Risk-Aware Active Perception and Control in Sensing-Constrained Environments

James Zhu, Thierry Simeon, Marco Cognetti

I. Introduction

Robotic manipulation in real-world environments requires operating under significant uncertainty. Objects may be partially observed due to occlusions, their physical properties may be unknown, and the environment may evolve during task execution.

This motivates the paradigm of *active perception*, in which a robot selects sensing actions (e.g., moving a camera by reorienting an end-effector) to reduce the effects of the uncertainty. Crucially, the robot's objective in this case is not to maximally reduce uncertainty, which is often unnecessary and inefficient. Instead, what matters is how uncertainty impacts task performance and safety, i.e., the *risk* induced by uncertainty.

In this work, we investigate a unified framework for *risk-aware active perception and control*, with a focus on sensing-constrained environments arising in manipulation and navigation. Our approach combines: (i) the Poisson Multi-Bernoulli Mixture (PMBM) filter to represent both detected and potentially undetected entities, and (ii) sensitivity-aware control to account for uncertainty during motion planning. Building on these components, we propose selecting sensing actions that reduce *uncertainty-induced risk*. We provide:

- A risk-aware active perception framework leveraging PMBM uncertainty representations,
- An integration with state-of-the-art sensitivity-aware controllers,
- A preliminary experimental validation in a sensing-constrained environment inspired by the game Frogger.

While demonstrated in a navigation setting, the proposed framework is relevant to manipulation scenarios involving occlusions, limited fields of view, and dynamic object interactions.

II. Sensitivity-Aware Control

Sensitivity-aware control addresses the problem of planning under parametric and state uncertainty by explicitly accounting for the effect of uncertainty on task performance [1]. Given an environment belief characterized by a mean and covariance, the controller adapts its behavior to mitigate the impact of uncertainty on safety and objective satisfaction.

This is done by propagating estimation uncertainties over a trajectory horizon and computing how closed-loop inputs and system state vary based on initial perturbations. This approach has demonstrated improved robustness in both model predictive control and learning-based policies.

In our context, sensitivity-aware control serves as the execution layer, taking uncertainty estimates provided by the perception module and producing risk-aware motion plans.

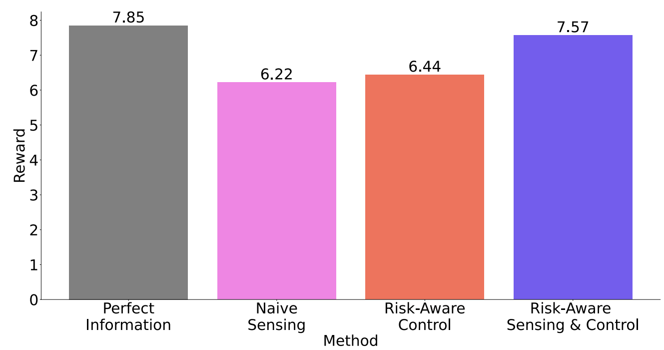

Fig. 1. Incorporating risk-awareness in both perception and control improves performance in a sensing-constrained task. A system with perfect information achieves an average reward of 7.85. Under partial observability, performance drops to 6.22. Sensitivity-aware control alone yields 6.44, while the proposed joint risk-aware perception and control approach recovers near-optimal performance (7.57).

III. The Poisson Multi-Bernoulli Mixture Filter

The Poisson Multi-Bernoulli Mixture (PMBM) filter provides a principled solution to the multi-target tracking problem, in which "a variable and unknown number of targets appear, move, and disappear from a scene of interest. At each time step, these targets are observed through noisy measurements..." [2]. Within the Random Finite Set (RFS) framework, the environment is modeled as a set with unknown and time-varying size [3].

The PMBM representation consists of two complementary components:

- A *Poisson process* representing the distribution of never-detected targets, encoding both their potential existence and state uncertainty,
- A *multi-Bernoulli mixture* capturing detected targets, including their state estimates and existence probabilities.

This dual structure enables reasoning not only about observed entities, but also about *potential* entities that may exist but have not yet been sensed. Such capability is critical in manipulation scenarios with occlusions, where unseen objects may still pose collision risks.

The authors are with the LAAS-CNRS, Université de Toulouse, CNRS, UPS, Toulouse, France. E-mail: {james.zhu, simeon, marco.cognetti}@laas.fr

This work was supported by the project ANR-20-CE33-0003 "CAMP" and by the Chaire de Professeur Junior Grant ANR-22-CPJ1-0064-01.

IV. Risk-Aware Active Perception

Standard PMBM formulations assume a fixed observation model. However, in many robotic systems, the observation model is *action-dependent*. For example, a robot may reposition its camera or manipulate an object to reveal occluded regions. This introduces a key decision problem: *where and how should the robot sense next?*

Rather than minimizing global uncertainty, we propose selecting sensing actions that reduce *task-relevant risk*. Specifically, we quantify how uncertainty in both detected and undetected targets affects future task execution, and prioritize sensing actions that mitigate high-risk interactions.

This perspective is especially relevant in manipulation, where: (i) sensing is often costly (e.g., regrasping or repositioning), (ii) full observability is rarely achievable, and (iii) only a subset of uncertainties are critical for task success.

V. Proposed Method

Step 1: Nominal Rollout

Starting from the current belief (mean estimate), we simulate a K-step nominal trajectory of the robot and environment. This rollout represents the behavior under perfect state knowledge and serves as a reference for evaluating risk.

Step 2: Risk Evaluation

We evaluate a risk metric over the nominal trajectory for both: (i) detected targets (Bernoulli components), and (ii) potential undetected targets (Poisson component).

The metric is based on the Mahalanobis distance, capturing the likelihood of unsafe interactions under uncertainty. To emphasize near-term interactions, we introduce a temporal discount factor δ, assigning higher weight to imminent risks.

Step 3: Active Sensing and Belief Update

Based on the computed risk distribution, the robot selects a sensing action that maximally reduces expected risk. New observations are incorporated into the PMBM, updating both detected and undetected target distributions.

Step 4: Sensitivity-Aware Control

The updated belief is passed to the sensitivity-aware controller, which generates a control action robust to the remaining uncertainty. Importantly, this action may deviate from the nominal rollout due to risk considerations.

VI. Toy Example

To illustrate the proposed framework, we consider a simplified environment inspired by the arcade game Frogger. A robot must traverse a grid with moving agents, under limited sensing capabilities.

At each timestep, the robot can sense only a single row with noise, reflecting sensing constraints such as limited field of view. A PPO-based controller is trained with rewards encouraging forward progress and penalizing collisions.

Under full observability, the agent achieves an average reward of 7.85. With naive sensing that senses the rows in a fixed sequence, performance drops to 6.22. Adding sensitivity-aware control yields a modest improvement to 6.44. In contrast, the proposed joint risk-aware perception and control approach achieves 7.57, approaching the fully observable upper bound (Fig. 1).

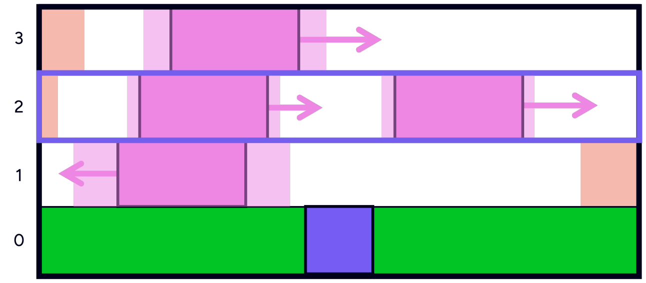

Fig. 2. The PMBM representation enables reasoning about both observed agents (pink) and potential unseen agents (orange). Sensitivity-aware control enables distributional robustness, while risk-aware perception prioritizes sensing regions that most influence future safety.

In this toy example, sensing can be chosen independently of the control action. This is in general not true for realistic robots, particularly when handling challenges like occlusions. Work from occlusion-aware autonomous driving [4] may be useful to bridging this gap.

VII. Connection to Manipulation Applications

Although evaluated in a navigation scenario, we believe the proposed framework can extend to manipulation tasks characterized by occlusions and partial observability.

Consider a grasping task where a target object is visible, but may be partially occluded by other objects. A nominal plan may suggest executing a grasp immediately. However, uncertainty in the target pose, as well as the possible existence of unseen objects along the approach trajectory, introduces a risk of collision or grasp failure.

Within our framework, this risk is explicitly quantified using the PMBM belief. If the estimated risk is low, the robot proceeds with execution. If the risk is high, the robot instead selects a sensing action aimed at reducing task-relevant uncertainty. Such actions may include repositioning the camera or performing a rearrangement action to reveal occluded regions.

This perspective provides a principled approach to active perception in manipulation, where sensing, estimation, and control are tightly integrated to optimize task performance under uncertainty.

References

[1] J. Zhu, T. Simeon, and M. Cognetti, “Robust sensitivity-aware chance-constrained mpc for efficient handling of multiple uncertainty sources,” *IEEE Robotics and Automation Letters*, vol. 10, no. 10, pp. 10 330–10 337, 2025.

[2] A. F. García-Fernández, J. L. Williams, K. Granström, and L. Svensson, “Poisson multi-bernoulli mixture filter: Direct derivation and implementation,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 54, no. 4, pp. 1883–1901, 2018.

[3] R. P. Mahler, *Advances in statistical multisource-multitarget information fusion*. Artech House, 2014.

[4] R. Trauth, K. Moller, and J. Betz, “Toward safer autonomous vehicles: Occlusion-aware trajectory planning to minimize risky behavior,” *IEEE Open Journal of Intelligent Transportation Systems*, vol. 4, pp. 929–942, 2023.