

Counter-Hypothetical Particle Filters for Single Object Pose Tracking

Elizabeth A. Olson

Jana Pavlasek

Jasmine A. Berry

Odest Chadwicke Jenkins

Abstract—Particle filtering is a common technique for six degree of freedom (6D) pose estimation due to its ability to tractably represent belief over object pose. However, the particle filter is prone to particle deprivation due to the high-dimensional nature of 6D pose. When particle deprivation occurs, it can cause mode collapse of the underlying belief distribution during importance sampling. If the region surrounding the true state suffers from mode collapse, recovering its belief is challenging since the area is no longer represented in the probability mass formed by the particles. Previous methods mitigate this problem by randomizing and resetting particles in the belief distribution, but determining the frequency of reinvigoration has relied on hand-tuning abstract heuristics. In this paper, we estimate the necessary reinvigoration rate at each time step by introducing a *Counter-Hypothetical likelihood function*, which is used alongside the standard likelihood. Inspired by the notions of plausibility and implausibility from Evidential Reasoning, the addition of our Counter-Hypothetical likelihood function assigns a level of doubt to each particle. The competing cumulative values of confidence and doubt across the particle set are used to estimate the level of failure within the filter, in order to determine the portion of particles to be reinvigorated. We demonstrate the effectiveness of our method on the rigid body object 6D pose tracking task.

I. INTRODUCTION

As robot assistants become tasked with accomplishing complex chores, such as preparing a meal or tidying a room, they must be able to interact with a variety of objects. Object pose estimation in unstructured scenes remains a challenge due to the ambiguity in perception which arises from occlusion and symmetries. Particle-based inference methods have been widely applied to six degree of freedom (6D) object pose estimation and tracking due to their ability to represent high-dimensional spaces with finite sample sets [1], [2], [3]. These methods model estimation uncertainty and are capable of maintaining multiple possible pose hypotheses, which provides robustness in challenging scenarios such as object occlusion and ambiguous symmetries.

Despite these promising properties, particle filter algorithms are forced to limit the size of their sample sets to ensure tractability for robotics applications. When applied to 6D pose tracking, particle filters typically can afford only a small number of samples when compared to the overall size of the continuous state space. Since the sample set is unable to completely cover the space, certain regions of the state space will contain no particles, making their representation in the belief distribution collapse. This phenomenon is called *particle deprivation*, and can occur due to poor initialization or the stochasticity of importance sampling. Regaining belief

E. A. Olson, J. Pavlasek, J. A. Berry, and O. C. Jenkins are with the Robotics Department, University of Michigan, Ann Arbor, MI, USA, {lizolson, pavlasek, jasab, ocj}@umich.edu.




Fig. 1: To estimate when the particle filter is failing, we measure the doubt associated with each sample through a Counter-Hypothetical likelihood function. We quantify both the evidence against a given estimate (gray), as well as in support of it (yellow). We estimate these quantities independently of one another, because they are not zero-sum due to ambiguity in the observation (blue). The relative magnitudes of these weightings fluctuate based on the quality of both the observations and estimates, as illustrated by a mug that is (A) unambiguously unlikely (B) plausible yet ambiguous due to the occluded handle (C) unambiguously likely (D) highly ambiguous.

in these regions is challenging and can cause the filter to converge to an incorrect local optimum.

One strategy for mitigating particle deprivation is *particle reinvigoration* [4], in which reinitialized samples are routinely added to the set. However, determining the portion of particles to be reinvigorated at a given iteration often requires tedious hand-tuning through trial and error. A number of adaptive approaches have been proposed to mitigate this challenge by leveraging the information provided in the likelihood function [5], [6], [7]. The likelihood function gives an importance weighting to each sample by measuring the correspondence of the hypothesis to the observed sensor data. However, it only provides a relative weighting of which samples are *better* or *worse*, and no indication of the absolute error in the sample set.

We introduce the Counter-Hypothetical Particle Filter (CH-PF) to counteract the problem of particle deprivation in high-dimensional state spaces. Our method proposes to quantify the confidence that the true state is unrepresented in the sample set. To this aim, we model the evidence *against* a particular hypothesis, termed the *Counter-Hypothetical likelihood*. Our work measures this weighting independently of the traditional likelihood through the lens of Evidential Reasoning (Dempster-Shafer Theory) [8]. This framework argues that the plausibility and implausibility of proving an outcome can be based on different factors, and are not zero-sum due to potential overlap and ambiguity of the underlying evidence. Each particle is given both a likelihood and a Counter-Hypothetical likelihood weight. Our method utilizes both likelihoods to quantify the cumulative confidence and

doubt across the sample set. The relationship between these values is used to reason about the likelihood that the true state is underrepresented in our sample set, and in turn used to compute an adaptive rate of particle reinvigoration.

In this paper, we propose the Counter-Hypothetical Particle Filter, a particle filtering algorithm designed to mitigate particle deprivation for 6D pose tracking in challenging environments. We introduce a Counter-Hypothetical likelihood and explain how its utilization alongside the traditional likelihood counteracts particle deprivation by leveraging implausibility information, as illustrated in Figure 1. We evaluate CF-PF on 6D single object pose tracking on the YCB Video Dataset [9]. Our method achieves better accuracy for cases of high occlusion, particularly when depth data is not available.

II. RELATED WORK

A. Object Pose Estimation and Tracking for Robotics

Pose estimation and tracking have received considerable attention in the robotics community. In recent years, various works have demonstrated the capability of data-driven methods to provide discriminative pose estimates over a single view [9], [10], [11], [12], or pose tracking over a sequence of observations [13]. These methods have achieved impressive results but are prone to inaccuracies, particularly in challenging scenes such as those with heavy clutter. Probabilistic inference methods instead maintain belief over pose estimates in order to provide additional robustness for robotic manipulation applications [14]. We focus on probabilistic inference for object pose estimation and tracking, specifically on particle filtering.

Particle filtering is an iterative inference algorithm which can represent an arbitrary nonparametric belief distribution using a set of weighted particles sampled from the state space [5]. Particle filtering is a common technique for 6D pose estimation and tracking [2], [15], [16], [17] due to its ability to efficiently approximate high-dimensional state spaces and represent multiple competing hypotheses in the belief. More recently, deep convolutional neural networks (CNNs) have been applied to particle filtering for pose estimation. Deng et al. [3] create an observation model from autoencoder embeddings which is used within a Rao-Blackwellized particle filter. Recently, this method has been extended to category-level tracking [18]. Though the main contribution of our work is the introduction of the Counter-Hypothetical likelihood function, we also demonstrate how it could be similarly learned in an end-to-end fashion.

To mitigate challenges associated with fully sampling the high-dimensional 6D pose space, previous works have leveraged domain-specific knowledge such as physical constraints [19], robotic arm joint angles [20], [1], or context information [21]. Belief propagation using particle belief representations [22] has been applied to parts-based object models in order to factor the high-dimensional articulated object localization task [23], [24]. Similarly, our work addresses the problem of particle deprivation in sampling-based methods for manipulation tasks, but we aim to do so through

adaptive reinvigoration, and do not rely on provided object-environment interaction models or parts-based object models.

B. Robust Particle Filtering

Many works have focused on mitigating particle deprivation. One approach is annealing, in which the distribution of importance weights is smoothed according to a hand-tuned schedule to avoid collapsing modes during importance sampling [25]. Pfaff et al. propose an adaptive method to smooth the importance weights using local density estimation around each particle for mobile robot localization [26]. In contrast, CH-PF does not require modification to the sampling weights, but instead handles particle deprivation through reinvigoration.

In the global localization stage of Monte Carlo localization for mobile robots, particle deprivation is a common problem [5], motivating many works to reinitialize samples as needed. One approach is to sample from an inverse distribution based on the sensor readings [27] or “reset” a subset of particles when the average likelihood of the sample set is low [6]. Augmented Monte Carlo Localization [4], [28] extends this idea by performing particle reinvigoration from a uniform distribution at a rate proportional to the difference between the long- and short-term averages of the particle weights, instead of a fixed threshold. These methods require a sensor model from which samples can be efficiently drawn, which is challenging to model in RGB images. Fox et al. proposed modifying the size of the sample set based on the quality of the sample approximation [7]. Zhang et al. propose a self-adaptive method which maintains a fixed sample size which is augmented by samples from a “similar energy region” [29]. This method requires a discretization of the state space. Recent work in localization leverages the estimates of a neural network by sampling from this proposal at a fixed rate and fuses the particles into the distribution through importance sampling [30].

Each of these methods uses the likelihood weights of the particles to estimate the quality of the sample set. CH-PF instead uses a separate source of information, the Counter-Hypothetical likelihood, alongside the likelihood function, to provide an estimate of the overall quality of the particle set. We draw inspiration from Evidential Reasoning [8] for measuring the evidence disproving a hypothetical estimate separately from the likelihood function that looks for supportive evidence.

III. BACKGROUND: PARTICLE FILTERING

We consider the problem of tracking a known object over time. Given a sequence of RGB images or RGB-D data, $z_{1:t}$, we seek to localize the pose, $x_t \in \mathcal{X}$, of an object at time t . We also model any motion to the system, caused by either user input or jittering, with $u_{1:t}$. Here \mathcal{X} represents the space of 6D poses, comprised of 3D translation and 3D rotation.

The Bayes filter seeks to model the posterior distribution of the state, $p(x_t | x_{1:t-1}, z_{1:t}, u_{1:t})$ by iteratively updating the distribution at each timestep t . The posterior is called the *belief* of x_t , $bel(x_t)$. At each time step, the predicted belief,

$\widehat{bel}(x_t)$, is obtained by applying the action model to the prior belief distribution. Employing the Markov assumption:

$$\widehat{bel}(x_t) = \int p(x_t | x_{t-1}, u_t) bel(x_{t-1}) dx_{t-1} \quad (1)$$

We can then update this distribution based on the current observation, z_t , to estimate the posterior distribution:

$$bel(x_t) = p(z_t | x_t) \widehat{bel}(x_t) \quad (2)$$

A. Particle Filtering

The particle filter is a Bayes filtering algorithm in which the belief distribution $bel(x_t)$ is a nonparametric distribution approximated by a particle set, \mathbb{X}_t :

$$\mathbb{X}_t = \{(x_t^1, \pi_t^1), (x_t^2, \pi_t^2), \dots, (x_t^N, \pi_t^N)\} \quad (3)$$

Each particle, x_t^i has a corresponding weight, π_t^i . The predicted belief in Equation (1) is formed by applying action u_t to each particle in the previous sample set, \mathbb{X}_{t-1} .

The particle set, \mathbb{X}_t , is generated through *importance sampling*, where the target distribution is $bel(x_t)$ and the proposal distribution is $\widehat{bel}(x_t)$. Samples from the proposal are drawn with replacement, in which the probability of a particle being drawn is proportional to its weight, π_t^i . Typically, the weight is computed using a *likelihood function*, $\mathcal{L}(x_t^i)$, which represents the observation model:

$$\pi_t^i = \mathcal{L}(x_t^i) := p(z_t | x_t^i) \quad (4)$$

B. Particle Deprivation & Particle Reinvigoration

If the proposal distribution does not include samples close to the true value of the state, the probability of sampling values in this region is negligibly small. This phenomena is known as *particle deprivation*, and is illustrated in Figure 2 (left). It can occur due to poor initialization, unmodelled movements in the state, or a series of unfortunate draws in importance sampling, causing the particle set to converge to local optima.

A common approach to mitigating particle deprivation is *particle reinvigoration*, in which particles are drawn jointly from the predicted belief, $\widehat{bel}(x_t)$ and a candidate distribution, $\phi_{cand}(x_t)$. Choices of candidate distribution might include a uniform distribution over the region of interest of the state space, or a wide Gaussian distribution around an initial estimate. This modification allows importance sampling to draw from outside of the sample set, reintroducing samples in underrepresented regions. A hyperparameter α , where $0 \leq \alpha \leq 1$, controls the proportion of samples to be drawn from $\phi_{cand}(x_t)$. The final particle set is defined as the union of particles drawn from each set:

$$\mathbb{X}_t = \{(x_t^1, \pi_t^1), \dots, (x_t^{\alpha N}, \pi_t^{\alpha N})\} \cup \{(x_t^{\alpha N+1}, \pi_t^{\alpha N+1}), \dots, (x_t^N, \pi_t^N)\} \quad (5)$$

where $x_t^i \sim \phi_{cand}$ for $1 \leq i \leq \alpha N$, and $x_t^j \sim \widehat{bel}$ for $\alpha N < j \leq N$. Note that in practice, αN is constrained to be an integer.

IV. COUNTER-HYPOTHETICAL PARTICLE FILTER

Selecting the reinvigoration rate, α , is challenging in practice. Sampling from the candidate distribution too frequently can discard key information from the belief distribution, while sampling from the belief distribution too frequently could lead to particle deprivation. *Adaptive particle reinvigoration* mitigates this challenge by determining the frequency with which to draw from each distribution online at each time step. The Counter-Hypothetical Particle Filter (CH-PF) adaptively selects the reinvigoration rate, α , such that it fluctuates in accordance with the portion of the true belief distribution estimated to be underrepresented by \mathbb{X}_t .

One method of achieving adaptive particle reinvigoration is Sensor Resetting Localization (SRL) [6]. This method defines a probability threshold, β , which represents a threshold on “good” unnormalized likelihood values. The reinvigoration rate is defined as:

$$\alpha = 1 - \left(\frac{1}{\beta N} \sum_{i=1}^N \mathcal{L}(x_t^i) \right) \quad (6)$$

CH-PF builds off this equation, using a Counter-Hypothetical likelihood to compute the reinvigoration rate instead of a probability threshold.

A. Counter-Hypothetical Resampling

To motivate our proposed method for determining the reinvigoration rate, we first rewrite Equation (6):

$$\alpha = 1 - \frac{\sum_{i=1}^N \mathcal{L}(x_t^i)}{(\sum_{i=1}^N f(x_t^i)) + (\sum_{i=1}^N \mathcal{L}(x_t^i))} \quad (7)$$

where $f(x_t^i) := \beta - \mathcal{L}(x_t^i)$. With this notation, the numerator and right hand side of the denominator are an aggregate measurement of the likelihood of the sample set. The left hand side of the denominator, $\sum_{i=1}^N f(x_t^i)$, measures the poor performance across the sample set. In this way, calculating the rate of particle reinvigoration in SRL can be seen as simultaneously measuring the positive performance and poor performance of the sample set. However, the measure of poor performance, $f(x_t^i)$, is dependent on the measure of positive performance, as it is defined by $\mathcal{L}(x_t^i)$. This dependency is due to traditional Bayesian probability, in which the observed probability of a state being true and the observed probability of a state being false are always zero-sum.

Our method relaxes this assumption by taking inspiration from Evidential Reasoning, also known as Dempster-Shafer Theory [8]. This paradigm allows for the evidence discounting an event, *generalized disbelief*, to be quantified independently of the evidence associated with supporting an event, *generalized belief*. As shown in Figure 3, Evidential Reasoning models this concepts, as well as the presence of ambiguity or ignorance [32]. Generalized belief is a measurement of all evidence that undeniably supports an event and is bounded from above by plausibility, because plausibility includes the ambiguity in the belief. Similarly, generalized disbelief quantifies the evidence that works to




Fig. 2: An illustration of the proposed modification to the traditional resampling step of the particle filter, visualized as was presented in the Condensation Algorithm work [31]. In the traditional particle filter (left), each iteration begins with a set of weighted particles (top). The samples are drawn with replacement, then the action model and diffusion is applied to create the prediction distribution (blue). Each sample is then passed through the likelihood function (yellow) to receive a weighting. This posterior distribution then becomes the prior for the next iteration. In our proposed modification (right), each iteration begins with a set of weighted particles (top), as well as a weighting for the Counter-Hypothetical (black). In the resampling stage, the prior distribution is only sampled off of for five of the eight particles (blue), and the Counter-Hypothetical weighting causes three samples to be randomly sampled from the candidate distribution (pink). All of these samples are passed through the Counter-Hypothetical likelihood to be assigned a Counter-Hypothetical weighting (gray). These raw weightings are summed to create a new, singular Counter-Hypothetical weighting representing the doubt across the set. All samples are also passed through the traditional likelihood function (yellow). The posterior distribution and Counter-Hypothetical weighting then become the inputs for the next iteration.




Fig. 3: Relationship of the quantities in Evidential Reasoning [8]. This paradigm inspires our Counter-Hypothetical likelihood function visualized in Figure 1

disprove an event, while implausibility is an upper bound that considers ambiguity.

We posit that this framework is apt for our application of evaluating object poses based on images. The occlusions and geometric symmetries present suggest that there is ambiguity in how a given pose is supported or unsupported by the evidence, motivating us to measure these quantities independently. We consider the likelihood function to be analogous to Evidential Reasoning’s notion of generalized belief, and therefore introduce a Counter-Hypothetical likelihood to function in accordance with generalized disbelief.¹

We design the Counter-Hypothetical likelihood to measure how the observed image provides evidence *against* the hypothetical proposed state. Our approach replaces $f(x_t^i)$ from Equation (6) with a Counter-Hypothetical likelihood, which is estimated independently of $\mathcal{L}(x_t^i)$. We introduce a function to reason about the confidence of a state counter to our given hypothesis, the *Counter-Hypothetical likelihood*, $\mathcal{C}(x_t)$. By comparing the quantities of the (unnormalized) likelihoods, $\mathcal{L}(x_t^i)$ and $\mathcal{C}(x_t^i)$, across the proposal distribution, we can

¹A common critique of Evidential Reasoning is its appliance to true probability, instead of the probability of provability [33]. We do not directly use generalized belief to reason about the true underlying probability distribution, we merely take inspiration from Evidential Reasoning by measuring doubt independently of confidence.

reason about the cumulative confidence and doubt in our sample set. To this end, we redefine α , the reinvigoration ratio, as follows by modifying Equation (7):

$$\alpha = 1 - \frac{\sum_{i=1}^N \mathcal{L}(x_t^i)}{\sum_{i=1}^N \mathcal{C}(x_t^i) + \sum_{i=1}^N \mathcal{L}(x_t^i)} \quad (8)$$

We then sample αN particles from ϕ_{cand} in accordance with our doubt in the set, and sample the remaining $(1 - \alpha)N$ particles from \widehat{bel} based on our confidence in the set. As such, the Counter-Hypothetical likelihood quantifies our notion of generalized disbelief, which controls the amount of particle reinvigoration to be performed.

B. Counter-Hypothetical Likelihood for 6D Pose Estimation

In the context of our application, we design the Counter-Hypothetical likelihood to signal when the pose tracking is in a failure mode. Unlike traditional likelihood functions that must be precisely crafted or trained to ensure the most accurate samples have the highest weights, the Counter-Hypothetical likelihood function can be more crudely or intuitively constructed.

For a simple example, consider how the captured depth data can be compared against a rendered depth image of a sample at a candidate pose. A traditional likelihood function might assign weightings based on the number of pixels of the object that have rendered and captured depths within a given threshold. This heuristic can be noisy due to the presence of occlusions and difficult to tune. On the other hand, the Counter-Hypothetical likelihood could just measure the number of pixels in which the rendered depth is *less* than the captured depth. Measured depth being significantly closer than the rendered depth can be explained away by

a potential occlusion, while the reverse indicates the given pose is wrong.

In this work we use deep learning for the Counter-Hypothetical likelihood function, and use the encoder architecture of PoseRBPF [9]. We also use their synthetic training data setup. However, instead of training the representation in an auto-encoder manner, we leverage the ground truth pose information present in synthetic data. For our pairs of synthetically rendered scenes and rendered pose candidates, half the training data are positive samples, where the true pose is only slightly perturbed in the candidate rendering. The negative samples are generated by the candidate pose rendering being randomized. These crops are the inputs to the encoder network, the output embeddings from which are passed together through three fully connected layers. The network is trained with binary cross entropy loss. Through this fashion, the classifier learns to estimate when a hypothetical pose is misaligned with the given observation. At test time, the scores are used as the Counter-Hypothetical weightings.

V. EXPERIMENTS

To evaluate the proposed Counter-Hypothetical likelihood, we measure key performance metrics on the YCB-Video Dataset, a benchmarked real-world dataset [9]. We implement a standard particle filter to estimate the 6D pose of a given object across the video sequences. Our results test on both RGB and RGB-D data. All of the results are variants of the same particle filter, using the same likelihood function, provided by PoseRBPF [3] and use the same number of particles (50). However each baseline has a different strategy for combating particle deprivation, such as reinvigoration or Rao-Blackwellization. Whenever a candidate distribution is needed for initialization or reinvigoration, a uniform distribution of orientations located in the estimated 2D bounding box from PoseCNN [9] is used. Depth values are sampled from a uniform distribution, but when depth data is present, it is sampled off the measured depth for the object’s location.

A. Baselines

We compare against methods for particle deprivation specifically designed for 6D pose estimation, as well as adopting other techniques common in mobile robot localization.

Annealing [25] does not use any particle reinvigoration, but rather has an annealed likelihood function that cyclically smooths the likelihood weightings.

SRL [6] performs adaptive particle reinvigoration from the candidate distribution by comparing the average unnormalized likelihood weighting to a predetermined user threshold. This is a minimum cosine similarity between the embeddings in PoseRBPF.

Aug. MCL [4], [28] also performs adaptive particle reinvigoration from the candidate distribution, but the threshold is determined at each time step by user defined decay rates.

MCL + E2E [30] has a fixed number of samples coming from the predicted distribution with the remaining being




Fig. 4: Area Under the Curve scores for all methods for object 6D pose tracking on the YCB Video Dataset. ADD scores and the symmetric version (ADD-S) are presented. Our presented method, CH-PF, outperforms other methods focused on alleviating particle deprivation in instances of occlusion, when the information from the likelihood function is the most noisy (RGB data).

sampled off a neural network estimate. We sample off a Gaussian distribution centered at the full 6D pose estimate provided by PoseCNN for the current frame.

PoseRBPF [3] is ran as described in the publication, but with disabling any ground truth information used in the system. The original implementation ensures the initialization is close to the ground truth orientation, or it is completely reset. In our experiment, it is only reset when the likelihood weighting of the estimate drops below a threshold. We also include its suggested variant, **PoseRBPF++**, in which half the samples are reinvigorated at each time step from the candidate distribution.

CH-PF, Annealing, SRL, Aug. MCL, and MCL + E2E are filtering across all six dimensions to better test their ability to withstand particle deprivation. Through their Rao-Blackwellized implementation, PoseRBPF and PoseRBPF++ are filtering across only three dimensions of continuous state space (translation), because the orientation space is discretized.

VI. RESULTS

We present results with the absolute and symmetric point-wise matching errors between the estimated and ground truth pose (commonly referred to as ADD and ADD-S respectively) [9]. These errors are analyzed by viewing the Area Under the Curve (AUC) score of each method up to 10 cm error, as in other YCB works. Full quantitative results are shown in Figure 4.

Our analysis shows that for variants of a particle filter using the same weights but different methods for maintaining particle diversity, there is little deviation in performance

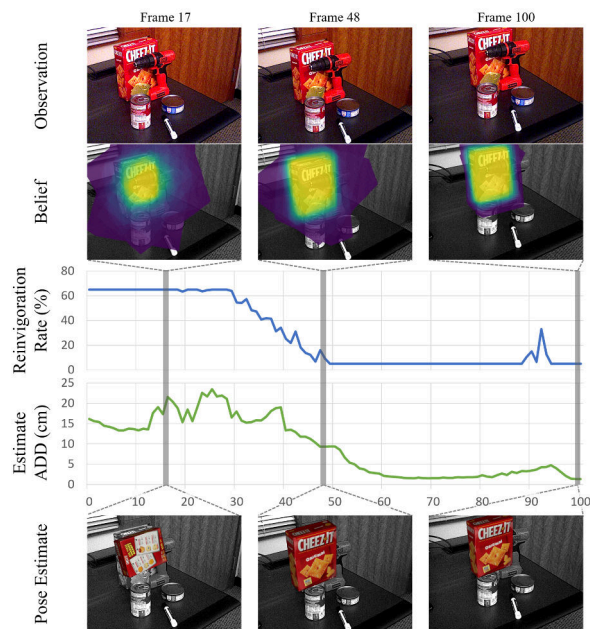


Fig. 5: Selected qualitative results for the Counter-Hypothetical Particle Filter. The cracker box is significantly occluded by other objects in the scene. In early iterations (left), the cracker box belief is not converged and the estimate has high error. The reinivigoration rate, calculated from the belief, is high. In later iterations (middle), the belief converges to the ground truth state, and the reinivigoration rate drops. Once the belief has converged (right), the reinivigoration rate is low. This figure is best viewed in color.

across the entire dataset. The Rao-Blackwellized implementations, PoseRBPf and PoseRBPf++, have the best performance. However, when looking at specifically the sequences where the given object is occluded, a disparity in accuracy is visible. The orientation filtering mechanism in PoseRBPf makes the system very sensitive to the quality of the initialization; an initial orientation that is flipped is difficult to correct, even with reinivigoration. For this reason, the poor initialization occurring in partially observable scenes lowers its accuracy ranking with respect to the other methods. Aug. MCL can work well when the filter loses track of the pose, but in the case of poor initialization, it does not record a value for an ideal likelihood threshold, and therefore performs little reinivigoration during occlusion. SRL and MCL + E2E perform on-par with or without occlusions being present. Our method performs similar to the rest across the entire dataset, but has the highest AUC accuracy for RGB data with occlusions.

Selected qualitative results for the Counter-Hypothetical Particle Filter method are shown in Figures 5 and 6. In these examples, the particle filter converged to an incorrect estimate at the beginning of the sequence. While the error was high, the Counter-Hypothetical Particle Filter was able to perform continued global localization with a high percentage of particles performing a coarse search. Once the error dropped, and a plausible region was found, it then reduced the reinivigoration rate to focus its resources on exploring the space more closely.

The main limitation of our work is additional computation time used at test time to evaluate each sample through an

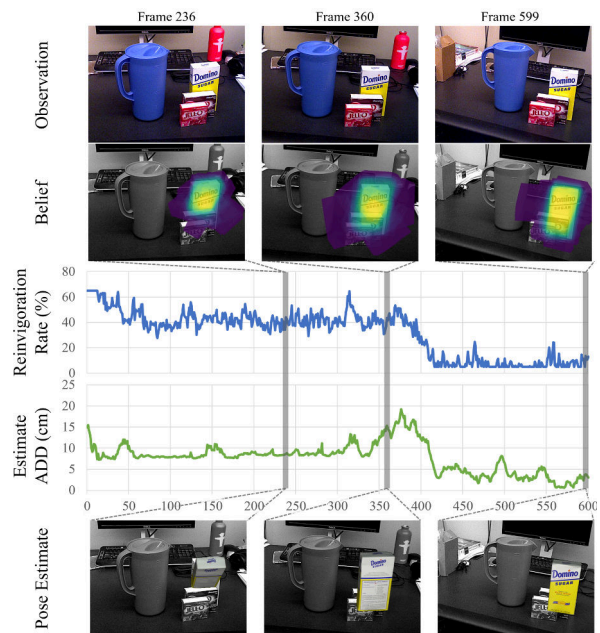


Fig. 6: Selected qualitative results for the Counter-Hypothetical Particle Filter. The belief of the sugar box converges to a local maximum in early frames (left). CH-PF applies a higher reinivigoration rate to mitigate this. The error in the estimate briefly increases (middle), but the belief eventually converges to the correct estimate (right).

additional likelihood function. For simple heuristics, this would not add much time. In our case of another neural network, it doubles the inference time. Our performance was similar to that of using a single likelihood function for most of the dataset, but improvements in performance for heavily occluded scenes are promising.

VII. CONCLUSION

This work aims to improve the accuracy of particle filters in tracking the 6D pose of rigid objects by adapting the rate of particle reinivigoration based on the estimated incompleteness of the current belief distribution. We propose independently estimating the potential error in each samples through a novel Counter-Hypothetical likelihood function. This modification allows us to reason over the cumulative doubt in our particle set, and use this estimate to apply particle reinivigoration as needed. This paper demonstrates the effectiveness of this modification as it matches overall performance when compared to standards methods of overcoming particle deprivation. Moreover, our particle filter proposed modification improves performance for scenes in heavy occlusion when only RGB data is present.

Future work will explore more cases for which our proposed incorporation of Evidential Reasoning into Bayesian filters could have more significant increases in performance, such as higher dimensional tracking cases or more dynamic objects and scenes. It will also investigate more computationally effective ways of incorporating this reasoning for real-time inference.

REFERENCES

- [1] M. Wuthrich, P. Pastor, M. Kalakrishnan, J. Bohg, and S. Schaal, "Probabilistic object tracking using a range camera," in *International*

- Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2013, pp. 3195–3202.
- [2] Z. Sui, O. C. Jenkins, and K. Desingh, “Axiomatic particle filtering for goal-directed robotic manipulation,” in *International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2015, pp. 4429–4436.
 - [3] 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.
 - [4] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics*. MIT Press, 2005.
 - [5] 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.
 - [6] S. Lenser and M. Veloso, “Sensor resetting localization for poorly modelled mobile robots,” in *International Conference on Robotics and Automation (ICRA)*, vol. 2, 2000, pp. 1225–1232.
 - [7] D. Fox, “Adapting the sample size in particle filters through kld-sampling,” *The international Journal of robotics research*, vol. 22, no. 12, pp. 985–1003, 2003.
 - [8] G. Shafer, *A mathematical theory of evidence*. Princeton university press, 1976, vol. 42.
 - [9] Y. Xiang, T. Schmidt, V. Narayanan, and D. Fox, “PoseCNN: A convolutional neural network for 6D object pose estimation in cluttered scenes,” in *Robotics: Science and Systems (RSS)*, 2018.
 - [10] J. Tremblay, T. To, B. Sundaralingam, Y. Xiang, D. Fox, and S. Birchfield, “Deep object pose estimation for semantic robotic grasping of household objects,” in *Conference on Robot Learning (CoRL)*, 2018.
 - [11] C. Wang, D. Xu, Y. Zhu, R. Martín-Martín, C. Lu, L. Fei-Fei, and S. Savarese, “DenseFusion: 6D object pose estimation by iterative dense fusion,” in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019, pp. 3343–3352.
 - [12] Y. He, H. Huang, H. Fan, Q. Chen, and J. Sun, “FFB6D: A full flow bidirectional fusion network for 6D pose estimation,” in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021, pp. 3003–3013.
 - [13] B. Wen, C. Mitash, B. Ren, and K. E. Bekris, “se(3)-TrackNet: Data-driven 6D pose tracking by calibrating image residuals in synthetic domains,” in *International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 10 367–10 373.
 - [14] M. Stoiber, M. Sundermeyer, and R. Triebel, “Iterative corresponding geometry: Fusing region and depth for highly efficient 3D tracking of textureless objects,” in *Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE/CVF, 2022, pp. 6855–6865.
 - [15] Z. Sui, Z. Zhou, Z. Zeng, and O. C. Jenkins, “SUM: Sequential scene understanding and manipulation,” in *International Conference on Intelligent Robots and Systems (IROS)*, 2017, pp. 3281–3288.
 - [16] S. Li, S. Koo, and D. Lee, “Real-time and model-free object tracking using particle filter with joint color-spatial descriptor,” in *International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2015, pp. 6079–6085.
 - [17] C. Choi and H. I. Christensen, “Robust 3d visual tracking using particle filtering on the special euclidean group: A combined approach of keypoint and edge features,” *The International Journal of Robotics Research*, vol. 31, no. 4, pp. 498–519, 2012.
 - [18] X. Deng, J. Geng, T. Bretl, Y. Xiang, and D. Fox, “iCaps: Iterative category-level object pose and shape estimation,” *Robotics and Automation Letters (RA-L)*, 2022.
 - [19] K. Desingh, O. C. Jenkins, L. Reveret, and Z. Sui, “Physically plausible scene estimation for manipulation in clutter,” in *International Conference on Humanoid Robots (Humanoids)*. IEEE, 2016, pp. 1073–1080.
 - [20] C. G. Cifuentes, J. Issac, M. Wüthrich, S. Schaal, and J. Bohg, “Probabilistic articulated real-time tracking for robot manipulation,” *Robotics and Automation Letters*, vol. 2, no. 2, pp. 577–584, 2016.
 - [21] Z. Zeng, A. Röfer, and O. C. Jenkins, “Semantic linking maps for active visual object search,” in *International Conference on Robotics and Automation (ICRA)*. IEEE, 2020.
 - [22] A. Ihler and D. McAllester, “Particle belief propagation,” in *Artificial intelligence and statistics*. PMLR, 2009, pp. 256–263.
 - [23] K. Desingh, S. Lu, A. Opiari, and O. C. Jenkins, “Efficient non-parametric belief propagation for pose estimation and manipulation of articulated objects,” *Science Robotics*, vol. 4, no. 30, 2019.
 - [24] J. Pavlasek, S. Lewis, K. Desingh, and O. C. Jenkins, “Parts-based articulated object localization in clutter using belief propagation,” in *International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020.
 - [25] J. Deutscher, A. Blake, and I. Reid, “Articulated body motion capture by annealed particle filtering,” in *Conference on Computer Vision and Pattern Recognition (CVPR)*, vol. 2. IEEE, 2000, pp. 126–133.
 - [26] P. Pfaff, W. Burgard, and D. Fox, “Robust Monte-Carlo localization using adaptive likelihood models,” in *European robotics symposium 2006*. Springer, 2006, pp. 181–194.
 - [27] S. Thrun, D. Fox, W. Burgard, and F. Dellaert, “Robust monte carlo localization for mobile robots,” *Artificial intelligence*, vol. 128, no. 1-2, pp. 99–141, 2001.
 - [28] J.-S. Gutmann and D. Fox, “An experimental comparison of localization methods continued,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, vol. 1. IEEE, 2002, pp. 454–459.
 - [29] L. Zhang, R. Zapata, and P. Lepinay, “Self-adaptive Monte Carlo localization for mobile robots using range finders,” *Robotica*, vol. 30, no. 2, pp. 229–244, 2012.
 - [30] N. Akai, T. Hirayama, and H. Murase, “Hybrid localization using model- and learning-based methods: Fusion of monte carlo and e2e localizations via importance sampling,” in *International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 6469–6475.
 - [31] M. Isard and A. Blake, “Condensation—conditional density propagation for visual tracking,” *International journal of computer vision*, vol. 29, no. 1, pp. 5–28, 1998.
 - [32] J. Y. Halpern, *Reasoning about uncertainty*. MIT press, 2017.
 - [33] J. Pearl, *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Morgan kaufmann, 1988.