

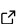
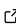

# 1 CompressedBeliefMDPs.jl: A Julia Package for 2 Solving Large POMDPs with Belief Compression

3 Logan Mondal Bhamidipaty <sup>1</sup> and Mykel J. Kochenderfer <sup>1</sup>

4 <sup>1</sup> Stanford University

DOI: [10.xxxxxx/draft](https://doi.org/10.xxxxxx/draft)

## Software

- [Review](#) 
- [Repository](#) 
- [Archive](#) 

Editor: [Open Journals](#) 

## Reviewers:

- [@openjournals](#)

Submitted: 01 January 1970

Published: unpublished

## License

Authors of papers retain copyright,  
and release the work under a  
Creative Commons Attribution 4.0  
International License ([CC BY 4.0](#)).

## 5 Summary

6 Partially observable Markov decision processes (POMDPs) are a standard mathematical model  
7 for sequential decision making under state and outcome uncertainty ([Kochenderfer et al., 2022](#)).  
8 They commonly feature in reinforcement learning research and have applications spanning  
9 medicine ([Zhou et al., 2019](#)), sustainability ([Wang et al., 2023](#)), and aerospace ([Folsom et  
10 al., 2021](#)). Unfortunately, real-world POMDPs often require bespoke solutions, because they  
11 are too large to be tractable with traditional methods ([Madani et al., 2003](#); [Papadimitriou  
12 & Tsitsiklis, 1987](#)). Belief compression ([Roy et al., 2005](#)) is a general-purpose technique  
13 that focuses planning on relevant belief states, thereby making it feasible to solve complex,  
14 real-world POMDPs more efficiently.

## Statement of Need

### 15 Research Purpose

17 [CompressedBeliefMDPs.jl](#) is a Julia package ([Bezanson et al., 2012](#)) for solving large POMDPs  
18 in the POMDPs.jl ecosystem ([Egorov et al., 2017](#)) with belief compression (described below).  
19 It offers a simple interface for efficiently sampling and compressing beliefs and for constructing  
20 and solving belief-state MDPs. The package can be used to benchmark techniques for sampling,  
21 compressing, and planning. It can also solve complex POMDPs to support applications in a  
22 variety of domains.

### 23 Relation to Prior Work

#### 24 Other Methods for Solving Large POMDPs

25 While traditional tabular methods like policy and value iteration scale poorly, there are modern  
26 methods such as point-based algorithms ([Kurniawati et al., 2008](#); [Pineau et al., 2003](#); [Smith &  
27 Simmons, 2012](#); [Spaan & Vlassis, 2005](#)) and online planners ([Kocsis & Szepesvári, 2006](#); [Ross  
28 et al., 2007](#); [Silver & Veness, 2010](#); [Somani et al., 2013](#); [Sunberg & Kochenderfer, 2018](#)) that  
29 perform well on real-world POMDPs in practice. Belief compression is an equally powerful but  
30 often overlooked alternative that is especially potent when belief is sparse.

31 [CompressedBeliefMDPs.jl](#) is a modular generalization of the original algorithm. It can be used  
32 independently or in conjunction with other planners. It also supports *both* continuous and  
33 discrete state, action, and observation spaces.

#### 34 Belief Compression

35 [CompressedBeliefMDPs.jl](#) abstracts the belief compression algorithm of Roy et al. ([2005](#))  
36 into four steps: sampling, compression, construction, and planning. The Sampler abstract

37 type handles belief sampling; the Compressor abstract type handles belief compression; the  
38 CompressedBeliefMDP struct handles constructing the compressed belief-state MDP; and  
39 the CompressedBeliefSolver and CompressedBeliefPolicy structs handle planning in the  
40 compressed belief-state MDP.

41 Our framework is a generalization of the original belief compression algorithm. Roy et al.  
42 (2005) uses a heuristic controller for sampling beliefs; exponential family principal component  
43 analysis with Poisson loss for compression (Collins et al., 2001); and local approximation  
44 value iteration for the base solver. CompressedBeliefMDPs.jl, on the other hand, is a modular  
45 framework, meaning that belief compression can be applied with *any* combination of sampler,  
46 compressor, and MDP solver.

## 47 Related Packages

48 To our knowledge, no prior Julia or Python package implements POMDP belief compression.  
49 A similar package exists for MATLAB (Chambrier, 2016), but it focuses on Poisson exponential  
50 family principal component analysis and not general belief compression.

## 51 Sampling

52 The Sampler abstract type handles sampling. CompressedBeliefMDPs.jl supports sampling  
53 with policy rollouts through PolicySampler and ExplorationSampler which wrap Policy and  
54 ExplorationPolicy from POMDPs.jl respectively. These objects can be used to collect beliefs  
55 with a random or  $\epsilon$ -greedy policy, for example.

56 CompressedBeliefMDPs.jl also supports fast *exploratory belief expansion* on POMDPs with  
57 discrete state, action, and observation spaces. Our implementation is an adaptation of  
58 Algorithm 21.13 in Kochenderfer et al. (2022). We use  $k$ -d trees (Bentley, 1975) to efficiently  
59 find the furthest belief sample.

## 60 Compression

61 The Compressor abstract type handles compression in CompressedBeliefMDPs.jl. Compressed-  
62 BeliefMDPs.jl provides seven off-the-shelf compressors:

- 63 1. Principal component analysis (PCA) (Hotelling, 1933),
- 64 2. Kernel PCA (Schölkopf et al., 1998),
- 65 3. Probabilistic PCA (Tipping & Bishop, 2002),
- 66 4. Factor analysis (Thurstone, 1931),
- 67 5. Isomap (Tenenbaum et al., 2000),
- 68 6. Autoencoder (Kramer, 1991), and
- 69 7. Variational auto-encoder (VAE) (Kingma & Welling, 2013).

70 The first four are supported through [MultivariateState.jl](#); Isomap is supported through [Mani-  
71 foldLearning.jl](#); and the last two are implemented in [Flux.jl](#) (Innes, 2018).

## 72 Compressed Belief-State MDPs

### 73 Definition

74 First, recall that any POMDP can be viewed as a belief-state MDP (Åström, 1965), where  
75 states are beliefs and transitions are belief updates (e.g., with Bayesian or Kalman filters).  
76 Formally, a POMDP is a tuple  $\langle S, A, T, R, \Omega, O, \gamma \rangle$ , where  $S$  is the state space,  $A$  is the  
77 action space,  $T : S \times A \times S \rightarrow \mathbb{R}$  is the transition model,  $R : S \times A \rightarrow \mathbb{R}$  is the reward model,  
78  $\Omega$  is the observation space,  $O : \Omega \times S \times A \rightarrow \mathbb{R}$  is the observation model, and  $\gamma \in [0, 1)$  is

79 the discount factor. The POMDP is said to induce the belief-state MDP  $\langle B, A, T', R', \gamma \rangle$ ,  
 80 where  $B$  is the POMDP belief space,  $T' : B \times A \times B \rightarrow \mathbb{R}$  is the belief update model, and  
 81  $R' : B \times A \rightarrow \mathbb{R}$  is the reward model.  $A$  and  $\gamma$  remain the same.

82 We define the corresponding *compressed belief-state MDP* (CBMDP) as  $\langle \tilde{B}, A, \tilde{T}, \tilde{R}, \gamma \rangle$   
 83 where  $\tilde{B}$  is the compressed belief space obtained from the compression  $\phi : B \rightarrow \tilde{B}$ . Then  
 84  $\tilde{R}(\tilde{b}, a) = R(\phi^{-1}(\tilde{b}), a)$  and  $\tilde{T}(\tilde{b}, a, \tilde{b}') = T(\phi^{-1}(\tilde{b}), a, \phi^{-1}(\tilde{b}'))$ . When  $\phi$  is lossy,  $\phi$  may  
 85 not be invertible. In practice, we circumvent this issue by caching items on a first-come,  
 86 first-served basis (or under an arbitrary ranking over  $B$  if the compression is parallel), so that  
 87 if  $\phi(b_1) = \phi(b_2) = \tilde{b}$  we have  $\phi^{-1}(\tilde{b}) = b_1$  if  $b_1$  was ranked higher than  $b_2$  for  $b_1, b_2 \in B$  and  
 88  $\tilde{b} \in \tilde{B}$ .

## 89 Implementation

90 The CompressedBeliefMDP struct contains a [GenerativeBeliefMDP](#), a Compressor, and a  
 91 cache  $\phi$  that recovers the original belief. The default constructor handles belief sampling,  
 92 compressor fitting, belief compressing, and cache management. Any POMDPs.jl Solver can  
 93 solve a CompressedBeliefMDP.

```
using POMDPs, POMDPModels, POMDPTools
```

```
using CompressedBeliefMDPs
```

```
# construct the CBMDP
```

```
pomdp = BabyPOMDP()
```

```
sampler = BeliefExpansionSampler(pomdp)
```

```
updater = DiscreteUpdater(pomdp)
```

```
compressor = PCACompressor(1)
```

```
cbmdp = CompressedBeliefMDP(pomdp, sampler, updater, compressor)
```

```
# solve the CBMDP
```

```
solver = MyMDPSolver()::POMDPs.Solver
```

```
policy = solve(solver, cbmdp)
```

## 94 Solvers

95 CompressedBeliefSolver and CompressedBeliefPolicy wrap the belief compression  
 96 pipeline, meaning belief compression can be applied without explicitly constructing a  
 97 CompressedBeliefMDP.

```
using POMDPs, POMDPModels, POMDPTools
```

```
using CompressedBeliefMDPs
```

```
pomdp = BabyPOMDP()
```

```
base_solver = MyMDPSolver()
```

```
solver = CompressedBeliefSolver(
```

```
    pomdp,
```

```
    base_solver;
```

```
    updater=DiscreteUpdater(pomdp),
```

```
    sampler=BeliefExpansionSampler(pomdp),
```

```
    compressor=PCACompressor(1),
```

```
)
```

```
policy = POMDPs.solve(solver, pomdp) # CompressedBeliefPolicy
```

```
s = initialstate(pomdp)
```

```
v = value(policy, s)
```

```
a = action(policy, s)
```

98 Following Roy et al. (2005), we use local value approximation as our default base solver,  
99 because it bounds the value estimation error (Gordon, 1995).

```
using POMDPs, POMDPTools, POMDPModels  
using CompressedBeliefMDPs
```

```
pomdp = BabyPOMDP()  
solver = CompressedBeliefSolver(pomdp)  
policy = solve(solver, pomdp)
```

100 To solve a continuous-space POMDP, simply swap the base solver. More details, examples,  
101 and instructions on implementing custom components can be found in the [documentation](#).

## 102 Circular Maze

103 CompressedBeliefMDPs.jl also includes the Circular Maze POMDP from Roy et al. (2005)  
104 and scripts to recreate figures from the original paper. Additional details can be found in the  
105 [documentation](#).

```
using CompressedBeliefMDPs
```

```
n_corridors = 2  
corridor_length = 100  
pomdp = CircularMaze(n_corridors, corridor_length)
```

## 106 Acknowledgments

107 We thank Arc Jamgochian, Robert Moss, Dylan Asmar, and Zachary Sunberg for their help  
108 and guidance.

## 109 References

- 110 Åström, K. J. (1965). Optimal control of Markov processes with incomplete state information.  
111 *Journal of Mathematical Analysis and Applications*, 10(1), 174–205. [https://doi.org/10.1016/0022-247X\(65\)90154-X](https://doi.org/10.1016/0022-247X(65)90154-X)  
112
- 113 Bentley, J. L. (1975). Multidimensional binary search trees used for associative searching.  
114 *Communications of the ACM*, 18(9), 509–517. <https://doi.org/10.1145/361002.361007>
- 115 Bezanson, J., Karpinski, S., Shah, V. B., & Edelman, A. (2012). *Julia: A fast dynamic*  
116 *language for technical computing*. <https://doi.org/10.48550/arXiv.1209.5145>
- 117 Chambrier, G. de. (2016). *E-PCA*. <https://github.com/gpldecha/e-pca>
- 118 Collins, M., Dasgupta, S., & Schapire, R. E. (2001). A generalization of principal components  
119 analysis to the exponential family. *Advances in Neural Information Processing Systems*.  
120 <https://doi.org/10.7551/mitpress/1120.003.0084>
- 121 Egorov, M., Sunberg, Z. N., Balaban, E., Wheeler, T. A., Gupta, J. K., & Kochenderfer, M.  
122 J. (2017). POMDPs.jl: A framework for sequential decision making under uncertainty.  
123 *Journal of Machine Learning Research*, 18(1), 831–835.
- 124 Folsom, L., Ono, M., Otsu, K., & Park, H. (2021). Scalable information-theoretic path planning  
125 for a rover-helicopter team in uncertain environments. *International Journal of Advanced*  
126 *Robotic Systems*, 18(2), 1729881421999587. <https://doi.org/10.1177/1729881421999587>
- 127 Gordon, G. J. (1995). Stable function approximation in dynamic programming. In A. Prieditis

- 128 & S. Russell (Eds.), *Machine Learning* (pp. 261–268). Morgan Kaufmann. <https://doi.org/10.1016/B978-1-55860-377-6.50040-2>  
129
- 130 Hotelling, H. (1933). Analysis of a complex of statistical variables into principal components.  
131 *Journal of Educational Psychology*, 24, 498–520. <https://doi.org/10.1037/h0070888>
- 132 Innes, M. (2018). Flux: Elegant machine learning with Julia. *Journal of Open Source Software*.  
133 <https://doi.org/10.21105/joss.00602>
- 134 Kingma, D. P., & Welling, M. (2013). Auto-encoding variational Bayes. *arXiv Preprint*  
135 *arXiv:1312.6114*. <https://doi.org/10.48550/arXiv.1312.6114>
- 136 Kochenderfer, M. J., Wheeler, T. A., & Wray, K. H. (2022). *Algorithms for Decision Making*.  
137 MIT Press. ISBN: 9780262370233
- 138 Kocsis, L., & Szepesvári, C. (2006). Bandit based Monte-Carlo planning. *European Conference*  
139 *on Machine Learning*, 282–293. [https://doi.org/10.1007/11871842\\_29](https://doi.org/10.1007/11871842_29)
- 140 Kramer, M. A. (1991). Nonlinear principal component analysis using autoassociative neural  
141 networks. *AIChE Journal*, 37(2), 233–243. <https://doi.org/10.1002/aic.690370209>
- 142 Kurniawati, H., Hsu, D. H., & Lee, W. S. (2008). SARSOP: Efficient point-based POMDP  
143 planning by approximating optimally reachable belief spaces. *Robotics: Science and*  
144 *Systems*. <https://doi.org/10.15607/RSS.2008.IV.009>
- 145 Madani, O., Hanks, S., & Condon, A. (2003). On the undecidability of probabilistic planning  
146 and related stochastic optimization problems. *Artificial Intelligence*, 147(1), 5–34. [https://doi.org/10.1016/S0004-3702\(02\)00378-8](https://doi.org/10.1016/S0004-3702(02)00378-8)  
147
- 148 Papadimitriou, C. H., & Tsitsiklis, J. N. (1987). The complexity of Markov decision processes.  
149 *Mathematics of Operations Research*, 12(3), 441–450. <https://doi.org/10.1287/moor.12.3.441>  
150
- 151 Pineau, J., Gordon, G., & Thrun, S. (2003). Point-based value iteration: An anytime algorithm  
152 for POMDPs. *International Joint Conference on Artificial Intelligence*, 1025–1030.
- 153 Ross, S., Chaib-Draa, B., & others. (2007). AEMS: An anytime online search algorithm for  
154 approximate policy refinement in large POMDPs. *IJCAI*, 2592–2598.
- 155 Roy, N., Gordon, G., & Thrun, S. (2005). Finding approximate POMDP solutions through  
156 belief compression. *Journal of Artificial Intelligence Research*, 23, 1–40. <https://doi.org/10.1613/jair.1496>  
157
- 158 Schölkopf, B., Smola, A., & Müller, K.-R. (1998). Nonlinear component analysis as a kernel  
159 eigenvalue problem. *Neural Computation*, 10(5), 1299–1319. <https://doi.org/10.1162/089976698300017467>  
160
- 161 Silver, D., & Veness, J. (2010). Monte-Carlo planning in large POMDPs. *Advances in Neural*  
162 *Information Processing Systems*, 23.
- 163 Smith, T., & Simmons, R. (2012). Point-based POMDP algorithms: Improved analysis and  
164 implementation. *arXiv Preprint arXiv:1207.1412*. <https://doi.org/10.48550/arXiv.1207.1412>  
165
- 166 Somani, A., Ye, N., Hsu, D., & Lee, W. S. (2013). DESPOT: Online POMDP planning  
167 with regularization. *Advances in Neural Information Processing Systems*, 26. <https://doi.org/10.1613/jair.5328>  
168
- 169 Spaan, M. T., & Vlassis, N. (2005). Perseus: Randomized point-based value iteration for  
170 POMDPs. *Journal of Artificial Intelligence Research*, 24, 195–220.
- 171 Sunberg, Z., & Kochenderfer, M. (2018). Online algorithms for POMDPs with continuous  
172 state, action, and observation spaces. *International Conference on Automated Planning*  
173 *and Scheduling*, 28, 259–263. <https://doi.org/10.1609/icaps.v28i1.13882>

- 174 Tenenbaum, J. B., Silva, V. de, & Langford, J. C. (2000). A global geometric framework for  
175 nonlinear dimensionality reduction. *Science*, 290(5500), 2319–2323. <https://doi.org/10.1126/science.290.5500.2319>  
176
- 177 Thurstone, L. L. (1931). Multiple factor analysis. *Psychological Review*, 38(5), 406. <https://doi.org/10.1037/h0069792>  
178
- 179 Tipping, M. E., & Bishop, C. M. (2002). Probabilistic Principal Component Analysis. *Journal*  
180 *of the Royal Statistical Society*, 61(3), 611–622. <https://doi.org/10.1111/1467-9868.00196>
- 181 Wang, Y., Zechner, M., Wen, G., Corso, A. L., Mern, J. M., Kochenderfer, M. J., & Karel  
182 Caers, J. (2023). Optimizing Carbon Storage Operations for Long-Term Safety. *arXiv*  
183 *e-Prints*, arXiv:2304.09352. <https://doi.org/10.48550/arXiv.2304.09352>
- 184 Zhou, Z., Kearnes, S., Li, L., Zare, R. N., & Riley, P. (2019). Optimization of mole-  
185 cules via deep reinforcement learning. *Scientific Reports*, 9(1). <https://doi.org/10.1038/s41598-019-47148-x>  
186

DRAFT