

A modified particle filter that reduces weight collapse

Shay Gilpin*

Department of Mathematics, University of Arizona, Tucson, AZ, USA

Michael Herty[†]

RWTH Aachen University, Institute Geometry and

Practical Mathematics, 52056 Aachen, Germany and

Extraordinary Professor, Department of Mathematics and Applied Mathematics,

University of Pretoria, Private Bag X20, Hatfield 0028, South Africa

(Dated: October 29, 2025)

Abstract

Particle filters are a widely used Monte Carlo based data assimilation technique that estimates the probability distribution of a system's state conditioned on observations through a collection of weights and particles. A known problem for particle filters is weight collapse, or degeneracy, where a single weight attains a value of one while all others are close to zero, thereby collapsing the estimated distribution. We address this issue by introducing a novel modification to the particle filter that is simple to implement and inspired by energy-based diversity measures. Our approach adjusts particle weights to minimize a two-body energy potential, promoting balanced weight distributions and mitigating collapse. We demonstrate the performance of this modified particle filter in a series of numerical experiments with linear and nonlinear dynamical models, where we compare with the classical particle filter and ensemble Kalman filters in the nonlinear case. We find that our new approach improves weight distributions compared to the classical particle filter and thereby improve state estimates.

I. INTRODUCTION

Monte Carlo based data assimilation algorithms are a popular class of statistical estimation techniques that update an ensemble, or collection, of state variables with observations that are noisy and possibly sparse. One such method is the *particle filter* [see e.g., 1–7]. In general, particle filters update a posterior distribution conditioned on a set of observations, where this posterior distribution is estimated from a sequence of weights $\{w_i\}$ associated with an ensemble of particles $\{x_i\}$ using Bayes theorem. Many different forms have been proposed in recent years [see 6, 8–10, and references therein].

While particle filters are simple to implement, their application is often limited to low-dimensional systems due to the phenomena known as weight collapse, or degeneracy [e.g., 11–15]. Weight collapse occurs when a single weights attains a value close to one, while all other weights are nearly zero, effectively concentrating all information onto one particle and rendering the posterior distribution approximated by these weights useless. One way to avoid this collapse is to increase the number of particles, however it has been shown that

* sgilpin@arizona.edu

† herty@igpm.rwth-aachen.de

the number of particles must grow exponentially with the dimension of the observations, therefore limiting its applicability in high-dimensional systems [16]. Several alternative approaches have been proposed to counteract weight collapse. For example, a suitable number of particles can be chosen by estimating the effective dimension of the system, [17]. Other approaches include localization, which limits the impact of observation information and thereby the increase in particle weights, [15]. We refer the reader to [7] for an overview of suitable alternatives.

Motivated by recent progress in the introduction of energy-based diversity measures in multi-objective optimization [18], we propose a new modification to the particle filter that employs a similar concept to adjust the posterior weights. In this formulation, the particle filter weights are adjusted in such a way that they correspond to a minimal configuration of a suitable two-body energy potential. We are able to obtain weight distributions by balancing their value within a potential landscape. This technique has been successfully applied in the context of particle swarm optimization methods for multi-objective minimization [see e.g., 19, 20]. This approach yields a simple modification to the classical particle filter that, and as we demonstrate in a series of numerical experiments, improves weight distributions and reduces the frequency of collapse relative to the classical particle filter. This work provides a proof of concept for this modified particle filter through a series of numerical experiments, which will be followed by analysis in future work.

This paper is organized as follows. We begin with the standard formulation of the particle filter in Sec. II A, followed by the modification of the weights in Sec II B. In Sec II C provides a motivating example to illustrate the modification’s impact on weight computations relative to the classical particle filter. We conclude this section with a discussion of the ensemble Kalman filter in Sec. II D, which is another class of Monte Carlo estimation techniques we will compare with in our numerical experiments. We present two sets of numerical experiments, the first for a linear dynamical system with a known solution in Sec. III, and the second for the nonlinear Lorenz ’63 dynamical system in Sec. IV. This is followed by a summary and discussion in Sec. V.

II. MONTE CARLO BASED DATA ASSIMILATION: PARTICLE FILTERS, THE MODIFIED PARTICLE FILTER, AND ENSEMBLE KALMAN FILTERS

In this section, we introduce the standard formulation of the particle filter, which we refer to as the classical particle filter, and our modification using potential functions. This is followed by a motivating example where we illustrate that the introduced potential actually allows to shift the weight distribution of the filter. We conclude this section with a brief discussion of ensemble Kalman filters, which are an alternative class of Monte Carlo based data assimilation algorithm we will use in our nonlinear numerical experiments.

A. The classical particle filter

To introduce the standard formulation of the particle filter, we follow the presentation in [7]. Consider an ensemble of N_e model states $x_i \in \mathbb{R}^{N_x}$, for $i = 1, \dots, N_e$ called *particles*. These particles represent the empirical measure of the prior probability density function $p(x)$,

$$p(x) = \frac{1}{N_e} \sum_{j=1}^{N_e} \delta(x - x_j). \quad (1)$$

Between observations, the particles are propagated (forecasted) using deterministic model $f : \mathbb{R}^{N_x} \rightarrow \mathbb{R}^{N_x}$ given by

$$x_i^f = f(x_i), \quad (2)$$

for all i , which may be nonlinear. Random forcing can be added to the dynamics in (2). The observation $y \in \mathbb{R}^{N_y}$ is given by

$$y = H(x_{true}) + \epsilon, \quad (3)$$

where x_{true} is the true state of the system, the operator $H : \mathbb{R}^{N_x} \rightarrow \mathbb{R}^{N_y}$, and ϵ is measurement, or observation, error. It is often assumed that ϵ is normally distributed with zero mean and $N_y \times N_y$ observation error covariance matrix $\mathbf{R}$, which is typically a scalar multiple of the identity matrix.

The observation y and the predicted state of the system x^f are assimilated using a likelihood function, i.e., the probability density $p(y|x)$ of an observation y given a possible model

state x . The posterior probability distribution $p(x^f|y)$ is obtained using Bayes' theorem

$$p(x^f|y) = \frac{p(y|x^f)p(x^f)}{p(y)}. \quad (4)$$

Using the representation of the prior probability for the particles x_i^f for $i = 1, \dots, N_e$, by (1), we can approximate the posterior distribution with the ansatz as

$$p(x^f|y) \approx \sum_{j=1}^N w_j \delta(x_j^f - x^f), \quad \sum_{j=1}^{N_e} w_j = 1. \quad (5)$$

The unknown weights w_i are obtained using Bayes' theorem (4) as

$$w_i = \frac{p(y|x_i^f)}{\sum_{j=1}^{N_e} p(y|x_j^f)}. \quad (6)$$

Note that all quantities in (6) are known, that is, the position of the particles x_i^f as well as the observation y , leading to an explicit computation of the weights w_i . This leads to an approximation of the posteriori distribution given by a weighted empirical measures (5).

The dynamical evolution of the particles in (2) and the assimilation of observation information in (5) and (6) define the two steps of a *data assimilation cycle* with the classical particle filter [7]. Repeating this cycle to assimilate more observations will result in a skewed distribution of weights w_i , ultimately causing one particle to have a nonzero weight while the others are zero or close to zero [2]. Several different algorithms have been designed to combat this degeneracy ranging in complexity, such as resampling methods, proposal densities, and others [e.g., 2, 5, 7, and reference therein]. In the next section, we propose an alternative method to prevent such weight collapse inspired by energy diversity measures.

B. Modified particle filters

As outlined in the introduction, we propose a simple procedure to update the weight distribution based on recent considerations for multi-objective minimization [19]. The idea is to slightly modify the weights in (6) to guarantee an equi-distribution of the weights. In order to measure the clustering of the weights, we introduce a diversity measure on the weight distribution. Consider the probability density $\varrho^{N_e} \in \mathbb{P}(\mathbb{R})$ of the weights $\{w_i\}$ given by the empirical measure $\varrho^{N_e}(w) = \frac{1}{N_e} \sum_{j=1}^{N_e} \delta(w - w_j)$. In general, the diversity $\mathcal{U}$ of a

probability measure $\varrho \in \mathcal{P}(\mathbb{R})$ is defined by a two-body potential

$$\mathcal{U}(\varrho) = \int \int U(x - y) \varrho(dy) \varrho(dx), \quad (7)$$

where $U : \mathbb{R} \rightarrow \mathbb{R}$ is, for example, the Morse potential $U(z) = \exp(-C\|z\|)$ for some constant $C > 0$. Other potentials, such as a Newtonian potential, are possible [see e.g., 19, 21, 22]. Obtaining an equi-distribution then amounts to obtaining a configuration that is minimal with respect to the diversity measure $\mathcal{U}$. Solving the minimization problem $\min_{\varrho} \mathcal{U}(\varrho)$ on the space of empirical measures of size N_e is, however, non-trivial and computationally expensive. Therefore, we follow a heuristic strategy similar to [19]: We consider a parameterized family $t \rightarrow \varrho(t)$ of probability measures $\varrho(t) \in \mathcal{P}(\mathbb{R})$, such that

$$\frac{d}{dt} \mathcal{U}(\varrho(t)) \leq 0, \quad \varrho(0) = \varrho^{N_e}, \quad (8)$$

and evolve these over a fixed time period. The decay of $\mathcal{U}$ can be achieved if ϱ is a gradient flow with respect to $\mathcal{U}$ and fulfills weakly

$$\partial_t \varrho + \partial_x \int U'(x - y) \varrho(t, dy) \varrho(t, dx) = 0, \quad \varrho(0) = \varrho^{N_e}. \quad (9)$$

Note that the time scale t here has no physical meaning. We, therefore, consider an explicit Euler discretization in time with a single(!) time step $\Delta t > 0$. Furthermore, ϱ^{N_e} is an empirical measure and therefore the evolution of ϱ can be simply reformulated in an update of the initial weights w_i . More precisely, we have

$$w_i(\Delta t) = w_i - \frac{\Delta t}{N_e} \sum_{j=1}^{N_e} U'(w_j - w_i). \quad (10)$$

Denoting the arbitrary time step $\Delta t = \alpha > 0$ as our algorithmic parameter, we propose to obtain the posterior distribution (5) as

$$p(x^f | y) = \sum_{j=1}^N w_j(\alpha) \delta(x_j^f - x^f), \quad (11)$$

where

$$w_i(\alpha) = w_i - \frac{\alpha}{N_e} \sum_{j=1}^{N_e} U'(w_j - w_i) \quad (12)$$

and the weights w_i are given by (6). If necessary, the weights $w_i(\alpha)$ are projected on $[0, 1]$. By the previous considerations, our proposed method balances the weight computation

according to Bayes' theorem with a gradient descent step on the diversity measure of the weight distribution. The balancing is controlled by the parameter $\alpha > 0$. In the case where $\alpha = 0$, we recover (6).

As an illustration, we show the Morse potential on an equi-distribution of $N_e = 10^3$ weights as well as for a normal distribution of weights close to one. The histogram of the weights are given in blue and red, respectively. The corresponding values of the Morse potential $U(z) = \exp(-\frac{1}{2}z)$ are also indicated in Figure 1. It is clearly visible that the clustering leads to higher potential values, motivating a gradient descent approach for the potential as modification for the filter.




FIG. 1. Morse potential for an equi-distribution of weights (blue) and a distribution clustered at weight one (red). Histograms of the weight distributions are also reported.

C. Motivating example

To illustrate the impact of the added potential to the weight computation in the particle filter, we consider the example suggested in Sec. 3 of [11]. For convenience of the reader, we repeat the setup of the problem using the notation of [11]. We consider a state $x \in \mathbb{R}^{N_x}$ for $N_x \in \{10, 30, 100\}$. The observation model is given by $y = x + \epsilon$ where ϵ and x are both normally distributed with zero mean and unit variance. The ensemble size is set to $N_e = 10^3$. For a classical particle filter, the weight of the posterior of the ensemble member i is given by Equation (3) of [11] or our (5), respectively. In this example, each x_i^f is sampled from a Gaussian distribution with unit variance. The likelihood is a multivariate Gaussian

distribution with mean $y - x$ and unit variance as in the example in Sec. 3 of [11]. Reported is the histogram of $\max_{i=1,\dots,N_e} w_i$ in Figure 2 (blue part of the histogram). As in [11] we observe the clustering of weights at weight $w \approx 1$ that becomes particularly pronounced as the dimension N_x increases, noting that N_e remains fixed.

For the modified particle filter in (11) and (12), the weights of the posterior are computed with the weighted potential $U'(x) = (1 - e^{-|x|})^2 - 1$, i.e., for $i = 1, \dots, N_e$ we consider

$$w_i = \frac{p(y|x_i^f)}{\sum_{j=1}^{N_e} p(y|x_j^f)}, \quad w_i(\alpha) = w_i - \alpha \frac{1}{N_e} \sum_{j=1}^{N_e} U'(w_j - w_i). \quad (13)$$

The weight $\alpha > 0$ is fixed in all subsequent computations to $\alpha = \frac{1}{2}$. The potential used is the Morse potential, i.e.,

$$U(z) = -\frac{1}{2} \exp\left(-\frac{1}{2}|z|\right). \quad (14)$$

Again, for different dimensions N_x , a histogram of the maximum weights is reported in Figure 2 (orange). As expected, the additional forcing due to the potential U leads to a more uniform weight distribution, and helps to offset the weight collapse at weight $w \approx 1$ observed in the classical particle filter as the dimension N_x increases.




FIG. 2. Histogram of the maximal weights as in the example of Sec. 3 of [11]. Computation of the weights according to Bayes' formula (5) (blue) and using the proposed Morse potential (13) (red). The state space dimension is $N_x = 10, 30$, and $N_x = 100$ as in [11]. The sample size is $N_e = 10^3$ and the weight parameter $\alpha = \frac{1}{2}$.

D. Ensemble Kalman filters

To conclude this section, we introduce another class of Monte Carlo based data assimilation algorithms known as ensemble Kalman filters [23–25], which we will compare with in our numerical experiment with a nonlinear dynamical model. Ensemble Kalman filters (EnKFs) are an extension of the Kalman filter formalism introduced by [26], where Kalman filters apply Bayes’ theorem to assimilate observations assuming that (i) the prior distribution and likelihood are both Gaussian, and (ii) both the observation operator H in (3) and the model dynamics f in (2) are linear. The posterior distribution is therefore a Gaussian distribution whose mean and covariance can be computed explicitly [26]. Ensemble Kalman filters extend this formalism by allowing for nonlinear observation operators, nonlinear model dynamics, and by approximating covariances using Monte Carlo techniques. This approach is done by updating an ensemble of states using the same Kalman filter formalism. For an ensemble of states $x_i^f \in \mathbb{R}^{N_x}$ for $i = 1, 2, \dots, N_e$ evolved (forecasted) by (2), the ensemble is updated to generate a new ensemble $x_i^a \in \mathbb{R}^{N_x}$ for $i = 1, 2, \dots, N_e$ given an observation y defined by (3),

$$x_i^a = x_i^f + \mathbf{K} \left(y - (\mathbf{H}x_i^f + \eta_i) \right), \quad i = 1, 2, \dots, N_e, \quad (15a)$$

$$\mathbf{K} = \mathbf{P}^f \mathbf{H}^T (\mathbf{H} \mathbf{P}^f \mathbf{H}^T + \mathbf{R})^{-1}. \quad (15b)$$

The matrix $\mathbf{K}$ is referred to as the Kalman gain. The $N_x \times N_y$ matrix $\mathbf{H}$ is a linearization of the observation operator H in (3), the η_i are independent samples from a Gaussian distribution with mean zero and covariance $\mathbf{R}$, and the $N_x \times N_x$ matrix $\mathbf{P}^f$ is the (forecast) error covariance matrix estimated from the ensemble $\{x_i^f\}$ using the empirical sample covariance estimator,

$$\mathbf{P}^f = \frac{1}{N_e - 1} \sum_{i=1}^{N_e} (x_i^f - \bar{x}^f)(x_i^f - \bar{x}^f)^T, \quad \bar{x}^f = \frac{1}{N_e} \sum_{i=1}^{N_e} x_i^f. \quad (16)$$

Equation 15a defines the “stochastic” EnKF [24], where the term “stochastic” arises from generating an ensemble of observations through the addition of η_i , which ensures the correct error statistics [24]. Variants of the EnKF include the ensemble adjustment Kalman filter [27] and ensemble transform filters [28], for example.

Unlike particle filters, the EnKF assumes that the prior distribution and likelihood are both Gaussian in order to apply the Kalman filter formalism to update the forecasted particles in (15a). Ensemble Kalman filters, therefore rely on the accuracy of the mean state and

covariance estimates from the empirical estimators in (16). The empirical sample covariance in (16) is unbiased and consistent, however produces large errors when $N_e \ll N_x$ [29]. In these contexts, additional covariance estimation techniques are employed to mitigate these errors due to limited ensemble size, for example by reducing spurious correlations through tapering [25, 30, Ch. 10 and references therein]. Since we are interested in comparing the EnKF with the particle filter, we will consider the case where $N_e \gg N_x$ so that we can compute (16) directly with sufficient accuracy.

III. NUMERICAL EXPERIMENTS FOR LINEAR DYNAMICS

To assess the performance of the modified particle filter in a cycling data assimilation context, we first test the method for a linear dynamical model in which the true state is known. In this example, we consider a simple time-discrete linear model with parameter $\lambda > 0$ and time step $\Delta t > 0$. The evolution of the true state is given by

$$x_{true}^n = x^0 \exp(-\lambda \Delta t n). \quad (17)$$

for an initial data $x_{true}^0 \in \mathbb{R}$. The time horizon is $n\Delta t = 2$. The observation model is

$$y = x_{true} + \sigma \epsilon \quad (18)$$

where ϵ is the normally distributed noise with unit variance and the positive parameter $\sigma > 0$. Observations are recorded at each time $n = 0, 1, \dots$. The dynamics for the state estimate (2) is defined by a time-discrete forward model with states x_n^f at time $n\Delta t$. As forward model we use an explicit Euler discretization, i.e.,

$$x^{f,n+1} = x^{f,n} - \lambda \Delta t x^{f,n}. \quad (19)$$

The same time step as above is used and the initialization $x^{f,0}$ is distributed according to a prior distribution $p(x)$ specified below.

We compare the performance of the classical particle filter and modified particle filter for estimating the current state of the system at time $n\Delta t$. The initial data $x^{f,0}$ is distributed according to p , which is taken to be a normal distribution with mean $\bar{x}^{f,0}$ and variance σ^f . An iterative procedure is applied to identify the true state x_{true}^n . At each time step, we apply Bayes' theorem and compute the weights w_i^n using equation (6) or equation (12),

respectively. The posteriori distribution $p(x^{f,n}|y)$ is given by equation (5) or equation (11), respectively. The likelihood is a multivariate Gaussian distribution with mean $y - x^{f,n}$ and unit variance. For the next time step $n + 1$, this posteriori is re-sampled to obtain a new prior of the type given by equation (1). At each point in time we record the mean and variance of the posterior $p(x^{f,n}|y)$ and we compare this numerically to the true state x_{true}^n .

We expect that the variance of the particle filter decays for increasing number n as well as convergence of the mean of the posterior towards the true state. Since in the modified version of the particle filter (11) and (12) we modify the weights, we expect a slower decay of the variance for the modified algorithm. In both cases, we observe the expected decay towards the true mean in Figure 3. The modification applied to the computation of the weights (12) is expected to introduce an additional variance to the sample. The example shows that this additional variance does not deter the overall properties of the particle filter, namely, the convergence to the true state. In the right part of Figure 3 it is shown that even so the variance is larger on the first time steps, it decays as in the classical particle filter (left part).

The following parameters are used in numerical simulations: $\Delta t = \frac{1}{10}$, $\lambda = \frac{1}{2}$, $\sigma = \frac{1}{10}$, $x_{true}^0 = 10$, $\bar{x}^{f,0} = 8$ and $\sigma^f = 4$. We use $N_e = 10^3$ particles and a weight $\alpha = 10^{-3}$. The potential is the same as in the first example given by equation (14). The number of time steps for the simulation is $T/\Delta t$.

IV. NUMERICAL EXPERIMENTS WITH LORENZ '63

We extend the linear example of the previous section to a nonlinear dynamical system. In this example, we compare the classical particle filter with the particle filter with modified weights and the ensemble Kalman filter (EnKF) in a series of cycling data assimilation experiments with the Lorenz '63 dynamical model [31]. The Lorenz '63 dynamical model is system of three nonlinear, coupled ordinary differential equations for scalar variables $x(t)$, $y(t)$ and $z(t)$,

$$\begin{aligned}x' &= \sigma(y - x), \\y' &= \rho x - xz - y, \\z' &= xy - bz.\end{aligned}\tag{20}$$




FIG. 3. Particle filter for the linear model (17), with forecast model (19). Noise is added to the observation according to (18). In red the true solution is depicted. In blue, we present the mean of the ensemble where we add and subtract the variance of the ensemble. In part a) of the figure the classical particle filter defined through the weights (6) is shown. In part b) of the figure the modified particle filter defined through (12).

We fix the parameters $\sigma = 10, \rho = 28, b = 8/3$, and define the transpose of the state vector $\mathbf{x}^T = [x \ y \ z]^T$.

The set up for these experiments is similar to that described in Sec. 4b of [32]. We define our reference solution $\mathbf{x}_{\text{ref}}$ as the solution to (20) approximated by a fourth-order explicit Runge-Kutta scheme with time step $\Delta t = 0.01$ and initial condition

$$\mathbf{x}_{\text{ref}}^0 = \begin{bmatrix} 1.508870 \\ -1.531271 \\ 25.46091 \end{bmatrix}. \quad (21)$$

We generate the initial ensemble of states $\mathbf{x}_i^0, i = 1, 2, \dots, N_e$ from the reference initial condition $\mathbf{x}_{\text{ref}}^0$ by adding independent and identically distributed (IID) noise that is Gaussian distributed with mean zero and variance two. The observations y are generated from reference solution using (3) with an observation operator $H = \mathbf{H} = [1 \ 0 \ 0]$ and ϵ are IID Gaussian with mean zero and observation error variance $r = 1$. Observations are assimilated every 10 integration time steps (i.e., every $10\Delta t$). Assimilation experiments are run for a total of 100 cycles, starting each cycle with a filter update, followed by a forecasting step. For the forecasting step, we evolve $i = 1, 2, \dots, N_e$ particles according to

$$\mathbf{x}_i^{k+1} = f(\mathbf{x}_i^k) + \boldsymbol{\omega}^k, \quad (22)$$

where $f(\cdot)$ is the same explicit Runge-Kutta scheme used to generate the reference solution with index k denoting one integration time step Δt . The vector $\boldsymbol{\omega}$ is a stochastic forcing term which we define to be

$$\boldsymbol{\omega} := \sigma \sqrt{\Delta t} \begin{bmatrix} \sqrt{2} \\ \sqrt{12.13} \\ \sqrt{12.31} \end{bmatrix}, \quad (23)$$

where σ are independent draws from a standard normal distribution. This corresponds to adding uncorrelated model errors with a diagonal model error matrix $\mathbf{Q} = \text{diag}(2, 12.13, 12.31)$.

In these experiments, we compare the state estimates after a filter update, which we refer to as the analysis state, from the classical particle filter, modified particle filter, and EnKF. We consider two ensemble sizes of $N_e = 50$ and $N_e = 100$. For both particle filters we apply (uniform) sequential importance resampling to generate the analysis particles [33, p. 1993]. Therefore, the only difference between the two particle filters is in the calculation of the weights. The stepsize $\alpha > 0$ in the modified particle filter is chosen via a line search, where for a fixed number of particles N_e , we choose the optimal α that minimizes the average root mean square (RMS) error in the analysis state over a series of 500 data assimilation experiments.

Figures 4 and 5 summarize the results for 1000 cycling experiments with the Lorenz '63 model for $N_e = 50$ and $N_e = 100$, respectfully. Panel (a) of Figs. 4 and 5 are histograms of the average state analysis root mean square (RMS) error for all these experiments, where in each experiment the first 30 cycles are removed from the averaging for “spin-up.” We see that the modified particle filter generally reduces the state analysis RMS error, in particular reducing the frequency of experiments with large RMS error. This is reflected in the mean, median, and standard deviation of these averaged state analysis errors in Table I. We use the average analysis RMS error from the EnKF as a reference for comparison with the two particle filters. Even though the statistics over the 1000 experiments indicate larger errors in both particle filters relative to the EnKF, individual experiments have state analysis RMS that are comparable to that of the EnKF. The modified particle filter overall improves the mean state analysis error relative to the classical particle filter, and reduces the standard deviation in these errors, most notably in the $N_e = 100$ case. While the median state RMS error for the modified particle filter is larger than the classical particle filter in the $N_e = 100$ case, the magnitude of the large errors is reduced, as reflected in the lower standard deviation

and histogram in Fig. 5.

The improvements in the state estimates from the modified particle filter can be related to the improved distributions of weights. Panel (b) of Figs. 4 and 5 plots histograms of the maximum weights w_i for each of the 100 cycles and over all 1000 experiments. As observed in the motivating example, the modified particle filter shifts the distribution of maximum weights towards smaller values, reducing the instances of weight collapse to values of one. This is also reflected in the weight distributions for a single experiment in Fig. 6, in which the modified particle filter has slightly smaller state analysis RMS errors than the classical particle filter.




FIG. 4. Panel (a): Average state analysis root mean square (RMS) errors over 1000 cycling experiments with the Lorenz '63 model and $N_e = 50$. Panel (b): Histogram of the maximum weight w_i of each cycle for the 1000 cycling experiments. Orange corresponds to the classical particle filter, and blue the modified particle filter. Vertical scales in both panels are logarithmic scales.

In addition to the distributions of the maximum weights, we also compute two additional quantities for comparison. The first is the variance in the weights, τ^2 ,

$$\tau^2 = \text{var} \log(w), \quad (24)$$




FIG. 5. Same as Fig. 4 for $N_e = 100$.

TABLE I. State analysis RMS Error over 1000 experiments for the classical particle filter (PF), modified particle filter (PF), and ensemble Kalman filter (EnKF).

	$N_e = 50$			$N_e = 100$		
	PF	mPF	EnKF	PF	mPF	EnKF
Mean	3.077	2.375	0.642	1.584	1.030	0.635
Median	1.066	0.939	0.620	0.699	0.782	0.626
Standard Deviation	3.393	2.931	0.104	2.431	1.170	0.102

derived in Sec. 4a of [11]. Larger values of τ correspond to higher variance in the weights and can be indicative of collapse. Table II summarizes the mean, median, and standard deviation of τ for the classical particle filter, and modified particle filter over the 1000 experiments with $N_e = 50$ and $N_e = 100$. For both ensemble sizes, the modified particle filter reduces the variances in the computed weights and significantly lowers the standard deviation. As a consequence, the modified particle filter decreases the frequency of experiments with large state analysis errors. We also see that the modified particle filter reduces the effective




FIG. 6. Maximum weights w_i for the classical particle filter (orange) and modified particle filter (blue) over a single cycling experiment of 100 cycles. Panel (a) corresponds to $N_e = 50$ and panel (b) $N_e = 100$. The average state analysis RMS errors (removing the first 30 cycles) for the classical particle filter is 0.506 for $N_e = 50$ and 0.608 for $N_e = 100$; for the modified particle filter 0.424 for $N_e = 50$ and 0.568 for $N_e = 100$.

TABLE II. Variance in the computed weights (τ) over 1000 experiments for the classical particle filter (PF) and modified particle filter (mPF).

	$N_e = 50$		$N_e = 100$	
	PF	mPF	Pf	mPF
Mean	5.558	0.316	5.141	0.164
Median	0.837	0.252	0.596	0.123
Standard Deviation	19.39	0.246	22.26	0.130

ensemble size, N_{eff} ,

$$N_{\text{eff}} = \left(\sum_{i=1}^{N_e} w_i^2 \right)^{-1}, \quad (25)$$

[25, their (9.5)], given in Table III. Small effective ensemble sizes N_{eff} indicate collapse and the need for resampling. The modified particle filter increases the effective ensemble size relative to the classical particle filter, particularly in the $N_e = 100$ case.

TABLE III. Effective ensemble size (N_{eff}) for the 1000 experiments for the classical particle filter (PF) and modified particle filter (mPF).

	$N_e = 50$		$N_e = 100$	
	PF	mPF	PF	mPF
Mean	31	36	69	82
Median	36	41	79	90
Standard Deviation	16	14	28	20

V. SUMMARY AND DISCUSSION

Particle filters are a class of Monte Carlo data assimilation techniques that are known to suffer from weight collapse, or weight degeneracy, where the weight of one of the particles concentrates at one while all others become very close to zero. We introduce a small modification to the particle filter, inspired by recent developments in energy-based diversity measures, to mitigate this collapse. The proposed modification introduces a potential on the weight distribution of the filter and adjusts the weight values based on a minimization of the corresponding potential. Through a series of numerical experiments with linear and nonlinear dynamics, we compare the performance of this new, modified particle filter, with the classical particle filter, and with the ensemble Kalman filter in the nonlinear case. We find that this modification helps to reduce weight collapse by improving weight distributions while still yielding the correct Bayes's posterior distribution. In the numerical experiments with the nonlinear Lorenz '63 model, we find both improved state estimates and reduced frequency of data assimilation experiments with significant errors when including the potential function in the weight computations.

The findings in our numerical experiments suggest that modifying the weight computation with a potential function can result in improved particle filter performance for a moderately small number of particles. In future work we plan to investigate the numerical impact of different potential functions, like e.g. the Riesz potential, as well as analytical considerations on the large particle limit of the proposed method.

Acknowledgments MH is funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) - SPP 2183: Eigenschaftsgeregelte Umformprozesse with the Project(s)

HE5386/19-3 Entwicklung eines flexiblen isothermen Reckschmiedeprozesses für die eigenschafts-geregelte Herstellung von Turbinenschaufeln aus Hochtemperaturwerkstoffen (424334423) and through HE5386/33-1 Control of Interacting Particle Systems, and Their Mean-Field, and Fluid-Dynamic Limits (560288187), HE5386/34-1 Partikelmethoden für unendlich dimensionale Optimierung (561130572) and through 442047500/SFB1481 within the projects B04 (Sparsity fördernde Muster in kinetischen Hierarchien), B05 (Sparsifizierung zeitabhängiger Netzwerkflußprobleme mittels diskreter Optimierung) and B06 (Kinetische Theorie trifft algebraische Systemtheorie) is acknowledged. SG is supported in part by the Data Driven Discovery RTG at the University of Arizona under NSF Grant No. DMS-1937229.

Author Contributions S.G. and M.H. contributed equally to this work.

-
- [1] N. J. Gordon, D. J. Salmond, and A. F. Smith, Novel approach to nonlinear/non-gaussian bayesian state estimation, in *IEEE proceedings F (radar and signal processing)*, Vol. 140 (IEE, 1993) pp. 107–113.
 - [2] A. Doucet, N. De Freitas, and N. Gordon, An introduction to sequential monte carlo methods, in *Sequential Monte Carlo methods in practice* (Springer, 2001).
 - [3] P. Fearnhead and H. R. Künsch, Particle filters and data assimilation, *Annual Review of Statistics and Its Application* **5**, 421 (2018).
 - [4] S. Vetra-Carvalho, P. J. Van Leeuwen, L. Nerger, A. Barth, M. U. Altaf, P. Brasseur, P. Kirchgessner, and J.-M. Beckers, State-of-the-art stochastic data assimilation methods for high-dimensional non-gaussian problems, *Tellus A: Dynamic Meteorology and Oceanography* **70**, 1 (2018).
 - [5] P. J. Van Leeuwen, Particle filtering in geophysical systems, *Monthly Weather Review* **137**, 4089 (2009).
 - [6] P. J. Van Leeuwen, Nonlinear data assimilation for high-dimensional systems, *Nonlinear data assimilation* **778**, 1 (2015).
 - [7] P. J. Van Leeuwen, H. R. Künsch, L. Nerger, R. Potthast, and S. Reich, Particle filters for high-dimensional geoscience applications: A review, *Quarterly Journal of the Royal Meteorological Society* **145**, 2335 (2019).
 - [8] V. Zaritskii, V. B. Svetnik, and L. Šimelevič, Monte-carlo technique in problems of optimal

- information processing, *Avtomatika i telemekhanika* , 95 (1975).
- [9] M. Morzfeld, X. Tu, E. Atkins, and A. J. Chorin, A random map implementation of implicit filters, *Journal of Computational Physics* **231**, 2049 (2012).
 - [10] E. Vanden-Eijnden and J. Weare, Data assimilation in the low noise regime with application to the kuroshio, *Monthly Weather Review* **141**, 1822 (2013).
 - [11] C. Snyder, T. Bengtsson, P. Bickel, and J. Anderson, Obstacles to high-dimensional particle filtering, *Monthly Weather Review* **136**, 4629 (2008).
 - [12] C. Snyder, T. Bengtsson, and M. Morzfeld, Performance bounds for particle filters using the optimal proposal, *Monthly Weather Review* **143**, 4750 (2015).
 - [13] B. Li, T. Bengtsson, and P. Bickel, Curse-of-dimensionality revisited: Collapse of importance sampling in very large scale systems, *Rapport technique* **85**, 205 (2005).
 - [14] P. Bickel, B. Li, and T. Bengtsson, Sharp failure rates for the bootstrap particle filter in high dimensions, in *Pushing the limits of contemporary statistics: Contributions in honor of Jayanta K. Ghosh*, Vol. 3 (Institute of Mathematical Statistics, 2008) pp. 318–330.
 - [15] M. Morzfeld, D. Hodyss, and C. Snyder, What the collapse of the ensemble kalman filter tells us about particle filters, *Tellus A: Dynamic Meteorology and Oceanography* **69**, 1283809 (2017).
 - [16] T. Bengtsson, P. Bickel, and B. Li, Curse-of-dimensionality revisited: Collapse of the particle filter in very large scale systems, in *Probability and statistics: Essays in honor of David A. Freedman*, Vol. 2 (Institute of Mathematical Statistics, 2008) pp. 316–335.
 - [17] A. J. Chorin and M. Morzfeld, Conditions for successful data assimilation, *Journal of Geophysical Research: Atmospheres* **118**, 11 (2013).
 - [18] C. A. Coello Coello, S. Gonzalez Brambila, J. Figueroa Gamboa, M. G. Castillo Tapia, and R. Hernandez Gomez, Evolutionary multiobjective optimization: open research areas and some challenges lying ahead, *Complex & Intelligent Systems* **6**, 221 (2020).
 - [19] G. Borghi, M. Herty, and L. Pareschi, An adaptive consensus based method for multi-objective optimization with uniform pareto front approximation, *Applied Mathematics & Optimization* **88**, 58 (2023).
 - [20] G. Borghi, M. Herty, and L. Pareschi, A consensus-based algorithm for multi-objective optimization and its mean-field description, in *2022 IEEE 61st Conference on Decision and Control (CDC)* (IEEE, 2022) pp. 4131–4136.

- [21] C. M. Fonseca, X. Gandibleux, J.-K. Hao, and M. Sevaux, *Evolutionary multi-criterion optimization* (Springer, 2009).
- [22] M. A. Braun, P. K. Shukla, and H. Schmeck, Obtaining optimal pareto front approximations using scalarized preference information, in *Proceedings of the 2015 Annual Conference on Genetic and Evolutionary Computation* (2015) pp. 631–638.
- [23] G. Evensen, Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics, *J. Geophys. Res.* **99**, 10143 (1994).
- [24] G. Burgers, P. van Leeuwen, and G. Evensen, Analysis scheme in the ensemble kalman filter, *Monthly Weather Review* **126**, 1719 (1998).
- [25] G. Evensen, F. C. Vossepoel, and P. J. Van Leeuwen, *Data assimilation fundamentals: A unified formulation of the state and parameter estimation problem* (Springer Nature, 2022).
- [26] R. Kalman, A new approach to linear filtering and prediction problems, *Journal of Basic Engineering* **82**, 35 (1960).
- [27] J. Anderson, An ensemble adjustment Kalman filter for data assimilation, *Monthly Weather Review* **129**, 2884 (2001).
- [28] M. Tippett, J. Anderson, C. Bishop, T. Hamill, and J. Whitaker, Ensemble square root filters, *Mon. Weather Rev.* **131**, 1485 (2003).
- [29] M. J. Wainwright, *High-Dimensional Statistics: A Non-Asymptotic Viewpoint*, Cambridge Series in Statistical and Probabilistic Mathematics (Cambridge University Press, 2019).
- [30] P. L. Houtekamer and H. L. Mitchell, A sequential ensemble Kalman filter for atmospheric data assimilation., *Mon. Weather Rev.* **128**, 123 (2001).
- [31] E. Lorenz, Deterministic nonperiodic flow, *Journal of the Atmospheric Sciences* **20**, 130 (1963).
- [32] G. Evensen and P. J. Van Leeuwen, An ensemble kalman smoother for nonlinear dynamics, *Monthly Weather Review* **128**, 1852 (2000).
- [33] P. J. Van Leeuwen, Nonlinear data assimilation in geosciences: an extremely efficient particle filter, *Quarterly Journal of the Royal Meteorological Society* **136**, 1991 (2010).