Possibilistic Filtering for Reliable Robot Localization

Jan Schneider[1], Tom Könecke[1], Henrik Ebel[2], and Michael Hanss[1]
[1] *Institute of Engineering and Computational Mechanics, University of Stuttgart, Germany*
[2] *Department of Mechanical Engineering, LUT University, Lappeenranta, Finland*

Abstract. Uncertainty in engineering applications is commonly addressed within a probabilistic framework that assumes an underlying aleatory nature. However, in many practical scenarios, uncertainty can arise from incomplete knowledge rather than inherent randomness, reflecting epistemic rather than purely stochastic effects. Classical probabilistic approaches do not explicitly distinguish between these types of uncertainty, which can lead to overly confident or diluted probability assessments. This motivates the use of more expressive frameworks for uncertainty quantification. Possibility theory, as a theory of imprecise probabilities, provides a unified and computationally efficient framework for representing both epistemic and aleatory uncertainty. In particular, possibilistic filtering techniques have recently been developed for dynamic state estimation problems. In this contribution, we apply a particle-based possibilistic filtering approach to a real-world robot localization problem. The considered setup involves significant epistemic uncertainty arising from ambiguous landmark configurations and measurement limitations. The results demonstrate that the proposed method yields robust state estimates that reflect the underlying uncertainty. Furthermore, we introduce modifications to the original filter implementation that substantially reduce the computational cost.

Keywords: uncertainty quantification, possibility theory, state estimation, filtering, measurement uncertainty, robot localization

1. Introduction

In many engineering applications, uncertainties are not purely random but stem also from a lack of knowledge. Classical probabilistic approaches struggle to distinguish between these types of uncertainty and can suffer from probability dilution or false confidence, where critical event probabilities counter-intuitively decrease with fewer observations, as convincingly shown for the problem of satellite conjunction analysis by Balch et al. (2019). The problem of dynamic state estimation is frequently handled with probabilistic filtering techniques, where knowledge about past states of a system is propagated forward in time and updated with knowledge from new measurements. Also in this context, the susceptibility of probabilistic methods to probability dilution and false confidence can lead to systematic overconfidence in the state estimates, which can be critical in safety-relevant applications such as automation or autonomy stacks, e.g., in certain robotics applications. Possibility theory, as a theory of imprecise probabilities, i.e., a non-additive framework of uncertainty quantification, does not suffer from these issues.

Another challenge in localization arises when measurements are ambiguous, i.e., when the available sensor information allows for multiple equally plausible state estimates. This is a well-known issue to which considerable effort has been devoted to handling such ambiguity in the context of probabilistic filtering, e.g., by joint compatibility branch-and-bound techniques (Neira and Tardos, 2001), multiple hypothesis tracking (Blackman, 2004) or joint probabilistic data association filters (Bar-Shalom et al., 2009). However, these approaches require explicit algorithmic machinery, while possibility theory, by contrast, handles such ambiguity natively, as we demonstrated in the context of static state-estimation in Könecke et al. (2024).

Recent developments in possibilistic filtering techniques provide the opportunity to obtain robust state estimates that appropriately reflect the uncertainty inherent to the estimation process, as addressed by several authors over the recent years. Benferhat et al. (2000) present a qualitative counterpart to Kalman filtering, where beliefs about the state are represented by possibilistic ranking functions and iteratively updated. Matía et al. (2006) investigate a robot localization problem in which odometry and laser-based distance and angle measurements to landmarks and wall segments are processed as trapezoidal fuzzy variables and propagated through nonlinear models. They improve their approach later (Matía et al., 2021) by splitting core and support of the fuzzy variables into separate, parallel algorithms. Ferrero et al. acknowledge the limitation of the Gaussian noise assumption in Kalman filtering and propose an alternative approach using random-fuzzy variables (that we call possibility distributions) to more appropriately model measurement noise while keeping process noise probabilistic (Ferrero et al., 2019; Ferrero et al., 2021). A converse assumption is made by Hose and Hanss (2021a), who argue that primarily process noise is epistemic, i.e., non-probabilistic in nature, and present a particle-based possibilistic filtering technique applied to a nonlinear two-state batch chemical reaction problem. It is the particle-based character of this method that allows for arbitrary, non-parametric possibility distributions and nonlinear models, which is crucial for the application example presented in this contribution. The authors share the assessment that, in most engineering applications, uncertainty in the process description is often to a large part epistemic, as hard-to-model, highly detailed, time-variant, or expensive-to-evaluate deterministic, physical phenomena are often purposefully or knowingly neglected or simplified in process models. In essence, every model is, at least partly, systematically and deterministically not perfectly accurate. Hence, the contribution builds upon the foundational methodology set forward in Hose and Hanss (2021a).

This contribution consists of two main parts. First, some elementaries of possibility theory are introduced and the main steps of the possibilistic filtering procedure as originally presented by Hose and Hanss (2021a) are recalled. Then, the filter is applied to a robot localization task, where it improves on an earlier static state estimation approach by Könecke et al. (2024). Critically, in addition to the application to a real-world estimation problem, the contribution of this paper consists of some key modifications to the originally proposed filter implementation, which enable a real-time application in localization problems, e.g., in mobile robotics.

2. Foundations of Possibilistic Filtering

In this section, we provide a brief introduction to possibility theory and the key concepts of possibilistic filtering, offering a compact yet comprehensive overview of the underlying principles and methodology. In the following, the quantity of interest is modeled as a random variable $\boldsymbol{X} : \Omega \to E$

on the measurable space (Ω, Σ), where Ω is the sample space, E is the state space, and Σ is the associated σ-algebra. All subsequently introduced random variables are assumed to be well defined on their respective measurable spaces. On this basis, a possibility distribution $\pi_{\boldsymbol{X}}$ is defined as any mapping $\pi_{\boldsymbol{X}} : \Omega \to [0,1]$ satisfying $\sup_{x\in\Omega} \pi_{\boldsymbol{X}}(x) = 1$. This function assigns a degree of possibility to each element in the sample space and induces the corresponding possibility measure

$$\Pi_{\boldsymbol{X}}(A) = \sup_{\boldsymbol{x}\in A} \pi_{\boldsymbol{X}}(\boldsymbol{x}) \tag{1}$$

for any event $A \in \Sigma$. With its fundamental axioms

$$\Pi_{\boldsymbol{X}}(\emptyset) = 0, \quad \Pi_{\boldsymbol{X}}(\Omega) = 1, \quad \Pi_{\boldsymbol{X}}\left(\bigcup_{i=1}^{\infty} A_i\right) = \sup_{i\in\mathbb{N}} \Pi_{\boldsymbol{X}}(A_i), \tag{2}$$

it represents a complete, normal and maxitive measure of uncertainty. Crucially, the possibility measure $\Pi_{\boldsymbol{X}}(A)$ is not self-dual, but rather has the dual necessity measure

$$\mathrm{N}_{\boldsymbol{X}}(A) = 1 - \Pi_{\boldsymbol{X}}(\neg A), \tag{3}$$

with the corresponding counter event $\neg A$.

More precisely, a probability measure $\mathrm{P}_{\boldsymbol{X}}$ is said to be consistent with a possibility distribution $\pi_{\boldsymbol{X}}$ if the probability of any event $A \in \Sigma$ is upper bounded by the possibility of that event, i.e., if $\mathrm{P}_{\boldsymbol{X}}(A) \leq \Pi_{\boldsymbol{X}}(A)$, and lower bounded by the necessity of that event, i.e., if $\mathrm{P}_{\boldsymbol{X}}(A) \geq \mathrm{N}_{\boldsymbol{X}}(A)$. The connection between probability and possibility theory is formalized through the notion of consistency (Delgado and Moral, 1987). A probability measure $\mathrm{P}_{\boldsymbol{X}}$ is said to be consistent with a possibility distribution $\pi_{\boldsymbol{X}}$ if, for all events $A \in \Sigma$,

$$\mathrm{N}_{\boldsymbol{X}}(A) \leq \mathrm{P}_{\boldsymbol{X}}(A) \leq \Pi_{\boldsymbol{X}}(A). \tag{4}$$

In this sense, the possibility and necessity measure act as upper and lower probability bounds, respectively. The set of all probability measures consistent with a given possibility distribution is called the credal set $\mathcal{P}_{\boldsymbol{X}}$, giving a possibilistic representation of imprecise probabilistic knowledge.

Crucially, consistency induces a second probabilistic interpretation of possibilistic modeling, namely the α-cut interpretation. In particular, the α-cut $C^{\alpha}_{\pi_{\boldsymbol{X}}} = \{\boldsymbol{x} \in \Omega \mid \pi_{\boldsymbol{X}}(\boldsymbol{x}) \geq \alpha\}$ of a possibility distribution $\pi_{\boldsymbol{X}}$ admits probabilistic guarantees with respect to all consistent probability measures. More precisely, for every $\alpha \in [0,1]$,

$$\mathrm{P}_{\boldsymbol{X}}\left(C^{\alpha}_{\pi_{\boldsymbol{X}}}\right) \geq 1 - \alpha \tag{5}$$

holds for all probability measures $\mathrm{P}_{\boldsymbol{X}}$ consistent with $\pi_{\boldsymbol{X}}$. Further details on this relation can be found in Hose and Hanss (2021b).

To complete the foundations required for a compact formulation of possibility theory, it remains to derive suitable possibility distributions. Although several construction methods exist depending on the available information, we adopt the so-called Imprecise Probability-to-Possibility (IP-Π) transformation

$$\pi_{\boldsymbol{X}}(\boldsymbol{x}) = \sup_{\mathrm{P}_{\boldsymbol{X}}\in\mathcal{P}_{\boldsymbol{X}}} \mathrm{P}_{\boldsymbol{X}}\left(\{\boldsymbol{\xi} \in \Omega : \rho(\boldsymbol{\xi}) \leq \rho(\boldsymbol{x})\}\right) \quad \forall \boldsymbol{x} \in \Omega, \tag{6}$$

based on a (subjective) plausibility ordering $\rho : \Omega \to [0, 1]$. This transformation yields a possibility distribution consistent with all probability measures $\mathrm{P}_{\boldsymbol{X}}$ corresponding to the imprecise probabilistic information set $\mathcal{P}_{\boldsymbol{X}}$. In this sense, the IP-Π transformation provides the least conservative possibilistic representation that preserves consistency with the underlying imprecise probabilistic information. The objective of this paper is to apply a possibilistic filtering technique to a real-world example. To this end, the filtering problem based on possibilistic inference is briefly introduced below, together with its main recursive steps. A complete derivation of the filter can be found in (Hose and Hanss, 2021a).

Evolution For a general filter problem, consider a sequence of state variables $\boldsymbol{X}_k$. The state evolution is governed by a (potentially nonlinear) mapping

$$\boldsymbol{X}_k = \boldsymbol{f}_k(\boldsymbol{X}_{k-1}, \boldsymbol{V}_k), \tag{7}$$

where $\boldsymbol{V}_k : \Omega_v \to \mathcal{V}$ denotes an uncertain perturbation, also denoted as process noise, and $\boldsymbol{f}_k : E \times \mathcal{V} \to E$ is a measurable state transition function.

Observation Additional information about the state is obtained through observations $\boldsymbol{Y}_k : \Omega_y \to \mathcal{Y}$, which are generated through a (potentially nonlinear) observable

$$\boldsymbol{Y}_k = \boldsymbol{h}_k(\boldsymbol{X}_k, \boldsymbol{W}_k), \tag{8}$$

where $\boldsymbol{W}_k : \Omega_w \to \mathcal{W}$ represents measurement uncertainty and $\boldsymbol{h}_k : E \times \mathcal{W} \to \mathcal{Y}$ is a measurable observation function. All uncertain variables are assumed to be strongly independent. More details on this restriction can be found in (Hose and Hanss, 2019, Section 2.4).

The goal of filtering is to combine these two sources of information and recursively determine the possibility distribution of the state variable $\boldsymbol{X}_k$ given the sequence of observations $\boldsymbol{y}_{1,\ldots,k}$ and the model information, i.e., to compute $\pi_{\boldsymbol{X}_k|\boldsymbol{y}_{1,\ldots,k}}$ for all time steps $k \in \mathbb{N}$. In the context of possibilistic inference, it is of major importance to preserve consistency with the underlying imprecise probabilistic information, which is achieved by the following four filter steps: Initialization, Prediction, Inversion and Updating.

Initialization The filter is initialized with a possibility distribution $\pi_{\boldsymbol{X}_0}$ representing any kind of available initial information. If little information is available, it suffices to choose a (quasi) vacuous possibility distribution, i.e., $\pi_{\boldsymbol{X}_0}(\boldsymbol{x}) = 1$ for all $\boldsymbol{x} \in \Omega_{\mathrm{qv}} \subseteq \Omega$. For a numerically efficient implementation, we approximate the desired distributions throughout the filter by suitable particle clouds. In particular, the initial distribution $\pi_{\boldsymbol{X}_0}$ is approximated by a set of N particles $\left\{\boldsymbol{x}_0^i \in \Omega\right\}_{i=1}^N$ with associated possibilistic memberships $\left\{\mu_0^i\right\}_{i=1}^N$, which are given by the evaluation of the possibility distribution for the corresponding particle, i.e., $\mu_0^i = \pi_{\boldsymbol{X}_0}(\boldsymbol{x}_0^i)$ for all $i \in \{1, \ldots, N\}$. Generally, a suitable number of particles with appropriate sampling strategies has to be chosen to ensure a representation of the underlying possibility distribution with sufficient accuracy, allowing a reconstruction by interpolation or outer approximations.

Prediction The prediction step propagates the past distribution $\pi_{\boldsymbol{X}_{k-1}|\boldsymbol{y}_{1,\ldots,k-1}}$ through the mapping in Eq. (7) to receive the possibility function $\pi_{\boldsymbol{X}_k|\boldsymbol{y}_{1,\ldots,k-1}}$ of the current state. Theoretically, this step can be performed by the utilization of the extension principle for independent variables given in (Hose and Hanss, 2019) in combination with Eq. (7). For the particle-based implementation, however, the prediction is carried out separately for each particle. In fact, we use

$$\boldsymbol{x}_k^i = \boldsymbol{f}_k(\boldsymbol{x}_{k-1}^i, \boldsymbol{v}_k^i) \quad \forall i \in \{1, \ldots, N\}\,, \tag{9}$$

where the set $\left\{\boldsymbol{v}_k^i \in \mathcal{V}\right\}_{i=1}^N$ is a suitable sampling representation of the underlying process noise. Since this noise is usually of epistemic nature, the corresponding possibility distribution $\pi_{\boldsymbol{V}_k}$ assigns unit possibility on the whole support on its domain $\mathcal{V}$ which has to be chosen appropriately. Moreover, it is argued in Hose and Hanss (2021a) that the predicted possibilistic memberships of the particles remain approximately invariant with respect to their previous values, i.e., $\mu_k^{i,\text{pred}} \approx \mu_{k-1}^i$ for all $i \in \{1, \ldots, N\}$.

Inversion The second source of information is the observation at iteration k, which is incorporated through the inversion step. The objective is to approximate a possibility function $\pi_{\boldsymbol{X}_k|\boldsymbol{y}_k}$ on the state space based solely on the current observation and Eq. (8). To this end, we first consider $\delta_k^i = \boldsymbol{y}_k - \boldsymbol{h}_k(\boldsymbol{x}_k^i)$, i.e., the deviation between the observed value $\boldsymbol{y}_k$ and the predicted observation $\boldsymbol{h}_k(\boldsymbol{x}_k^i)$ associated with particle i, neglecting the influence of measurement noise. The membership of each particle is then determined by the possibilistic characterization of the measurement noise $\boldsymbol{W}_k$, which quantifies a degree of plausibility of the observed deviation given the available noise information. In particular, since measurement noise is often of aleatory nature, precise probabilistic information may be available. In such cases, the corresponding possibility distribution $\pi_{\boldsymbol{W}_k}$ can be obtained via a special case of Eq. (6), the so-called optimal transform (Dubois et al., 2004; Hose and Hanss, 2021b). Consequently, the inversion-based memberships are given by

$$\mu_k^{i,\text{inv}} = \pi_{\boldsymbol{W}_k}(\delta_k^i) \quad \forall i \in \{1, \ldots, N\}\,. \tag{10}$$

Update The final step consists of updating the memberships of the particles by combining the information from the prediction and inversion steps. Since we have obtained two possibility distributions $\pi_{\boldsymbol{X}_k|\boldsymbol{y}_{1,\ldots,k-1}}$ and $\pi_{\boldsymbol{X}_k|\boldsymbol{y}_k}$, defined over the same domain for the same variable, an appropriate combination rule must be applied. With the assumption of strong independence, the joint possibility distribution is computed as

$$\pi_{\boldsymbol{X}_k|\boldsymbol{y}_{1,\ldots,k}}(\boldsymbol{x}) = 1 - \left(1 - \min\left(\pi_{\boldsymbol{X}_k|\boldsymbol{y}_{1,\ldots,k-1}}(\boldsymbol{x}),\ \pi_{\boldsymbol{X}_k|\boldsymbol{y}_k}(\boldsymbol{x})\right)\right)^2 \tag{11}$$

where the independence copula is used for combining two possibility distributions, as described in (Hose and Hanss, 2019). This directly translates to the particle-based implementation, yielding the updated membership of particle i as

$$\mu_k^i = 1 - \left(1 - \min\left(\mu_k^{i,\text{pred}},\ \mu_k^{i,\text{inv}}\right)\right)^2 \quad \forall i \in \{1, \ldots, N\}\,. \tag{12}$$

Upon completion of this step, the filter is fully updated and ready for the subsequent iteration.

3. On the Resampling and Perturbation of Particles

Two issues regarding the particle-based filtering procedure require special attention. The first issue is that particles, by the nature of the state evolution prescribed by the model, but also due to the perturbation by process noise, end up in regions of the state space with very low or zero membership. In order to maintain a sufficient representation of $\pi_{\boldsymbol{X}_k}$, such particles are discarded after every measurement update. That is,

$$\left\{\boldsymbol{x}_k^i\right\}_{i=1}^{\tilde{N}} \leftarrow \left\{\boldsymbol{x}_k^i,\ \forall i \in \{1,\ldots,N\} \mid \mu_k^i > \varepsilon\right\}, \quad \tilde{N} \leq N\,, \tag{13}$$

where ε is a small threshold value, leaving $\tilde{N}$ particles. To restore the particle count, $N - \tilde{N}$ samples are drawn uniformly from the survivors and added to the particle cloud.

The second issue is the one of perturbation of the particles by process noise. The *naïve* approach of simply drawing a random perturbation from the support of the process noise, as originally proposed by Leong and Nair (2016) and adopted by Hose and Hanss (2021a), is asymptotically conservative, but requires a substantial number of particles for even this three-dimensional example, making real-time feasibility intractable. Both publications propose an improved scheme of *redundancy removal* that divides the state space into a grid and retains only the particle of highest membership in each grid cell and subsequently resamples from the survivors to restore the particle count. This way, regions of high density are thinned out, giving the remaining particles more room to spread out and cover the state space more efficiently.

What both approaches lack is the exploitation of the fact that arbitrary—that is, independent of the uncertainty-describing distribution—sampling schemes are permitted in possibility theory. This is due to the fact that, in contrast to probability theory, the quantitative information on uncertainty is not contained in the density of samples, but rather in the image of the distribution itself. The strategy proposed here exploits this degree of freedom: half the particles are drawn classically with probabilistic uniform perturbation, while the remainder are produced in pairs from source particles—one displaced outward along the vector away from the centroid of higher-membership particles until the process-noise boundary is reached, and one displaced inward along the same vector by half that distance—which is why we denote this variant as *centrifugal sampling*. This is followed by a redundancy-removal pass per grid cell, as in the originally improved scheme. Qualitative sketches of the described strategies are given in Fig. 1. An evaluation of their computational efficiency is given in the application example in the next section.

4. Application Example

The main contribution of this paper is the validation of the proposed filter in a localization scenario for a mobile robot. The robot's pose is defined as $\boldsymbol{x} = (x\ y\ \theta)^{\mathsf{T}}$, where x and y denote the Cartesian position coordinates and θ represents the orientation. In the experiments, a mobile robot equipped with four Mecanum wheels, which permit the robot to accelerate in any direction independently of its orientation, is employed, see Fig. 2 for a photograph. To its top plate, a single-board computer with a wide-angle RGB camera is attached, allowing the robot to perceive

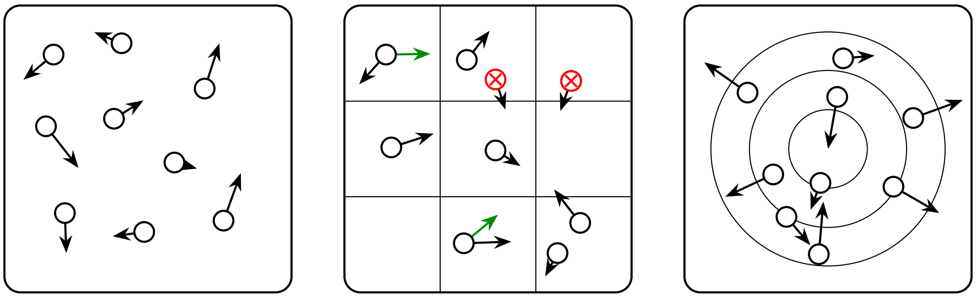

Figure 1. Illustration of different sampling strategies for perturbing particles by process noise.

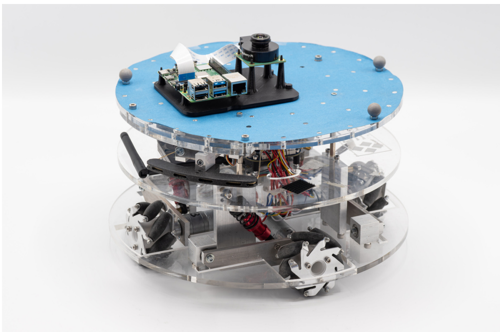

Figure 2. The mobile robot used in the experiments, equipped with four Mecanum wheels and a single-board computer with a wide-angle RGB camera.

visual landmarks in its environment via computer vision, compare Könecke et al. (2024). The setup records the directions, relative to the robot's own orientation, in which the (purposefully visually indistinguishable) landmarks are perceived. The robot's control input is given by the robot's velocity commands $\boldsymbol{u} = (v_x,\ v_y,\ \omega)^\mathsf{T}$, specifying the desired translational velocities in the body frame and the desired angular velocity, respectively. The positions of the visual landmarks are assumed to be known with respect to some inertial coordinate frame. Although the filter operates exclusively on the purely directional, visual measurement information from the indistinguishable landmarks, an additional high-precision motion capture system is available to provide ground-truth pose estimates. This enables a quantitative evaluation of the filter's performance at every time step. In this article, both the on-board directional data as well as the ground-truth data are recorded. Then, various filter variants are applied on the on-board dataset, presented time step by time step to the filter, in an offline manner. This allows to study various approaches, including less optimized ones and with very large numbers of particles, independently from on-board processing capabilities and real-time considerations. Real-time suitability can still be faithfully judged by analyzing per-step calculation times.

In general, the state filtering problem of an n-dimensional dynamical system represents a special case in which the underlying sample space is given by $\Omega = \mathbb{R}^n$, with the reduction of Σ to the corresponding Borel σ-field $\mathcal{B}(\mathbb{R}^n)$. In our case, the experimental validation is conducted on the physically meaningful subspace $\mathcal{D} \subseteq \mathbb{R}^2 \times [0, 2\pi]$, corresponding to planar position and heading angle of the robot.

Model The system dynamics are modeled in an approximative manner as a discrete-time nonlinear system of the form

$$\underbrace{\begin{pmatrix} x_k \\ y_k \\ \theta_k \end{pmatrix}}_{\boldsymbol{x}_k} = \underbrace{\begin{pmatrix} x_{k-1} \\ y_{k-1} \\ \theta_{k-1} \end{pmatrix}}_{\boldsymbol{x}_{k-1}} + \Delta \begin{pmatrix} \cos\theta_{k-1} & -\sin\theta_{k-1} & 0 \\ \sin\theta_{k-1} & \cos\theta_{k-1} & 0 \\ 0 & 0 & 1 \end{pmatrix} \underbrace{\begin{pmatrix} v_{x,k-1} \\ v_{y,k-1} \\ \omega_{k-1} \end{pmatrix}}_{\boldsymbol{u}_{k-1}} + \boldsymbol{v}_k \ , \tag{14}$$

where the pose $\boldsymbol{x}_k$ at time step k depends on the previous pose $\boldsymbol{x}_{k-1}$, the control input $\boldsymbol{u}_{k-1}$, and the process noise $\boldsymbol{v}_k$ at the current time step k. The system operates at a sampling rate of 10 Hz, i.e. with the time step $\Delta = 0.1\,\mathrm{s}$. This dynamic and uncertain system evolution $\boldsymbol{x}_k = f(\boldsymbol{x}_{k-1}, \boldsymbol{u}_{k-1}, \boldsymbol{v}_k)$ is equivalent to the general formulation in Eq. (7).

Measurement The measurement model of the system is defined as

$$\boldsymbol{y}_k = \boldsymbol{h}(\boldsymbol{x}_k, \boldsymbol{w}_k) \ , \tag{15}$$

where $\boldsymbol{y}_k$ is a set of angles that correspond to the relative angles between the robot's orientation and the lines of sight to up to four landmarks in the laboratory environment. Analogously, this formulation is consistent with Eq. (8) and provides the basis for the inversion step. An extensive description on the measurement system behind h as well as the calibration for an inverse possibilistic description $\pi_{\boldsymbol{W}}$ of possible states $\boldsymbol{x}_k$ given observed angles $\boldsymbol{y}_k$ is given in (Könecke et al., 2024). Essentially, the measurement setup is able to give a calibrated and conservative estimate $\pi_{\boldsymbol{X}_k|\boldsymbol{y}_k}(\boldsymbol{x}_k)$ of the robot's pose at each time step, but due to the ambiguity of the measurement information, the confidence sets are large, originally motivating the incorporation of the system dynamics to obtain more informative estimates.

Process Noise Process noise is modeled additively as an interval-representing possibility distribution with a half-width of 0.02 m for the positional coordinates and 0.03 rad for the orientation, which are known to safely bound the actual process noise in the system. No probabilistic distribution can be assumed for the process noise, as it is highly dependent on a multitude of factors such as the condition of the floor, the battery level, the wear of the robot's wheels, etc., which are not explicitly modeled and change over time. And while those systematic effects could be modeled in a more complex way, for example by a Gaussian process model (Eschmann et al., 2021; Eschmann et al., 2023), interval values are sufficient for now and are especially more robust to changes in the system's behavior.

Modeled possibilistically, the process noise is given by the quasi vacuous possibility distribution

$$\pi_{\boldsymbol{V}_k}(\boldsymbol{v}_k) = \begin{cases} 1 & \text{if } \boldsymbol{v}_k \in \mathcal{H} \\ 0 & \text{otherwise} \end{cases} \tag{16}$$

supported on the hyperrectangle $\mathcal{H} = [-0.02\,\mathrm{m}, 0.02\,\mathrm{m}] \times [-0.02\,\mathrm{m}, 0.02\,\mathrm{m}] \times [-0.03\,\mathrm{rad}, 0.03\,\mathrm{rad}]$.

Experiment Setup The laboratory setup is depicted in Fig. 3, illustrating a non-symmetric arrangement of four landmarks within the environment. The robot is manually driven through the laboratory by a human operator.

For the filter validation, it is assumed that there is no information about the robot's initial pose available, so that the filter is initialized with a vacuous possibility distribution, i.e., $\pi_{\boldsymbol{X}_0}(\boldsymbol{x}) = 1$ for all $\boldsymbol{x} \in \mathcal{D}$. Subsequently, the particle-based filter is applied with $N = 10^6$ particles following the steps described in the previous section, utilizing the given model and measurement information. Intuitively, each particle represents a possible pose of the robot, which is propagated through the system dynamics and updated with the measurement information at each time step. A video of the experiment together with the corresponding dataset is available via the QR code in Fig. 3, which links to (Schneider et al., 2026). To assess the performance of the filter, the empirical coverage curve given in Fig. 4 is considered, which shows the empirical coverage of the α-cuts of the estimated possibility distribution $\pi_{\boldsymbol{X}_k|\boldsymbol{y}_{1,\ldots,k}}$ with respect to the ground truth pose estimates from the motion capture system. The curve shows that the filter's estimates are well-calibrated, as the empirical coverage indicates the satisfaction of the probabilistic guarantees of α-cuts, i.e.,

$$\mathrm{P}_{\boldsymbol{X}}\left(C^{\alpha}_{\pi_{\boldsymbol{X}_k|\boldsymbol{y}_{1,\ldots,k}}}\right) \geq 1 - \alpha \quad \text{for all} \quad \alpha \in [0,1]\,. \tag{17}$$

A crucial factor influencing the filter's performance is the number of particles, which inherently represents a trade-off between computational cost and estimation accuracy. Thus, the improved sampling strategies introduced in the previous section are compared with respect to their convergence behavior in empirical coverage for varying particle numbers. The results, presented in Fig. 5, demonstrate that more advanced resampling strategies allow for a substantial reduction in the number of particles while maintaining comparable coverage performance, enabling real-time feasibility of the filter in this application example.

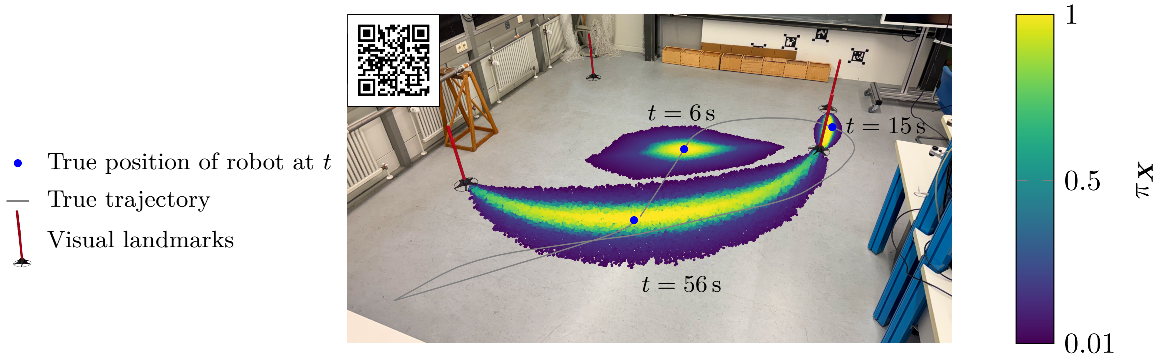

Figure 3. The laboratory setup with confidence distributions of the robot's position at three different time steps.

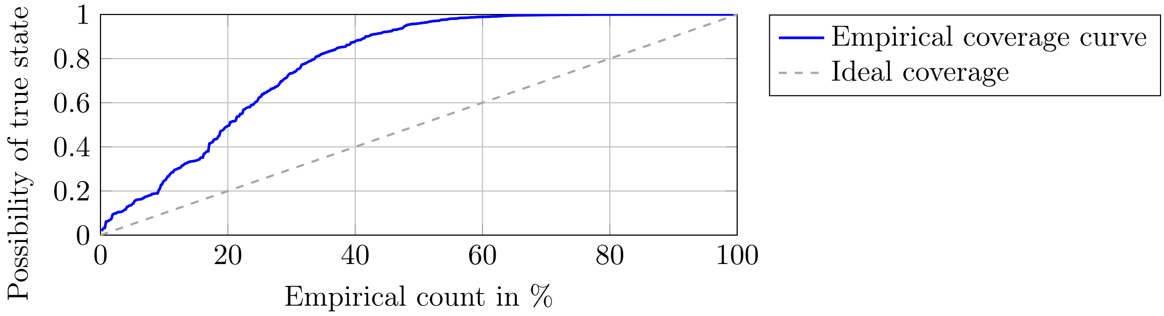

Figure 4. Empirical coverage curve.

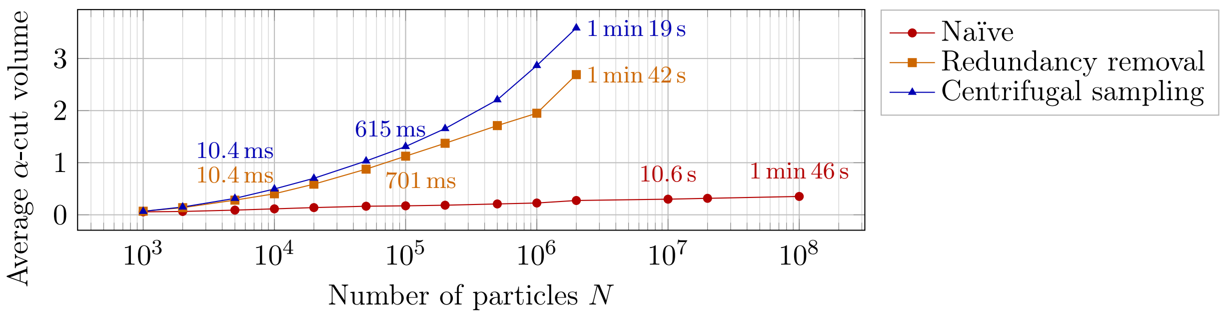

Figure 5. Average α-cut volume of the estimated possibility distribution as a function of the number of particles N, averaged across $\alpha \in \{0.05, 0.5, 0.95\}$. Average computation times per time step are given for selected runs.

5. Conclusions

We have demonstrated the applicability of a particle-based possibilistic filter to a real-world robot localization problem characterized by a challenging measurement setup and significant process noise. The results indicate that the filter successfully approximates a possibility distribution whose α-cuts satisfy the corresponding probabilistic guarantees, a property that is essential for reliable state estimation in safety-critical applications. In addition, we have introduced an improved sampling strategy for the evolution of the particle cloud. These strategies enable a substantial reduction in the number of particles while maintaining comparable estimation performance, thereby significantly enhancing the real-time feasibility of the proposed filtering approach.

Several promising directions for future research emerge from this work, including the investigation of more complex system models. In particular, questions of stability and observability constitute challenging aspects in the context of state estimation. Furthermore, if the filter is to be embedded within a feedback control loop, issues of closed-loop stability must be addressed. This represents a non-trivial problem that also demands highly efficient filter implementations to meet real-time requirements. Furthermore, set-based propagation and fusion techniques could be explored to approximate possibility distributions without relying on a particle-based implementation. Such approaches would trade estimation accuracy for improved computational efficiency. In particular, hyperellipsoidal set representations could be employed to approximate a fixed number of α-cuts of the underlying possibility distribution. In any case, the proposed filter may already be attractive in applications. For instance, it may be useful in mobile robotics in applications that require a more robust and more informed avoidance of collisions in the phase of different kinds of uncertainty, including ambiguity, compared to nominal approaches, as proposed by Rosenfelder et al. (2025), can offer.

References

Balch, M. S., R. Martin, and S. Ferson: 2019, 'Satellite Conjunction Analysis and the False Confidence Theorem'. *Proceedings of the Royal Society A* **475**(2227), 20180565.

Bar-Shalom, Y., F. Daum, and J. Huang: 2009, 'The Probabilistic Data Association Filter'. *IEEE Control Systems Magazine*.

Benferhat, S., D. Dubois, and H. Prade: 2000, 'Kalman-like Filtering and Updating in a Possibilistic Setting'. In: *14th European Conference on Artificial Intelligence*. pp. 8–12.

Blackman, S.: 2004, 'Multiple Hypothesis Tracking for Multiple Target Tracking'. *IEEE Aerospace and Electronic Systems Magazine* **19**(1), 5–18.

Delgado, M. and S. Moral: 1987, 'On the Concept of Possibility-Probability Consistency'. *Fuzzy Sets and Systems. An International Journal in Information Science and Engineering* **21**(3), 311–318.

Dubois, D., L. Foulloy, G. Mauris, and H. Prade: 2004, 'Probability-Possibility Transformations, Triangular Fuzzy Sets, and Probabilistic Inequalities'. *Reliable Computing* **10**(4), 273–297.

Eschmann, H., H. Ebel, and P. Eberhard: 2021, 'Data-Based Model of an Omnidirectional Mobile Robot Using Gaussian Processes'. *IFAC-PapersOnLine* **54**(7), 13–18.

Eschmann, H., H. Ebel, and P. Eberhard: 2023, 'Exploration-Exploitation-based Trajectory Tracking of Mobile Robots using Gaussian Processes and Model Predictive Control'. *Robotica* **41**(10), 3040–3058.

Ferrero, A., R. Ferrero, W. Jiang, and S. Salicone: 2019, 'The Kalman Filter Uncertainty Concept in the Possibility Domain'. *IEEE Transactions on Instrumentation and Measurement* **68**(11), 4335–4347.

Ferrero, A., H. V. Jetti, and S. Salicone: 2021, 'The Possibilistic Kalman Filter: Definition and Comparison with the Available Methods'. *IEEE Transactions on Instrumentation and Measurement* **70**, 1–11.

Hose, D. and M. Hanss: 2019, 'Possibilistic Calculus as a Conservative Counterpart to Probabilistic Calculus'. *Mechanical Systems and Signal Processing* **133**, 106290.

Hose, D. and M. Hanss: 2021a, 'A Recursive Formulation of Possibilistic Filters'. In: *Proceedings of the 12th International Symposium on Imprecise Probability: Theories and Applications*, Vol. 147 of *Proceedings of Machine Learning Research*. pp. 180–190.

Hose, D. and M. Hanss: 2021b, 'A Universal Approach to Imprecise Probabilities in Possibility Theory'. *International Journal of Approximate Reasoning* **133**, 133–158.

Könecke, T., H. Ebel, and M. Hanss: 2024, 'Possibilistic Robot Localization Using Visual Landmarks'. *Proceedings in Applied Mathematics and Mechanics* **25**(1), e70002.

Leong, P. H. and G. N. Nair: 2016, 'Set-Membership Filtering Using Random Samples'. In: *19th International Conference on Information Fusion (FUSION)*. pp. 1087–1094, IEEE.

Matía, F., A. Jiménez, B. M. Al-Hadithi, D. Rodríguez-Losada, and R. Galán: 2006, 'The Fuzzy Kalman Filter: State Estimation Using Possibilistic Techniques'. *Fuzzy Sets and Systems* **157**(16), 2145–2170.

Matía, F., V. Jiménez, B. P. Alvarado, and R. Haber: 2021, 'The Fuzzy Kalman Filter: Improving Its Implementation by Reformulating Uncertainty Representation'. *Fuzzy Sets and Systems* **402**, 78–104.

Neira, J. and J. Tardos: 2001, 'Data Association in Stochastic Mapping Using the Joint Compatibility Test'. *IEEE Transactions on Robotics and Automation* **17**(6), 890–897.

Rosenfelder, M., H. Carius, M. Herrmann-Wicklmayr, P. Eberhard, K. Flaßkamp, and H. Ebel: 2025, 'Efficient Avoidance of Ellipsoidal Obstacles with Model Predictive Control for Mobile Robots and Vehicles'. *Mechatronics* **110**, 103386.

Schneider, J., T. Könecke, H. Ebel, and M. Hanss: 2026, 'Visualization of Possibilistic Filtering for Reliable Robot Localization'.