

pomdp_py: A Framework to Build and Solve POMDP Problems

Kaiyu Zheng* Stefanie Tellex

Department of Computer Science,
Brown University
Providence, RI, USA
{kzheng10, stefie10}@cs.brown.edu

Abstract

In this paper, we present `pomdp_py`, a general purpose Partially Observable Markov Decision Process (POMDP) library written in Python and Cython. Existing POMDP libraries often hinder accessibility and efficient prototyping due to the underlying programming language or interfaces, and require extra complexity in software toolchain to integrate with robotics systems. `pomdp_py` features simple and comprehensive interfaces capable of describing large discrete or continuous (PO)MDP problems. Here, we summarize the design principles and describe in detail the programming model and interfaces in `pomdp_py`. We also describe intuitive integration of this library with ROS (Robot Operating System), which enabled our torso-actuated robot to perform object search in 3D. Finally, we note directions to improve and extend this library for POMDP planning and beyond.

Introduction

Partially Observable Markov Decision Processes (POMDP) are a sequential decision-making framework suitable to model many robotics problems, from localization and mapping (Ocaña et al. 2005) to human-robot interaction (Whitney et al. 2017). Early efforts in developing tools for POMDPs attempt to separate solvers from domain description by creating specialized file formats to specify POMDPs (Cassandra 2003; APPL 2009), which are not designed for large and complex problems. Among libraries under active development, Approximate POMDP Planning Toolkit (APPL) (Somani et al. 2013) and AI-Toolkit (Bargiacchi 2014) are implemented in C++ and contain numerous solvers. However, the learning curve for these libraries is steep as C++ is less accessible to current researchers in general compared to Python (Virtanen et al. 2020). POMDPs.jl (Egorov et al. 2017) is a POMDP library with a suite of solvers and domains, written in Julia. Though promising, Julia has yet to achieve a wide recognition and creates language barrier for many researchers. POMDPy (Emami, Hamlet, and Crane 2015) is implemented purely in Python. Yet with an original focus on POMCP implementation, it assumes a blackbox world model in its POMDP interface, limiting its extensibility. Finally, a promising toolchain is to use Relational Dynamic Influence Diagram Language (RDDL) (Sanner 2010) to describe factored POMDPs and solve them

*corresponding author

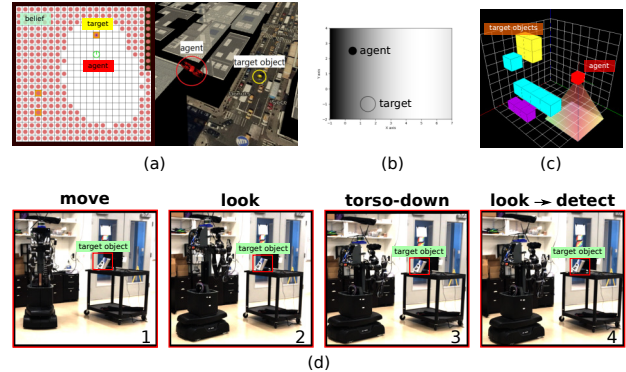


Figure 1: Example tasks implemented using `pomdp_py`. (a) Object search in 2D with a simulated drone in Unity controlled using ROS¹. (b) the light-dark domain with continuous spaces. (c) Object search in a 3D simulation with frustum-shaped field-of-view. (d) Object search in 3D implemented on a torso-actuated robot controlled using ROS; (d1-d4) shows a sequence of actions planned using an POUCT implemented in `pomdp_py`, where the robot decides to lower its torso to search, and finds the object on the table.

via ROSPlan (Cashmore et al. 2015), recently demonstrated for object fetching (Canal et al. 2019). Nevertheless, using this set of tools adds overhead of using a classical fluent-based planning paradigm, which is not required to describe and solve POMDPs in general.

This leads to our belief that there lacks a POMDP library with simple interfaces that brings together both accessibility and performance. We address this demand by presenting `pomdp_py`, a framework to build and solve POMDP problems written in Python and Cython (Behnel et al. 2011). It features simple and comprehensive interfaces to describe POMDP or MDP problems, and can be integrated with ROS (Quigley et al. 2009) intuitively through `rospy`. In the rest of this paper, we first review POMDPs, then illustrate the design principles and key features of `pomdp_py`, including integration with ROS. Finally, we note directions to improve and extend this library, in hope of cultivating an open-source community for POMDP-related research and devel-

¹We thank Rebecca Mathew for kindly providing this figure.

opment. The documentation of `pomdp.py` is available at: <https://h2r.github.io/pomdp-py/html/>. Tutorials on example domains can be found in the documentation. Figure 1 shows several different domains that are implemented on top of the `pomdp.py` framework. This library is currently actively developed as we continue our POMDP-related research.

POMDPs

POMDPs (Kaelbling, Littman, and Cassandra 1998) model sequential decision making problems where the agent must act under partial observability of the environment state (Figure 2). POMDPs consider both uncertainty in action effect (i.e. transitions) and observations, which are usually incomplete and noisy information related to the state. A POMDP is defined as a tuple $\langle \mathcal{S}, \mathcal{A}, \mathcal{O}, T, O, R, \gamma \rangle$. The problem domain is specified by $\mathcal{S}, \mathcal{A}, \mathcal{O}$: the state, action, and observation spaces. At each time step, the agent decides to take an action $a \in \mathcal{A}$, which may be sampled from $a \sim \pi(h_t, \cdot)$ according to a *policy model* $\pi(h_t, a) = \Pr(a|h_t)$. This leads to state change from s to $s' \sim T(s, a, s')$ according to the *transition model* T . Then, the agent receives an observation $o \sim O(s', a, o)$ according to the *observation model* O , and reward $r \sim R(s, a, s'), r \in \mathbb{R}$ according to the *reward model* R . Upon receiving o and r , the agent updates its history h_t and belief b_t to h_{t+1} and b_{t+1} . The goal of solving a POMDP is to find a policy $\pi(h_t, \cdot)$ which maximizes the expectation of future discounted rewards:

$$V^\pi(h_t) = \mathbb{E} \left[\sum_{k=t}^{\infty} \gamma^{k-t} R(s_k, a_k) \mid a_k = \pi(h_k, \cdot) \right],$$

where γ is the discount factor. In `pomdp.py`, a few key interfaces are defined to help organize the definition of POMDPs in a simple and consistent manner.

Solvers. Most recent POMDP solvers are *anytime algorithms* (Zilberstein 1996; Ross et al. 2008), due to the intractable computation required to solve POMDPs exactly (Madani, Hanks, and Condon 1999). There are currently two major camps of anytime solvers, point-based methods (Kurniawati, Hsu, and Lee 2008; Shani, Pineau, and Kaplow 2013) which approximates the belief space by a set of reachable α -vectors, and Monte-Carlo tree search-based methods (Silver and Veness 2010; Somani et al. 2013) that explores a subset of future action-observation sequences.

Currently, `pomdp.py` contains an implementation of POMCP and PO-UCT (Silver and Veness 2010), as well as a naive exact value iteration algorithm without pruning (Kaelbling, Littman, and Cassandra 1998). The interfaces of the library support implementation of other algorithms; We hope to cultivate a community to implement more solvers or create bridges between `pomdp.py` and other libraries.

Belief representation The partial observability of environment state implies that the agent has to maintain a posterior distribution over possible states (Thrun, Burgard, and Fox 2005). The agent should update this belief distribution through new actions and observations. The exact belief update is given by $b_{t+1}(s') = \eta \Pr(o|s', a) \sum_s \Pr(s'|s, a) b_t(s)$, where η is the normalizing factor. Hence, a naive tabular belief representation re-




Figure 2: POMDP model of agent-environment interaction. (1) Agent takes an action. (2) Environment state transitions. (3) Agent receives an observation and a reward signal. (4) Agent updates history and belief.

quires nested iterations over the state space to update the belief, which is computationally intractable in large domains. Particle belief representation is a simple and scalable belief representation which is updated through matching simulated and real observations exactly (Silver and Veness 2010). Different schemes of weighted particles have been proposed to handle large or continuous observation spaces where exact matching results in particle depletion (Sunberg and Kochenderfer 2018; Garg, Hsu, and Lee 2019).

`pomdp.py` does not commit to any specific belief representation. It provides implementations for basic belief representations and update algorithms, including tabular, particles, and multi-variate Gaussians, but more importantly allows the user to create their own new or problem-specific representation, according to the interface of a generative probability distribution.

Design Philosophy

Our goal is to design a framework that allows simple and intuitive ways of defining POMDPs at scale for both discrete and continuous domains, as well as solving them either through planning or through reinforcement learning. In addition, we implement this framework in Python and Cython to improve accessibility and prototyping efficiency without losing orders of magnitude in performance (Behnel et al. 2011; Smith 2015). We summarize the design principles behind `pomdp.py` below:

- Fundamentally, we view the POMDP scenario as the interaction between an *agent* and the *environment*, through a few important generative probability distributions (π, T, O, R or blackbox model G).
- The agent and the environment may carry different models to support learning, since for real-world problems especially in robotics, the agent generally does not know the true transition or reward models underlying the environment, and only acts based on a simplified or estimated model.
- The POMDP domain could be very large or continuous, thus explicit enumeration of elements in the spaces should be optional.
- The representation of belief distribution is decided by the

3. Environment state transitions $s_t \rightarrow s_{t+1}$ according to its transition model.
4. Agent receives an observation o_t and reward r_t from the environment.
5. Agent updates history and belief. $h_t, b_t \rightarrow h_{t+1}, b_{t+1}$, where $h_{t+1} = h_t(a_t, o_t)$.
6. Unless *termination condition* is true, repeat steps 2-5.

The `Planner` interface is as follows. The planner may be updated given a real action and a real observation, which is necessary for MCTS-based solvers.

```
class Planner:
    def plan(self, agent):
        """The agent carries the information:
        Bt, ht, O,T,R/G, pi, needed for planning"""
        raise NotImplementedError
    def update(self, agent, action, observation):
        """Updates the planner based on real action
        and observation. Updates the agent belief
        accordingly if necessary. """
        pass
```

Code Organization. In a more complicated problem such as the Light-Dark domain (Platt Jr et al. 2010) or Multi-Object Search with fan-shaped sensors (Wandzel et al. 2019), it may be tricky to organize the code base and be consistent across different problems. Below we provide a recommendation of the package structure to use `pomdp_py` to guide the development and facilitate code sharing:

```

- domain/
  - state.py          // State
  - action.py        // Action
  - observation.py   // Observation
  - ...
- models/
  - transition_model.py // TransitionModel
  - observation_model.py // ObservationModel
  - reward_model.py   // RewardModel
  - policy_model.py   // PolicyModel
  - ...
- agent/
  - agent.py         // Agent
  - ...
- env/
  - env.py          // Environment
  - ...
- problem.py       // POMDP
```

The recommendation is to separate code for domain, models, agent and environment, and have simple generic filenames. As in the above tree, files such as `state.py` or `transition_model.py` are self-evident in their role. The `problem.py` file is where the specific implementation of the POMDP class is defined, and where the logic of control flow is implemented. Refer to the Multi-Object Search example in the documentation for more detail⁴.

⁴<https://h2r.github.io/pomdp-py/html/examples.mos.html>

Object-Oriented POMDPs. OO-POMDP (Wandzel et al. 2019) is a particular kind of factored POMDP that factors the state and observation spaces into a set of n objects. For instance, $\Pr(s'|s, a) = \prod_i \Pr(s'_i|s, a)$, $i \in \{1, \dots, n\}$. The belief space is also factored, which allows the belief space to grow linearly instead of exponentially as the number of objects increases. Each object is of a certain class and has a set of attributes. The values of these attributes constitute the state of an object. In `pomdp_py`, we provide interfaces to implement OO-POMDPs, which serves as an example of extending the basic POMDP framework to create another class of model. These interfaces include `OOState`, `OOBelief`, `OOTransitionModel`, etc.

Integration with ROS. ROS (Quigley et al. 2009) is an open-source system that builds a network connecting computing stations and robots, where *nodes* interact with one another through publishing messages or making service requests. It is typical to separate nodes that manage resources and controls the robot from nodes that runs sophisticated algorithms. This is the case of `pomdp_py` as well. The POMDP-related computations can be done on a node that implements the POMDP control flow (see the six steps above). Inside this node, when an action is selected by the `Planner` (step 1), the node can publish a message to the nodes for robot control so that the robot can execute the action (step 2). The environment state automatically updates in the real world as a result of that action (step 3), and the node receives the sensor measurements or other forms of observations through subscribed topics (step 4), and performs belief update (step 5). This process is repeated until termination condition is met (step 6). ROS provides a package `rospy` which eases the integration of the POMDP control flow with the robot system.

3D Object Search with Torso-Actuated Robot. We developed a novel approach to model and solve an OO-POMDP for the task of multi-object search in 3D. Using `pomdp_py`, we implemented this approach in a simulated environment, and on Kinova MOVO, a torso-actuated mobile manipulator platform controlled with ROS. Figure 1 shows a sequence of actions that lead to object detection. In this demonstration, each planning step has a time budget of 3 seconds.

Conclusions & Future Work

We present a POMDP library, named `pomdp_py`, that brings together accessibility to programmers through Python as well as performance through Cython, with an intuitive design and straightforward integration with ROS. The programming model is designed to encourage organized development and code sharing within a community. We believe `pomdp_py` has potential to facilitate research besides POMDP planning, including reinforcement learning, transfer learning, and multi-agent systems. For example, multiple `Agent` objects could be instantiated, and different `RewardModel` classes can be created to represent different tasks. Finally, we call for support to create bridges between `pomdp_py` and other libraries to make use of existing algorithm implementations.

References

- [APPL 2009] APPL, P. 2009. Pomdpx file format (version 1.0). <https://bigbird.comp.nus.edu.sg/pmwiki/farm/appl/index.php?n=Main.PomdpXDocumentation>.
- [Bargiacchi 2014] Bargiacchi, E. 2014. Ai-toolbox: A c++ framework for mdps and pomdps with python bindings. <https://github.com/Svalorzen/AI-Toolbox>.
- [Behnel et al. 2011] Behnel, S.; Bradshaw, R.; Citro, C.; Dalcin, L.; Seljebotn, D. S.; and Smith, K. 2011. Cython: The best of both worlds. *Computing in Science & Engineering* 13(2):31–39.
- [Canal et al. 2019] Canal, G.; Cashmore, M.; Krivić, S.; Alenyà, G.; Magazzeni, D.; and Torras, C. 2019. Probabilistic planning for robotics with rosplan. In *Annual Conference Towards Autonomous Robotic Systems*, 236–250. Springer.
- [Cashmore et al. 2015] Cashmore, M.; Fox, M.; Long, D.; Magazzeni, D.; Ridder, B.; Carrera, A.; Palomeras, N.; Hurtos, N.; and Carreras, M. 2015. Rosplan: Planning in the robot operating system. In *Twenty-Fifth International Conference on Automated Planning and Scheduling*.
- [Cassandra 2003] Cassandra, A. R. 2003. Pomdp file format. <http://www.pomdp.org/code/pomdp-file-spec.html>.
- [Egorov et al. 2017] Egorov, M.; Sunberg, Z. N.; Balaban, E.; Wheeler, T. A.; Gupta, J. K.; and Kochenderfer, M. J. 2017. Pomdps.jl: A framework for sequential decision making under uncertainty. *The Journal of Machine Learning Research* 18(1):831–835.
- [Emami, Hamlet, and Crane 2015] Emami, P.; Hamlet, A. J.; and Crane, C. 2015. Pomdpy: An extensible framework for implementing pomdps in python.
- [Garg, Hsu, and Lee 2019] Garg, N. P.; Hsu, D.; and Lee, W. S. 2019. Despot- α : Online pomdp planning with large state and observation spaces. In *Robotics: Science and Systems*.
- [Kaelbling, Littman, and Cassandra 1998] 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.
- [Kurniawati, Hsu, and Lee 2008] Kurniawati, H.; Hsu, D.; and Lee, W. S. 2008. Sarsop: Efficient point-based pomdp planning by approximating optimally reachable belief spaces. In *Robotics: Science and systems*, volume 2008. Zurich, Switzerland.
- [Madani, Hanks, and Condon 1999] Madani, O.; Hanks, S.; and Condon, A. 1999. On the undecidability of probabilistic planning and infinite-horizon partially observable markov decision problems. In *AAAI/IAAI*, 541–548.
- [Ocaña et al. 2005] Ocaña, M.; Bergasa, L. M.; Sotelo, M.; and Flores, R. 2005. Indoor robot navigation using a pomdp based on wifi and ultrasound observations. In *2005 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2592–2597. IEEE.
- [Platt Jr et al. 2010] Platt Jr, R.; Tedrake, R.; Kaelbling, L.; and Lozano-Perez, T. 2010. Belief space planning assuming maximum likelihood observations.
- [Quigley et al. 2009] Quigley, M.; Conley, K.; Gerkey, B.; Faust, J.; Foote, T.; Leibs, J.; Wheeler, R.; and Ng, A. Y. 2009. Ros: an open-source robot operating system. In *ICRA workshop on open source software*, volume 3, 5. Kobe, Japan.
- [Ross et al. 2008] Ross, S.; Pineau, J.; Paquet, S.; and Chaib-Draa, B. 2008. Online planning algorithms for pomdps. *Journal of Artificial Intelligence Research* 32:663–704.
- [Sanner 2010] Sanner, S. 2010. Relational dynamic influence diagram language (rddl): Language description. *Unpublished ms. Australian National University* 32.
- [Shani, Pineau, and Kaplow 2013] Shani, G.; Pineau, J.; and Kaplow, R. 2013. A survey of point-based pomdp solvers. *Autonomous Agents and Multi-Agent Systems* 27(1):1–51.
- [Silver and Veness 2010] Silver, D., and Veness, J. 2010. Monte-carlo planning in large pomdps. In *Advances in neural information processing systems*, 2164–2172.
- [Smith 2015] Smith, K. W. 2015. *Cython: A Guide for Python Programmers*. ” O’Reilly Media, Inc.”.
- [Somani et al. 2013] Somani, A.; Ye, N.; Hsu, D.; and Lee, W. S. 2013. Despot: Online pomdp planning with regularization. In *Advances in neural information processing systems*, 1772–1780.
- [Sunberg and Kochenderfer 2018] Sunberg, Z. N., and Kochenderfer, M. J. 2018. Online algorithms for pomdps with continuous state, action, and observation spaces. In *Twenty-Eighth International Conference on Automated Planning and Scheduling*.
- [Thrun, Burgard, and Fox 2005] Thrun, S.; Burgard, W.; and Fox, D. 2005. Probabilistic robotics.
- [Virtanen et al. 2020] Virtanen, P.; Gommers, R.; Oliphant, T. E.; Haberland, M.; Reddy, T.; Cournapeau, D.; Burovski, E.; Peterson, P.; Weckesser, W.; Bright, J.; et al. 2020. Scipy 1.0: fundamental algorithms for scientific computing in python. *Nature methods* 1–12.
- [Wandzel et al. 2019] Wandzel, A.; Oh, Y.; Fishman, M.; Kumar, N.; LS, W. L.; and Tellex, S. 2019. Multi-object search using object-oriented pomdps. In *2019 International Conference on Robotics and Automation (ICRA)*, 7194–7200. IEEE.
- [Whitney et al. 2017] Whitney, D.; Rosen, E.; MacGlashan, J.; Wong, L. L.; and Tellex, S. 2017. Reducing errors in object-fetching interactions through social feedback. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, 1006–1013. IEEE.
- [Zilberstein 1996] Zilberstein, S. 1996. Using anytime algorithms in intelligent systems. *AI magazine* 17(3):73–73.