Integrating POMDP Models with PID Controllers with Application to Counting Doors on a Corridor

Rwaida Alssadi, Sara Zuiran, Jolie Elliott, Phi Doung, Chinedum Ajabor, Alwaleed AlSaleemi
Samuel Boddepalli, Vamshi Gattikoppula, Pavan Manigandla, Marius Silaghi
Florida Institute of Technology

Abstract

We propose a hybrid planning approach integrating Partially Observable Markov Decision Processes (POMDPs) with Proportional-Integral-Derivative (PID) controllers for mobile robot navigation. The method leverages PID feedback components as observation variables in a POMDP model to improve decision-making under uncertainty. This work is tested on a Create3 robot tasked with identifying and counting doors in a corridor using only onboard sensors. Our approach improves sensor noise robustness, achieving a **success rate of up to 85%**, demonstrating the feasibility of integrating AI-based planning with low-level control.

Introduction

Partially Observable Markov Decision Processes (POMDPs) models have been widely applied in robotics for decision-making under uncertainty (Kaelbling, Littman, and Cassandra 1998). Traditional POMDP formulations rely on external localization techniques or pre-defined environmental models to update belief states.

Our work extends the classical POMDP framework by incorporating Proportional-Integral-Derivative (PID) feedback as observation variables. This integration allows for real-time corrections in robot movement without relying solely on external sensors. Similar efforts to optimize POMDP-based navigation using sensor fusion have been explored (Kaelbling, Littman, and Cassandra 1998), but without leveraging PID dynamics for continuous corrections.

Unlike traditional POMDP-based robot navigation methods that rely on external localization techniques (Lauri, Hsu, and Pajarinen 2023), our approach **leverages internal PID controller states (proportional, integral, and derivative errors) as observation variables** within the POMDP.

This novel hybrid planning architecture allows robots to:

1. **Improve robustness in partially observable environments** (e.g., hallways with minimal localization infrastructure).
2. **Reduce reliance on high-fidelity sensors** by integrating PID error dynamics into the POMDP belief update process.
3. **Optimize decision-making at both high and low levels**—POMDP reasoning handles uncertainty while PID controllers provide real-time correction for smooth motion (Kaelbling, Littman, and Cassandra 1998), and input for high-level counting of doors.
4. **Demonstrate practical feasibility** with implementations balancing accuracy and computational efficiency.

By fusing **low-level control feedback with high-level AI planning**, our method bridges the gap between **reactive control systems and probabilistic reasoning**, making it suitable for real-world navigation tasks under uncertainty.

POMDP Belief Update Computation

The POMDP model maintains a **belief state** over the robot's location using Bayesian inference (Kaelbling, Littman, and Cassandra 1998). The belief update is computed as:

$$B(s') = \eta \sum_{s \in S} P(o|s')P(s'|s,a)B(s), \qquad (1)$$

where $B(s')$ is the updated belief, η is a normalization constant, $P(o|s')$ represents observation probability, $P(s'|s,a)$ is the transition probability, and $B(s)$ is the prior belief.

Integration of PID Feedback into POMDP Observations

To improve decision-making, the POMDP uses the PID feedback components as observation variables (Astrom, Hagglund, and Controllers 1995):

- **Proportional Error (e_p):** Distance from the target path, affecting wall-following precision.
- **Integral Error (e_i):** Accumulates over time, indicating drift and increasing probability of localization error.
- **Derivative Error (e_d):** Sudden changes in sensor readings, helping detect door boundaries.

These variables influence the POMDP belief updates, improving state estimation accuracy.

PID Parameter Selection Across Models

PID parameters were optimized based on empirical tuning for stability and response time.

Table 1: PID controller parameters for different models.

Model	K_p	K_i	K_d
A (Sonar-Based)	0.0065	0.00001	0.0065
B (Adaptive Proximity)	0.4	0.02	0.1
C (Multi-Sensor)	0.4	0.02	0.1

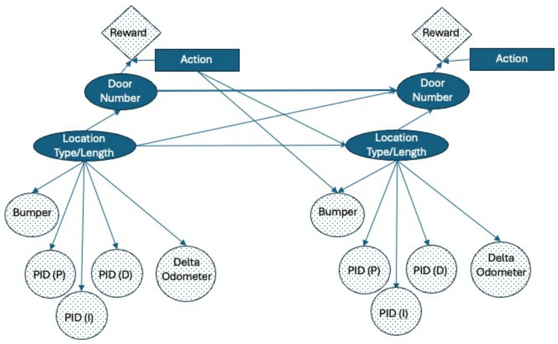

Figure 1: POMDP with PID: hidden states in filled ellipses and observation variables in light ellipses

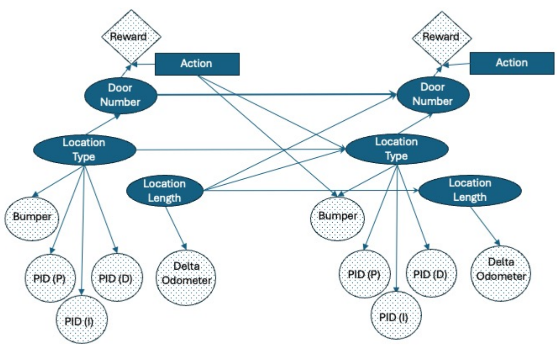

Figure 2: POMDP for Create 3 with PID: approximation with significant loss of robustness

The POMDP actions include: Unload, Return, and Ask Human Help. These actions are taken based on thresholds passed by the aforementioned linked belief state probability values. The door count used for the Unload action is computed as the expected value of the count given detection events and their probability.

Experimental Setup

To evaluate the proposed hybrid POMDP-PID approach, we conducted controlled experiments using the iRobot Create3 platform in an indoor corridor environment. The robot was tasked with identifying and counting doors while maintaining stable navigation along the wall.

The experimental conditions were standardized as follows:

- **Environment:** A 12-meter straight corridor with three evenly spaced doors on the right side.
- **Starting Position:** The robot started from a fixed initial location in all trials.
- **Lighting Conditions:** Experiments were conducted under consistent artificial lighting to minimize sensor noise.
- **Sensor Configuration:** The robot relied on infrared sensors, odometry, and collision detection.
- **Trial Count:** Each model (A, B, and C) was tested in **10 independent trials**, totaling **30 trials**.

Evaluation Metrics

Performance was assessed based on:

- **Success Rate:** Percentage of trials in which the robot correctly identified and stopped at the target door.
- **Door Count Accuracy:** Correctly counting the number of doors in the corridor.
- **Navigation Stability:** Measured by deviation from the planned path.
- **Computation Time:** Time taken by each model to process observations and decide on actions.

Conclusions

Besides proving the feasibility concept of the POMDP-PID integration, and validating its usefulness with the challenging counting doors application, three implementations were evaluated.

- **Model A** (Sonar-Based) struggled with accuracy due to sensor noise, leading to frequent miscounts.
- **Model B** (Adaptive Wall Proximity) improved stability and counting accuracy but suffered minor misalignment issues.
- **Model C** (Multi-Sensor Bayesian) achieved the highest accuracy, benefiting from sensor fusion but at the cost of increased computation time.

Model C (Multi-Sensor Bayesian) achieved the highest success rate (85%) and the counting precision (90%). This improvement is attributed to the integration of multiple sensor sources, which enhanced the robustness of POMDP belief updates. However, this model had the highest computational cost (0.25s per decision), making it less suitable for real-time applications with strict latency constraints.

References

Astrom, K. J.; Hagglund, T.; and Controllers, P. 1995. Theory, design and tuning. *Research Triangle Park: 2nd Ed. Instrumentation, Systems and Automatic Society*.

Kaelbling, L. P.; Littman, M. L.; and Cassandra, A. R. 1998. Planning and acting in partially observable stochastic domains. *Artificial Intelligence* 101(1-2):99–134.

Lauri, M.; Hsu, D.; and Pajarinen, J. 2023. Partially observable markov decision processes in robotics: A survey. *IEEE Transactions on Robotics* 39(1):21–40.