An ensemble adjustment Kalman filter for data assimilation. *Mon. Wea. Rev.*, **129**, 2884–2903, [https://doi.org/10.1175/1520-0493\(2001\)129<2884:AEAKFF>2.0.CO;2](https://doi.org/10.1175/1520-0493(2001)129<2884:AEAKFF>2.0.CO;2).
- , 2010: A non-Gaussian ensemble filter update for data assimilation. *Mon. Wea. Rev.*, **138**, 4186–4198, <https://doi.org/10.1175/2010MWR3253.1>.
- , 2020: A marginal adjustment rank histogram filter for non-Gaussian ensemble data assimilation. *Mon. Wea. Rev.*, **148**, 3361–3378, <https://doi.org/10.1175/MWR-D-19-0307.1>.
- Atlas, R., V. Tallapragada, and S. G. Gopalakrishnan, 2015: Advances in tropical cyclone intensity forecasts. *Mar. Technol. J.*, **49**, 149–160, <https://doi.org/10.4031/MTSJ.49.6.2>.
- Bannister, R. N., 2017: A review of operational methods of variational and ensemble-variational data assimilation. *Quart. J. Roy. Meteor. Soc.*, **143**, 607–633, <https://doi.org/10.1002/qj.2982>.
- Bengtsson, T., C. Snyder, and D. Nychka, 2003: Toward a nonlinear ensemble filter for high dimensional systems. *J. Geophys. Res.*, **108**, 8775, <https://doi.org/10.1029/2002JD002900>.
- Bishop, C. H., 2016: The GIGG-EnKF: Ensemble Kalman filtering for highly skewed non-negative uncertainty distributions. *Quart. J. Roy. Meteor. Soc.*, **142**, 1395–1412, <https://doi.org/10.1002/qj.2742>.
- , B. J. Etherton, and S. J. Majumdar, 2000: Adaptive sampling with the ensemble transform Kalman filter. *Mon. Wea. Rev.*, **129**, 420–436, [https://doi.org/10.1175/1520-0493\(2001\)129<0420:ASWTET>2.0.CO;2](https://doi.org/10.1175/1520-0493(2001)129<0420:ASWTET>2.0.CO;2).

- Biswas, M. K., and Coauthors, 2018: Hurricane Weather Research and Forecasting (HWRF) model: 2017 Scientific documentation. NCAR Tech. Note NCAR/TN-544+STR, 111 pp., <https://opensky.ucar.edu/islandora/object/technotes%3A563/datastream/PDF/view>.
- Bloom, S. C., L. L. Takacs, A. M. da Silva, and D. Ledvina, 1996: Data assimilation using incremental analysis updates. *Mon. Wea. Rev.*, **124**, 1256–1271, [https://doi.org/10.1175/1520-0493\(1996\)124<1256:DAUIAU>2.0.CO;2](https://doi.org/10.1175/1520-0493(1996)124<1256:DAUIAU>2.0.CO;2).
- Buehner, M., P. L. Houtekamer, C. Charette, H. Mitchell, and B. He, 2010a: Intercomparison of variational data assimilation and the ensemble Kalman filter for global deterministic NWP. Part I: Description and single-observation experiments. *Mon. Wea. Rev.*, **138**, 1550–1566, <https://doi.org/10.1175/2009MWR3157.1>.
- , —, —, —, and —, 2010b: Intercomparison of variational data assimilation and the ensemble Kalman filter for global deterministic NWP. Part II: One-month experiments with real observations. *Mon. Wea. Rev.*, **138**, 1567–1586, <https://doi.org/10.1175/2009MWR3158.1>.
- Chen, Y., and C. Snyder, 2007: Assimilating vortex position with an ensemble Kalman filter. *Mon. Wea. Rev.*, **135**, 1828–1845, <https://doi.org/10.1175/MWR3351.1>.
- Courtier, P., J.-N. Thepaut, and A. Hollingsworth, 1994: A strategy for operational implementation of 4D-Var, using an incremental approach. *Quart. J. Roy. Meteor. Soc.*, **120**, 1367–1387, <https://doi.org/10.1002/qj.49712051912>.
- Dirren, S., M. Didone, and H. C. Davies, 2003: Diagnosis of “forecast-analysis” differences of a weather prediction system. *Geophys. Res. Lett.*, **30**, 2060, <https://doi.org/10.1029/2003GL017986>.
- Doucet, A., N. de Freitas, and N. Gordon, 2001: An introduction to sequential Monte Carlo methods. *Sequential Monte Carlo Methods in Practice*, A. Doucet, N. de Freitas, and N. Gordon, Eds., Springer-Verlag, 2–14.
- Emerick, A. A., and A. C. Reynolds, 2012: History matching time-lapse seismic data using the ensemble Kalman filter with multiple data assimilations. *Comput. Geosci.*, **16**, 639–659, <https://doi.org/10.1007/s10596-012-9275-5>.
- Evensen, G., 1994: Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics. *J. Geophys. Res.*, **99**, 10 143–10 162, <https://doi.org/10.1029/94JC00572>.
- Farchi, A., and M. Bocquet, 2018: Review article: Comparison of local particle filters and new implementations. *Nonlinear Processes Geophys.*, **25**, 765–807, <https://doi.org/10.5194/npg-25-765-2018>.
- Feng, J., X. Wang, and J. Poterjoy, 2020: A comparison of two local moment-matching nonlinear filters: Local particle filter (LPF) and local nonlinear ensemble transform filter (LNETF). *Mon. Wea. Rev.*, **148**, 4377–4395, <https://doi.org/10.1175/MWR-D-19-0368.1>.
- Fletcher, S. J., and M. Zupanski, 2006: A data assimilation method for log-normally distributed observational errors. *Quart. J. Roy. Meteor. Soc.*, **132**, 2505–2519, <https://doi.org/10.1256/qj.05.222>.
- Gaspari, G., and S. E. Cohn, 1999: Construction of correlation functions in two and three dimensions. *Quart. J. Roy. Meteor. Soc.*, **125**, 723–757, <https://doi.org/10.1002/qj.49712555417>.
- Gopalakrishnan, S. G., Q. Liu, T. Marchok, D. Sheinin, N. Surgi, R. Tuleya, R. Yablonsky, and X. Zhang, 2010: Hurricane Weather Research and Forecasting (HWRF) model scientific documentation. NCAR Tech. Note, 75 pp.
- Gordon, N. J., D. J. Salmond, and A. F. M. Smith, 1993: Novel approach to nonlinear/non-Gaussian Bayesian state estimation. *IEE Proc.*, **140**, 107–113, <https://doi.org/10.1049/ip-f-2.1993.0015>.
- Greybush, S. J., E. Kalnay, T. Miyoshi, K. Ide, and B. Hunt, 2011: Balance and ensemble Kalman filter localization techniques. *Mon. Wea. Rev.*, **139**, 511–522, <https://doi.org/10.1175/2010MWR3328.1>.
- Hodyss, D., 2012: Accounting for skewness in ensemble data assimilation. *Mon. Wea. Rev.*, **140**, 2346–2358, <https://doi.org/10.1175/MWR-D-11-00198.1>.
- , and P. A. Reinecke, 2013: Skewness of the prior through position errors and its impact on data assimilation. *Data Assimilation for Atmospheric, Oceanic, and Hydrologic Applications*, S. K. Park and L. Xu, Eds., Vol. II, Springer, 147–175.
- Hoffman, R. N., Z. Liu, J. Louis, and C. Grassoti, 1995: Distortion representation of forecast errors. *Mon. Wea. Rev.*, **123**, 2758–2770, [https://doi.org/10.1175/1520-0493\(1995\)123<2758:DROFE>2.0.CO;2](https://doi.org/10.1175/1520-0493(1995)123<2758:DROFE>2.0.CO;2).
- Jazwinski, A. H., 1970: *Stochastic Processes and Filtering Theory*. Academic Press, 376 pp.
- Keper, J. D., 2009: Covariance localisation and balance in an ensemble Kalman filter. *Quart. J. Roy. Meteor. Soc.*, **135**, 1157–1176, <https://doi.org/10.1002/qj.443>.
- Kleist, D. T., and K. Ide, 2015: An OSSE-based evaluation of hybrid variational–ensemble data assimilation for the NCEP GFS. Part II: 4DEnVar and hybrid variants. *Mon. Wea. Rev.*, **143**, 452–470, <https://doi.org/10.1175/MWR-D-13-00350.1>.
- Knaff, J. A., S. P. Longmore, R. T. DeMaria, and D. A. Molenaar, 2015: Improved tropical cyclone flight-level wind estimates using routine infrared satellite reconnaissance. *J. Appl. Meteor. Climatol.*, **54**, 463–478, <https://doi.org/10.1175/JAMC-D-14-0112.1>.
- Kurosawa, K., and J. Poterjoy, 2021: Data assimilation challenges posed by nonlinear measurement operators: A comparative study using a simplified model. *Mon. Wea. Rev.*, **149**, 2369–2389, <https://doi.org/10.1175/MWR-D-20-0368.1>.
- Landsea, C. W., and J. L. Franklin, 2013: Atlantic hurricane database uncertainty and presentation of a new database format. *Mon. Wea. Rev.*, **141**, 3576–3592, <https://doi.org/10.1175/MWR-D-12-00254.1>.
- Lawson, W. G., and J. A. Hansen, 2005: Alignment error models and ensemble-based data assimilation. *Mon. Wea. Rev.*, **133**, 1687–1709, <https://doi.org/10.1175/MWR2945.1>.
- Lee, Y., and A. J. Majda, 2016: State estimation and prediction using clustered particle filters. *Proc. Natl. Acad. Sci. USA*, **113**, 14 609–14 614, <https://doi.org/10.1073/pnas.1617398113>.
- Leeuwen, V., H. R. Künsch, L. Nerger, R. Potthast, and S. Reich, 2019: Particle filters for high-dimensional geoscience applications: A review. *Quart. J. Roy. Meteor. Soc.*, **145**, 2335–2365, <https://doi.org/10.1002/qj.3551>.
- Lorenz, E. N., 1969: The predictability of a flow which possesses many scales of motion. *Tellus*, **21**, 289–307, <https://doi.org/10.3402/tellusa.v21i3.10086>.
- Lynch, P., and X.-Y. Huang, 1992: Initialization of the HIRLAM model using a digital filter. *Mon. Wea. Rev.*, **120**, 1019–1034, [https://doi.org/10.1175/1520-0493\(1992\)120<1019:IOTHMU>2.0.CO;2](https://doi.org/10.1175/1520-0493(1992)120<1019:IOTHMU>2.0.CO;2).
- Marks, F., N. Kurkowski, M. DeMaria, and M. Brennan, 2019: Hurricane forecast improvement program five-year plan: 2019–2024. NOAA, 86 pp., <https://hftp.org/sites/default/files/documents/hftp-strategic-plan-20190625-final.pdf>.

- Miyoshi, T., K. Kondo, and T. Imamura, 2014: The 10,240-member ensemble Kalman filtering with an intermediate AGCM. *Geophys. Res. Lett.*, **41**, 5264–5271, <https://doi.org/10.1002/2014GL060863>.
- Molinari, J., S. Skubis, D. Vollaro, F. Alsheimer, and H. E. Willoughby, 1998: Potential vorticity analysis of tropical cyclone intensification. *J. Atmos. Sci.*, **55**, 2632–2644, [https://doi.org/10.1175/1520-0469\(1998\)055<2632:PVAOTC>2.0.CO;2](https://doi.org/10.1175/1520-0469(1998)055<2632:PVAOTC>2.0.CO;2).
- Morzfeld, M., and D. Hodyss, 2019: Gaussian approximations in filters and smoothers for data assimilation. *Tellus*, **71A**, 1600344, <https://doi.org/10.1080/16000870.2019.1600344>.
- , —, and J. Poterjoy, 2018: Variational particle smoothers and their localization. *Quart. J. Roy. Meteor. Soc.*, **144**, 806–825, <https://doi.org/10.1002/qj.3256>.
- Nehrkorn, T., B. Woods, R. N. Hoffman, and T. Auligné, 2015: Correcting for position errors in variational data assimilation. *Mon. Wea. Rev.*, **143**, 1368–1381, <https://doi.org/10.1175/MWR-D-14-00127.1>.
- NWS, 2018: Average NHC Atlantic track forecast errors: 2010–2018. National Hurricane Center Forecast Verification, accessed 20 June 2019, <https://www.nhc.noaa.gov/verification/verify5.shtml>.
- Penny, S. G., and T. Miyoshi, 2016: A local particle filter for high dimensional geophysical systems. *Nonlinear Processes Geophys.*, **23**, 391–405, <https://doi.org/10.5194/npg-23-391-2016>.
- Posselt, D. J., and C. H. Bishop, 2012: Nonlinear parameter estimation: Comparison of an ensemble Kalman smoother with a Markov chain Monte Carlo algorithm. *Mon. Wea. Rev.*, **140**, 1957–1974, <https://doi.org/10.1175/MWR-D-11-00242.1>.
- Poterjoy, J., 2016: A localized particle filter for high-dimensional nonlinear systems. *Mon. Wea. Rev.*, **144**, 59–76, <https://doi.org/10.1175/MWR-D-15-0163.1>.
- , 2022: Regularization and tempering for a moment-matching localized particle filter. *Quart. J. Roy. Meteor. Soc.*, in press.
- , and F. Zhang, 2011: Dynamics and structure of forecast error covariance in the core of a developing hurricane. *J. Atmos. Sci.*, **68**, 1586–1606, <https://doi.org/10.1175/2011JAS3681.1>.
- , and —, 2014: Intercomparison and coupling of ensemble and variational data assimilation approaches for the analysis and forecasting of Hurricane Karl (2010). *Mon. Wea. Rev.*, **142**, 3347–3364, <https://doi.org/10.1175/MWR-D-13-00394.1>.
- , and J. L. Anderson, 2016: Efficient assimilation of simulated observations in a high-dimensional geophysical system using a localized particle filter. *Mon. Wea. Rev.*, **144**, 2007–2020, <https://doi.org/10.1175/MWR-D-15-0322.1>.
- , R. A. Sobash, and J. L. Anderson, 2017: Convective-scale data assimilation for the Weather Research and Forecasting Model using the local particle filter. *Mon. Wea. Rev.*, **145**, 1897–1918, <https://doi.org/10.1175/MWR-D-16-0298.1>.
- , L. Wicker, and M. Buehner, 2019: Progress toward the application of a localized particle filter for numerical weather prediction. *Mon. Wea. Rev.*, **147**, 1107–1126, <https://doi.org/10.1175/MWR-D-17-0344.1>.
- , G. Alaka, and H. Winterbottom, 2021: The irreplaceable utility of sequential data assimilation for model development: Lessons learned from an experimental HWRF system. *Weather Forecasting*, **36**, 661–677, <https://doi.org/10.1175/WAF-D-20-0204.1>.
- Potthast, R., A. Walter, and A. Rhodin, 2019: A localized adaptive particle filter within an operational NWP framework. *Mon. Wea. Rev.*, **147**, 345–362, <https://doi.org/10.1175/MWR-D-18-0028.1>.
- Ravela, S., K. Emanuel, and D. McLaughlin, 2007: Data assimilation by field alignment. *Physica D*, **230**, 127–145, <https://doi.org/10.1016/j.physd.2006.09.035>.
- Reich, S., and C. Cotter, Eds., 2015: Introduction to probability. *Probabilistic Forecasting and Bayesian Data Assimilation*. Cambridge University Press, 33–64.
- Robert, S., D. Leuenberger, and H. R. Kunsch, 2018: A local ensemble transform Kalman particle filter for convective-scale data assimilation. *Quart. J. Roy. Meteor. Soc.*, **144**, 1279–1296, <https://doi.org/10.1002/qj.3116>.
- Rotunno, R., and C. Snyder, 2008: A generalization of Lorenz's model for the predictability of flows with many scales of motion. *J. Atmos. Sci.*, **65**, 1063–1076, <https://doi.org/10.1175/2007JAS2449.1>.
- Slivinski, L., and C. Snyder, 2016: Exploring practical estimates of the ensemble size necessary for particle filters. *Mon. Wea. Rev.*, **144**, 861–875, <https://doi.org/10.1175/MWR-D-14-00303.1>.
- Snyder, C., T. Bengtsson, and M. Morzfeld, 2015: Performance bounds for particle filters using optimal proposal. *Mon. Wea. Rev.*, **143**, 4750–4761, <https://doi.org/10.1175/MWR-D-15-0144.1>.
- Stensrud, D. J., 2007: *Parameterization Schemes: Keys to Understanding Numerical Weather Prediction Models*. 1st ed. Cambridge University Press, 480 pp.
- Stratman, D. R., C. K. Potvin, and L. J. Wicker, 2018: Correcting storm displacement errors in ensemble using the Feature Alignment Technique (FAT). *Mon. Wea. Rev.*, **146**, 2125–2145, <https://doi.org/10.1175/MWR-D-17-0357.1>.
- Thepáut, J.-N., and P. Courtie, 1991: Four-dimensional variational data assimilation using the adjoint of a multilevel primitive-equation model. *Quart. J. Roy. Meteor. Soc.*, **117**, 1225–1254, <https://doi.org/10.1002/qj.49711750206>.
- Weng, Y., and F. Zhang, 2012: Assimilating airborne Doppler radar observations with an ensemble Kalman filter for convection-permitting hurricane initialization and prediction: Katrina (2005). *Mon. Wea. Rev.*, **140**, 841–859, <https://doi.org/10.1175/2011MWR3602.1>.
- Whitaker, J. S., and T. M. Hamill, 2002: Ensemble data assimilation without perturbed observations. *Mon. Wea. Rev.*, **130**, 1913–1924, [https://doi.org/10.1175/1520-0493\(2002\)130<1913:EDAWPO>2.0.CO;2](https://doi.org/10.1175/1520-0493(2002)130<1913:EDAWPO>2.0.CO;2).
- , and —, 2012: Evaluating methods to account for system errors in ensemble data assimilation. *Mon. Wea. Rev.*, **140**, 3078–3089, <https://doi.org/10.1175/MWR-D-11-00276.1>.
- Ying, Y., 2019: A multiscale alignment method for ensemble filtering with displacement errors. *Mon. Wea. Rev.*, **147**, 4553–4565, <https://doi.org/10.1175/MWR-D-19-0170.1>.
- Zhang, F., C. Snyder, and J. Sun, 2004: Impacts of initial estimate and observation availability on convective-scale data assimilation with an ensemble Kalman filter. *Mon. Wea. Rev.*, **132**, 1238–1253, [https://doi.org/10.1175/1520-0493\(2004\)132<1238:IOIEAO>2.0.CO;2](https://doi.org/10.1175/1520-0493(2004)132<1238:IOIEAO>2.0.CO;2).
- , M. Zhang, and J. A. Hansen, 2009: Coupling ensemble Kalman filter with four dimensional variational data assimilation. *Adv. Atmos. Sci.*, **26**, 1–8, <https://doi.org/10.1007/s00376-009-0001-8>.
- , Y. Weng, J. F. Gamache, and F. D. Marks, 2011: Performance of convection-permitting hurricane initialization and prediction during 2008–2010 with ensemble data assimilation of inner-core airborne Doppler radar observations. *Geophys. Res. Lett.*, **38**, L15810, <https://doi.org/10.1029/2011GL048469>.