

Resampling-free Particle Filters in High-dimensions

Akhilan Boopathy¹, Aneesh Muppidi², Peggy Yang¹, Abhiram Iyer¹, William Yue¹, Ila Fiete¹

Abstract—State estimation is crucial for the performance and safety of numerous robotic applications. Among the suite of estimation techniques, particle filters have been identified as a powerful solution due to their non-parametric nature. Yet, in high-dimensional state spaces, these filters face challenges such as ‘particle deprivation’ which hinders accurate representation of the true posterior distribution. This paper introduces a novel resampling-free particle filter designed to mitigate particle deprivation by forgoing the traditional resampling step. This ensures a broader and more diverse particle set, especially vital in high-dimensional scenarios. Theoretically, our proposed filter is shown to offer a near-accurate representation of the desired posterior distribution in high-dimensional contexts. Empirically, the effectiveness of our approach is underscored through a high-dimensional synthetic state estimation task and a 6D pose estimation derived from videos. We posit that as robotic systems evolve with greater degrees of freedom, particle filters tailored for high-dimensional state spaces will be indispensable.

I. INTRODUCTION

State estimation remains at the heart of many robotic applications, from autonomous navigation to manipulation tasks. Precise and timely estimates of the system’s state are critical to ensure the safe and effective operation of robotic systems. Among state estimation techniques, particle filters have emerged as a particularly potent tool [1]. Their non-parametric nature allows them to flexibly model a wide variety of distributions, making them ideal for capturing the intricacies of complex robotic systems.

However, using particle filters in high-dimensional state spaces has specific challenges. One issue is the so-called ‘particle deprivation’ or ‘sample impoverishment’ [2]. This refers to the scenario where large regions of the state space lose all representation during particle resampling, making it challenging to recover density in these regions later if necessary. In high-dimensional spaces, particle deprivation becomes increasingly problematic, often resulting in a particle filter that does not sufficiently represent the true posterior distribution.

In light of these challenges, this paper presents three primary contributions:

- 1) We introduce a novel resampling-free particle filter that mitigates the issue of particle deprivation. By avoiding the conventional resampling step, our approach ensures a more diverse set of particles, enhancing representation in high-dimensional spaces.
- 2) On the theoretical front, we rigorously show that our proposed particle filter provides a close approximation to the desired posterior distribution, even when particles lie in high-dimensional state spaces.

- 3) We empirically demonstrate the efficacy of our method on a high-dimensional synthetic state estimation task and 6D pose estimation from videos.

This paper aims to bridge the gap between particle filter theory and its applicability in high-dimensional robotic systems. By presenting a resampling-free approach, we hope to pave the way for more robust state estimation methods in complex robotic applications.

II. RELATED WORK

Particle filters have gained immense popularity in state estimation tasks due to their flexibility in handling nonlinear and non-Gaussian systems [1], [3]. More recent work has further expanded the applicability of particle filters to scenarios in which models of the environment are not available by *learning* the components of a particle filter [4], [5], [6].

An integral step in both traditional and modern particle filter algorithms is resampling, which ensures that particles with low weights are replaced by replicating particles with high weights [7]; resampling aims to increase the number of particles in areas intended to have high density. However, frequent resampling can lead to particle deprivation, a phenomenon where the diversity of particles diminishes over time, leaving the filter susceptible to divergence from the true state [2], [8].

Recognizing these challenges, researchers have explored resampling-free particle filters to mitigate the problem of particle deprivation [9], [10]. An underlying feature of these methods is the use of deterministic variational inference techniques that allow for estimation of the posterior distribution [11]. These methods have shown potential in maintaining particle diversity over extended periods of operation. However, unlike traditional particle filters, for which convergence guarantees have been extensively studied [12], [13], [14], resampling-free particle filters lack theoretical guarantees regarding their performance, especially in high-dimensional state spaces.

Our work addresses this gap, offering theoretical insights into the performance of resampling-free particle filters in high-dimensional scenarios and providing empirical evidence to support these claims. While building upon the foundations laid by previous works, we introduce key innovations that set our method apart, ensuring robust state estimation even in complex robotic systems.

III. DESIGNING A RESAMPLING-FREE PARTICLE FILTER

Here, we develop a resampling-free particle filter by designing a particular *flow* over particles that enables particles to track the true density we wish to estimate. We then prove

¹MIT, ²Harvard

that our method converges to the true density even in high-dimensional state spaces. Finally, we demonstrate that our method is computationally efficient in time and memory.

A. Using Flows to Track Posterior Density

a) Setup: We consider a continuous-time hidden Markov model with hidden state $x \in \mathbb{R}^d$ and observations y_t , where t denotes time and d is the state space dimensionality. For our presentation, we consider a continuous-time setting, but our analysis and algorithm apply to discrete-time as well. Let $P_t(y|x)$ denote the probability density of observing y at time t such that the probability of observing a constant y in an interval from t_1 to t_2 is $e^{\int_{t_1}^{t_2} \log P_t(y|x) dt}$. Denote the negative log-likelihood of observing y_t at time t under state x as:

$$L_t(x) = -\log P_t(y_t|x) \quad (1)$$

We may interpret $L_t(x)$ as a loss function. We wish to find the posterior distribution $p_t(x)$ over x given all observations from y_0 to y_t . By Bayes' rule, the posterior updates as:

$$p_t(x) + dp_t(x) \propto p_t(x) P_t(y_t|x)^{dt} \quad (2)$$

where the constant of proportionality is set such that $p_t(x) + dp_t(x)$ integrates to 1. Intuitively, this equation states that the posterior after time dt is proportional to both the prior distribution $p_t(x)$ and the likelihood of observing the data y_t under state x for an interval dt . For small dt , using the definition of $P_t(y|x)$, this likelihood is expressed as $e^{\int_t^{t+dt} \log P_t(y_t|x) dt} = e^{dt \log P_t(y_t|x)} = P_t(y_t|x)^{dt}$. After taking log of both sides and dividing by dt , this yields an expression for the time derivative of the log posterior:

$$\frac{d}{dt} \log p_t(x) = Z_t - L_t(x) \quad (3)$$

where Z_t is a normalizing constant. Note that for $p_t(x)$ to always remain a probability distribution, $\frac{d}{dt} \int p_t(x) dx = 0$, which implies:

$$\begin{aligned} \int \frac{d}{dt} p_t(x) dx &= \int p_t(x) \frac{d}{dt} \log p_t(x) dx \\ &= \int p_t(x) (Z_t - L_t(x)) dx = 0 \end{aligned} \quad (4)$$

Thus,

$$Z_t = \mathbb{E}_{x \sim p_t} [L_t(x)] \quad (5)$$

For notational convenience, we define a normalized loss $\tilde{L}_t(x) = L_t(x) - Z_t$ representing the loss at x relative to the losses on the full density $p_t(x)$.

b) Defining a Vector Field: Next, we design a vector field $F_t \in \mathbb{R}^d \rightarrow \mathbb{R}^d$ that when applied to posterior p_t , will yield a flow that tracks the posterior over time. We will first assume F_t yields a flow correctly tracking p_t , then select a *specific* F_t satisfying this property. The continuity equation [15] can be used to relate vector fields to the changes in density they induce:

$$\dot{p}_t = -\nabla \cdot (p_t F_t) \quad (6)$$

where $\nabla \cdot$ denotes divergence and $\dot{\cdot}$ denotes time derivative. We may rewrite the left-hand side as:

$$p_t \frac{d}{dt} \log p_t = -\nabla \cdot (p_t F_t) \quad (7)$$

Substituting in Eqn 3:

$$p_t \tilde{L}_t = \nabla \cdot (p_t F_t) \quad (8)$$

Now, suppose there exists a time-dependent potential function $\psi_t \in \mathbb{R}^d \rightarrow \mathbb{R}$ such that its gradient with respect to x equals $p_t F_t$: $\nabla \psi_t = p_t F_t$. Then:

$$p_t \tilde{L}_t = \nabla^2 \psi_t \quad (9)$$

Note that ∇^2 represents the Laplacian operator. We may solve for ψ_t by inverting the Laplacian as:

$$\psi_t = -[p_t \tilde{L}_t] \star K \quad (10)$$

where \star denotes convolution and $K \in \mathbb{R}^P \rightarrow \mathbb{R}$ is a kernel defined as [16]:

$$K(\Delta x) = C \|\Delta x\|_2^{2-d} \quad (11)$$

where $C = \frac{\Gamma(\frac{d}{2}+1)}{d(d-2)\pi^{\frac{d}{2}}}$ and $\Gamma(\cdot)$ denotes the gamma function.

Then, the gradient of ψ can be expressed as:

$$\nabla \psi_t = -\nabla [p_t \tilde{L}_t] \star K \quad (12)$$

The resulting F_t can be written as:

$$F_t = -\frac{\nabla [p_t \tilde{L}_t] \star K}{p_t} \quad (13)$$

If vector field F_t is applied to density p_t , this yields a flow on p_t that *exactly* tracks the true posterior as specified by Bayes' rule (as specified by Eqn 3). Note that there may be other F_t satisfying this property; however, this is the unique vector field such that $p_t F_t$ can be expressed as the gradient of a function ψ_t .

c) Applying the Flow to Particles: In practice, we cannot always exactly model the posterior density p_t since it often may not take an easily parametrically variable form. Instead, we use a set of n particles $x_t^1, x_t^2, \dots, x_t^n$ to approximate the posterior density:

$$p_t(x) \approx \frac{1}{n} \sum_i \delta(x - x_t^i) \quad (14)$$

where $\delta(\cdot)$ denotes the δ function. From Eqn 13, note that $p_t F_t$ may be written as:

$$p_t F_t = -\nabla [p_t \tilde{L}_t] \star K \quad (15)$$

Plugging in the particle approximation:

$$p_t(x) F_t(x) = -\nabla \left[\frac{1}{n} \sum_i \delta(x - x_t^i) \tilde{L}_t(x) \right] \star K \quad (16)$$

Applying the convolution on the right-hand side:

$$p_t(x) F_t(x) = -\nabla \left[\frac{1}{n} \sum_i \tilde{L}_t(x_t^i) K(x - x_t^i) \right] \quad (17)$$

Intuitively, $p_t(x)F_t(x)$ corresponds to the movement of probability density at x . We would like to set the movement of each particle x_t^j such that its movement captures the *local* value of $p_t(x)F_t(x)$ around x_t^j , thus capturing the correct overall movement of probability density. In other words, we would like to set $F_t(x_t^j)$ such that:

$$F_t(x_t^j) = n\mathbb{E}_\zeta[p_t(x_t^j + \zeta)F_t(x_t^j + \zeta)] \quad (18)$$

where ζ represents a random variable taking values close to 0. Using our expression for $p_t(x)F_t(x)$:

$$\begin{aligned} F_t(x_t^j) &= \mathbb{E}_\zeta \left[-\nabla_{x_t^j} \left[\sum_i \tilde{L}_t(x_t^i) K(x_t^j - x_t^i + \zeta) \right] \right] \\ &= -\nabla_{x_t^j} \left[\sum_i \tilde{L}_t(x_t^i) \mathbb{E}_\zeta [K(x_t^j - x_t^i + \zeta)] \right] \end{aligned}$$

Finally, we simply approximate $\mathbb{E}_\zeta[K(x_t^j - x_t^i + \zeta)]$ as:

$$\mathbb{E}_\zeta[K(x_t^j - x_t^i + \zeta)] \approx C(\|x_t^j - x_t^i\|_2^2 + \gamma^2)^{\frac{2-d}{2}} \quad (19)$$

where γ is a small constant added to approximate the fact that $x_t^j - x_t^i + \zeta$ will always be non-zero with probability 1 even when x_t^j and x_t^i are equal. This yields the following final flow for particle x_t^j :

$$\begin{aligned} F_t(x_t^j) &= -C\nabla_{x_t^j} \left[\sum_i \tilde{L}_t(x_t^i) (\|x_t^j - x_t^i\|_2^2 + \gamma^2)^{\frac{2-d}{2}} \right] \\ &= -C\gamma^{2-d}\nabla_{x_t^j} \tilde{L}_t(x_t^i) \\ &\quad - C \sum_i \frac{(d-2)\tilde{L}_t(x_t^i)}{(\|x_t^j - x_t^i\|_2^2 + \gamma^2)^{\frac{d}{2}}} (x_t^i - x_t^j) \end{aligned} \quad (20)$$

Intuitively, the first term of the update simply performs gradient descent on the loss \tilde{L}_t . The second term corresponds to an attraction-repulsion force that shifts point x_t^j towards point x_t^i if the loss on point x_t^i is favorable (i.e. < 0) and away from point x_t^i if the loss is unfavorable (i.e. > 0). Please see Fig 1 for an illustration of our update rule; note that the attraction-repulsion force may result in an update *opposite* to the direction of gradient descent.

B. Theoretical Guarantees of Convergence to Posterior

The particle ensemble approximation of the posterior introduces an error with the true posterior. In this section, we quantify this approximation error theoretically.

Let us denote the true posterior by p_t and our particle ensemble approximation by q_t . We assume p_t flows according to the true vector field F_t while q_t flows under an approximated vector field \tilde{F}_t . For simplicity, we assume that both p_t and q_t are discrete and consist of N points each, denoted $\{x_t^i\}_{i=1}^N$ for p and $\{z_t^i\}_{i=1}^N$ for q , though the bound can be generalized to the continuous case by taking the limit in N .

We use the Wasserstein distance as our distance measure. The Wasserstein distance between p_t and q_t is defined as

$$W(p_t, q_t) = \min_{\pi} \sum_{i=1}^N d(x_t^i, z_t^{\pi(i)}) \quad (21)$$




Fig. 1: Example of the particle updates computed by our particle filter. Black lines indicate lines of constant likelihood P_t , the black arrow indicates increasing likelihood, red dots indicate particles. Sample update directions are indicated at four points: orange arrows indicate the gradient descent on the negative log-likelihood, and blue arrows indicate the attraction-repulsion force.

where the expression is defined over a metric space with distance measure d , and the minimum is taken over all permutations π on $[N]$. Let π_t denote the minimizing π for the distributions at time t . Additionally, suppose that d has the shift-invariant property; that is,

$$d(x + \delta, z + \delta) = d(x, z) \quad (22)$$

for all x, z, δ .

We present the following theorem on the approximation error.

Theorem 1: Let L_F and L_d be the Lipschitz constants of F and d , respectively. Further, suppose that the maximum pointwise discrepancy between F and \tilde{F} is bounded by $\frac{\varepsilon}{N}$. Then,

$$W(p_t, q_t) \leq \left(W(p_0, q_0) + \frac{\varepsilon}{L_F} \right) e^{L_d L_F t} - \frac{\varepsilon}{L_F}. \quad (23)$$

Proof: Assume that p_t follows the true flow F_t and q_t follows the approximate flow \tilde{F}_t . We then bound the Wasserstein distance at time $t + dt$ using the Wasserstein distance at time t .

$$\begin{aligned} W(p_{t+dt}, q_{t+dt}) &\leq \sum_i d(x_t^i + F_t(x_t^i)dt, z_t^{\pi_t(i)} + \tilde{F}_t(z_t^{\pi_t(i)})dt) \\ &= \sum_i d(x_t^i + [F_t(x_t^i) - \tilde{F}_t(z_t^{\pi_t(i)})]dt, z_t^{\pi_t(i)}) \\ &\leq \sum_i d(x_t^i, z_t^{\pi_t(i)}) + L_d \|F_t(x_t^i) - \tilde{F}_t(z_t^{\pi_t(i)})\| dt \end{aligned}$$

where we applied the shift-invariant property of d . We then




Fig. 2: Comparison of the performance (measured as KL divergence with the true posterior distribution) of our particle filter with standard Monte Carlo Localization (MCL) on a synthetic localization problem with adjustable dimensionality. n denotes the number of particles. Margins indicate standard errors over 10 trials. [Left] Performance over the course of particle filter iterations on a 10-dimensional problem. [Right] Performance at the final iteration over varying dimensions.

rewrite the above as

$$\sum_i d(x_t^i, z_t^{\pi_t(i)}) + L_d \|F_t(x_t^i) - F_t(z_t^{\pi_t(i)}) + F_t(z_t^{\pi_t(i)}) - \tilde{F}_t(z_t^{\pi_t(i)})\| dt$$

which can be bounded by

$$\sum_i d(x_t^i, z_t^{\pi_t(i)}) + L_d(L_F d(x_t^i, z_t^{\pi_t(i)}) + \varepsilon) dt \quad (24)$$

and rewritten as

$$(1 + L_d L_F dt)W(p_t, q_t) + L_d \varepsilon dt. \quad (25)$$

We thus obtain a differential equation as our final inequality:

$$W(p_{t+dt}, q_{t+dt}) \leq (1 + L_d L_F dt)W(p_t, q_t) + L_d \varepsilon dt \quad (26)$$

which solves via Grönwall's inequality [17] to

$$W(p_t, q_t) \leq \left(W(p_0, q_0) + \frac{\varepsilon}{L_F} \right) e^{L_d L_F t} - \frac{\varepsilon}{L_F}. \quad (27)$$

Our result reveals that the error at time t scales linearly with the error at time 0 (as n or d is varied). Critically, this means that even in very high-dimensional state spaces, if a strong initial approximation of the prior distribution p_0 is provided, then the approximated posterior q_t is guaranteed to track p_t up to some constant error term and scale factor. In other words, if $W(p_0, q_0)$ is bounded by a dimension-independent constant K , then $W(p_t, q_t)$ can also be bounded by a dimension-independent constant: *our particle filter avoids the curse of dimensionality*. We note that $W(p_0, q_0)$ may generally scale with the dimensionality of the state space: in higher dimensional spaces, more particles may be necessary to approximate any distribution to the same precision. However, this dimensionality dependence is purely

due to the difficulty of approximating distributions in high-dimensional spaces, and is unavoidable for any particle filter. Our particle filter adds *no further error* in high-dimensional spaces relative to low-dimensional ones.

Moreover, unlike classical convergence results on standard (resampling-based) particle filters [12], due to the deterministic nature of our particle filter (it avoids any sampling), our convergence results are not probabilistic. Instead, convergence is always guaranteed to hold.

Interestingly, note that we may directly apply the classical convergence results [12] at time step 0 of our algorithm, and achieve the same convergence rate (up to constant terms and scale factors) at future timesteps t . For instance, the classical $O(\frac{1}{n})$ convergence rate of $(\mathbb{E}_p[f(x)] - \mathbb{E}_q[f(x)])^2$ for any bounded function f directly applies to our filter as well, including the dimension-independence of the convergence rate.

C. Computational Efficiency

We present our full algorithm in Algorithm 1. Observe that each step of the algorithm requires looping over all n particles to update each of them, and each particle's update step requires computing displacements with all other particles. Thus, our algorithm's runtime is $O(n^2 T)$ where T is the number of total timesteps and n is the number of particles. Note that the quadratic dependence on n can be computationally expensive for large numbers of particles. However, as we have previously shown, if the initial error of our particle filter is bounded, then the error in the posterior is *independent of dimension*; thus, the required number of particles to achieve a particular error rate also does not scale with dimension.

Storing each particle requires $O(d)$ memory, resulting in a memory consumption of $O(nd)$, which is the minimum that can be reasonably expected to represent the particles.




Fig. 3: Performance Analysis of 6D Pose Estimation Methods on the "Cracker Box" Model: This figure presents a comparative analysis of the translation and rotation mean error between our proposed method and the conventional gradient descent approach. The evaluation is based on 1912 RGB video sequences of the "cracker box" model. Both methods were tested under identical configurations, using 80 particles. The mean error is computed over the x, y, z coordinates and averaged for both translation and rotation. Specifically, the translation error measures the deviation in centimeters from the ground truth translation, while the (signed) rotation error quantifies the angular difference in degrees from the true rotation. Note that the rotation error can be positive or negative, with the optimal value being 0.

Algorithm 1 Resampling-free Particle Filter

Require: Initial particles $\{x_0^i\}_{i=1}^n$, time step Δt , small constant γ , total timesteps T , negative log-likelihood function L_t , constant C , dimensionality d

- 1: **for** $t = 1, 2, \dots, T$ **do**
 - 2: Compute $L_t(x_t^i)$ for all i
 - 3: $Z_t = \frac{1}{n} \sum_i L_t(x_t^i)$
 - 4: **for** $j = 1, 2, \dots, n$ **do**
 - 5: $\tilde{L}_t(x_t^j) = L_t(x_t^j) - Z_t$
 - 6: $x_{t+1}^j = x_t^j - C\gamma^{2-d} \nabla_{x_t^j} \tilde{L}_t(x_t^j)$
 - 7: $x_{t+1}^j = x_{t+1}^j - C \sum_i \frac{(d-2)\tilde{L}_t(x_t^i)}{(\|x_t^j - x_t^i\|_2^2 + \gamma^2)^{\frac{d}{2}}} (x_t^i - x_t^j)$
 - 8: **end for**
 - 9: **end for**
-

IV. EXPERIMENTS

To assess the effectiveness of our approach on high-dimensional localization, we conduct experiments on a synthetic high-dimensional localization task with a variable state-space dimensionality and 6D pose estimation.

A. Synthetic Localization

a) *Setting:* We consider the following log-likelihood function:

$$L_t(x) = \frac{1}{2}((x - u)^T \xi_t)^2 \quad (28)$$

where $u \in \mathbb{R}^d$ is an *unknown* parameter drawn from $N(0, I)$ and $\xi_t \in \mathbb{R}^d$ are drawn iid from $N(0, I)$. Due to the simple form of the likelihood, we may analytically solve for the true posterior over x . Specifically, assuming a prior of $N(0, I)$

over x , the corresponding expected log posterior after t timesteps may be expressed as:

$$\log p_t(x) = -\frac{1}{2}\|x\|_2^2 - \frac{t}{2}\|x - u\|_2^2 \quad (29)$$

We compare our particle filter and a baseline of standard Monte Carlo Localization (MCL) [18]. For MCL, we add Gaussian noise from $N(0, \epsilon I)$ as the motion model: $P(\xi_{t+1} - \xi_t | \xi_t) \sim N(0, \epsilon I)$, where ϵ is a hyperparameter. For both methods, we conduct grid searches over each method's relevant hyperparameter over 5 orders of magnitude and report results for the best setting. We run the particle filter for 50 iterations and evaluate performance at each iteration t using KL divergence between a multivariate Gaussian fit to the particles and the true posterior at time t . We also assess performance as a function of the state space dimensionality d .

b) *Results:* In Fig 2, we find that for both particle filters, the KL divergence initially increases as the true posterior shifts away from the prior distribution, but then decreases as the true posterior concentrates its density. Our method outperforms MCL over all iterations tested. Strikingly, while MCL worsens greatly as the number of particles decreases, our method maintains its performance. Moreover, our method much more gracefully handles higher dimensional state spaces, achieving roughly a $\times 10$ reduction in KL divergence compared to MCL at the highest dimension tested.

B. 6D Pose Estimation

a) *Setting:* 6D pose estimation, which involves estimating the 3D rotation and translation of objects relative to the

camera, is central to many robotic tasks such as manipulation and navigation, enabling efficient grasp planning, obstacle avoidance, and more. Historically, local-feature or template-matching techniques have been the backbone for estimating the 6D pose of objects. These methods typically involve matching image features against pre-generated templates or features of a 3D object model, thereby recovering the object’s pose [19].

We evaluate the performance of our particle filter to perform pose estimation. Here, note that each particle corresponds to a pose estimate for the object in question. We adapt the approach in [19] to factorize the posterior into two components: the 3D translation and the 3D rotation of the object; each particle corresponds to a specific translation and rotation. Thus, each particle has 6 dimensions represented as a combination of translation and rotation.

We use the same approach as in [19] to compute the likelihood of a particular observation y being consistent with a pose (T, R) consisting of translation T and rotation R . Critically, we assume access to a *reference* object for which we have observations under a variety of rotations. We use a pre-trained autoencoder that allows us to compare the similarity of two objects: the observation y under the candidate translation T and the appearance of reference object under candidate rotation R . If the pose is more likely, then the reference object at rotation R will be similar to the observation translated with T . Thus, we set generate a likelihood function as follows:

$$P(y|(T, R)) \propto S(g(\text{crop}(y, T)), g(\bar{y}^R)) \quad (30)$$

where g is a pre-trained encoder mapping images to a vector code, $\text{crop}(y, T)$ applies a crop to image y to isolate an object in the image assuming it is located at translation T , \bar{y}^R denotes a reference image of the object under rotation R , and S is a similarity function. Given this likelihood function, we run our particle filter as described in Algorithm 1 using randomly initialized initial pose estimates. We compare our particle filter with a baseline of simply applying gradient descent on the negative log-likelihood function (corresponding to ignoring the attraction-repulsion force in our update equation Eqn 20).

b) YCB dataset results: Our experiments were conducted on the YCB Video dataset [20], specifically on RGB video sequences of household objects, both textured and textureless. Each video sequence consists of a series of frames in which the object may have varying pose. The objects in this dataset are annotated with their 6D poses at each frame. For our experiment, we used the “cracker box” object which consisted of 1912 RGB video frames. We evaluated the performance of our particle filter at estimating the object’s pose over time using 80 particles. Specifically, we measured the error between the true pose of the object and mean pose estimated by our particles.

As seen in Figure 3, our results revealed a marked enhancement in 6D pose estimation accuracy using PoseRBPF as opposed to gradient descent. Our proposed method consistently yielded a diminished error when compared to the

gradient descent baseline. Both in terms of rotation and translation error rates, our algorithm exhibited superior efficiency in terms of number of iterations, arriving at values closer to the ground truth faster than gradient descent. This analysis underscores the potential of our approach in handling pose estimation tasks.

V. CONCLUSION

We develop a novel resampling-free particle filter designed to circumvent particle deprivation issues that arise when applying traditional particle filters in high-dimensional state spaces. Through theory and experiments, we demonstrate that our approach avoids the curse of dimensionality and is practically effective in high-dimensional localization. We believe particle filters explicitly designed for high-dimensional state spaces will be critical in practical applications as robots are increasingly designed with more degrees of freedom.

REFERENCES

- [1] S. Thrun, W. Burgard, and D. Fox, *Probabilistic robotics*. Cambridge, Mass.: MIT Press, 2005.
- [2] M. S. Arulampalam, S. Maskell, N. Gordon, and T. Clapp, “A tutorial on particle filters for online nonlinear/non-gaussian bayesian tracking,” *IEEE Transactions on signal processing*, vol. 50, no. 2, pp. 174–188, 2002.
- [3] A. Doucet, N. De Freitas, N. J. Gordon, *et al.*, *Sequential Monte Carlo methods in practice*, vol. 1. Springer, 2001.
- [4] R. Jonschkowski, D. Rastogi, and O. Brock, “Differentiable particle filters: End-to-end learning with algorithmic priors,” *Robotics: Science and Systems (RSS)*, 2018.
- [5] P. Karkus, S. Cai, and D. Hsu, “Differentiable slam-net: Learning particle slam for visual navigation,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2815–2825, June 2021.
- [6] A. Corenflos, J. Thornton, G. Deligiannidis, and A. Doucet, “Differentiable particle filtering via entropy-regularized optimal transport,” in *Proceedings of the 38th International Conference on Machine Learning* (M. Meila and T. Zhang, eds.), vol. 139 of *Proceedings of Machine Learning Research*, pp. 2100–2111, PMLR, 18–24 Jul 2021.
- [7] J. D. Hol, T. B. Schon, and F. Gustafsson, “On resampling algorithms for particle filters,” in *2006 IEEE nonlinear statistical signal processing workshop*, pp. 79–82, IEEE, 2006.
- [8] N. J. Gordon, D. J. Salmond, and A. F. Smith, “Novel approach to nonlinear/non-gaussian bayesian state estimation,” in *IEE proceedings F (radar and signal processing)*, vol. 140, pp. 107–113, IET, 1993.
- [9] M. Pulido and P. J. van Leeuwen, “Sequential monte carlo with kernel embedded mappings: The mapping particle filter,” *Journal of Computational Physics*, vol. 396, pp. 400–415, 2019.
- [10] F. A. Maken, F. Ramos, and L. Ott, “Stein particle filter for nonlinear, non-gaussian state estimation,” *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 5421–5428, 2022.
- [11] Q. Liu and D. Wang, “Stein variational gradient descent: A general purpose bayesian inference algorithm,” *Advances in neural information processing systems*, vol. 29, 2016.
- [12] D. Crisan and A. Doucet, “A survey of convergence results on particle filtering methods for practitioners,” *IEEE Transactions on signal processing*, vol. 50, no. 3, pp. 736–746, 2002.
- [13] N. Chopin, “Central limit theorem for sequential monte carlo methods and its application to bayesian inference,” *The Annals of Statistics*, 2004.
- [14] P. Del Moral, A. Doucet, and A. Jasra, “Sequential monte carlo samplers,” *Journal of the Royal Statistical Society Series B: Statistical Methodology*, vol. 68, no. 3, pp. 411–436, 2006.
- [15] H. Lamb, *Hydrodynamics*. University Press, 1924.
- [16] R. Cristofori, “The laplace and the poisson equations in the whole space,” 2017.
- [17] B. Pachpatte, “A note on gronwall-bellman inequality,” *Journal of Mathematical Analysis and Applications*, vol. 44, no. 3, pp. 758–762, 1973.

- [18] D. Fox, W. Burgard, F. Dellaert, and S. Thrun, "Monte carlo localization: Efficient position estimation for mobile robots," *Aaai/iaai*, vol. 1999, no. 343-349, pp. 2-2, 1999.
- [19] X. Deng, A. Mousavian, Y. Xiang, F. Xia, T. Bretl, and D. Fox, "Poserbpf: A rao-blackwellized particle filter for 6d object pose tracking," in *Robotics: Science and Systems (RSS)*, 2019.
- [20] Y. Xiang, T. Schmidt, V. Narayanan, and D. Fox, "Posecnn: A convolutional neural network for 6d object pose estimation in cluttered scenes," *Robotics: Science and Systems (RSS)*, 2018.