

Package ‘pomdpSolve’

January 7, 2023

Title Interface to 'pomdp-solve' for Partially Observable Markov Decision Processes

Version 1.0.2

Date 2022-12-27

Description Installs an updated version of 'pomdp-solve', a program to solve Partially Observable Markov Decision Processes (POMDPs) using a variety of exact and approximate value iteration algorithms. A convenient R infrastructure is provided in the separate package pomdp. Kaelbling, Littman and Cassandra (1998) <[doi:10.1016/S0004-3702\(98\)00023-X](https://doi.org/10.1016/S0004-3702(98)00023-X)>.

Classification/ACM G.4, G.1.6, I.2.6

URL <https://github.com/mhahsler/pomdpSolve>

BugReports <https://github.com/mhahsler/pomdpSolve/issues>

Depends R (>= 3.5.0)

Suggests pomdp

Encoding UTF-8

License GPL (>= 3)

Copyright pomdp-solve is Copyright (C) Anthony R. Cassandra; LASPack is Copyright (C) Tomas Skalicky; lp-solve is Copyright (C) Michel Berkelaar, Kjell Eikland, Peter Notebaert; all other code is Copyright (C) Michael Hahsler.

RoxygenNote 7.2.3

NeedsCompilation yes

Author Michael Hahsler [aut, cph, cre]
(<<https://orcid.org/0000-0003-2716-1405>>),
Anthony R. Cassandra [aut, cph]

Maintainer Michael Hahsler <mhahsler@lyle.smu.edu>

Repository CRAN

Date/Publication 2023-01-07 15:30:10 UTC

R topics documented:

pomdpSolve-package	2
find_pomdp_solve	2

Index	5
--------------	----------

pomdpSolve-package	<i>pomdpSolve: Interface to 'pomdp-solve' for Partially Observable Markov Decision Processes</i>
--------------------	--

Description

Installs an updated version of 'pomdp-solve', a program to solve Partially Observable Markov Decision Processes (POMDPs) using a variety of exact and approximate value iteration algorithms. A convenient R infrastructure is provided in the separate package **pomdp** ([pomdp::pomdp-package](#)).

Key functions

- Find and run the pomdp-solve executable using [find_pomdp_solve\(\)](#).
- Define a POMDP using [pomdp::POMDP](#)
- Solve a POMDP using [pomdp::solve_POMDP\(\)](#)

Author(s)

Michael Hahsler

find_pomdp_solve	<i>Find the executable for 'pomdp-solve'</i>
------------------	--

Description

Find the pomdp-solve executable to solve Partially Observable Decision Processes (POMDPs) (Kaelbling et al, 1998) installed by the **pomdpSolve** package.

Usage

find_pomdp_solve()

Details

This package only provides a direct interface to the executable. A more convenient and powerful interface is provided by the function `pomdp::solve_POMDP()` in package **pomdp**.

The executable of pomdp-solve in this direct interface needs to be called with `system2()` and runs in a separate process. This way, a failure in the solver will not compromise the R session. pomdp-solve creates files with the value function and the policy graph (see below).

Value Function

The value function is returned as files with the extension `.alpha` in the format:

```
A
V1 V2 V3 ... VN
```

```
A
V1 V2 V3 ... VN
```

...

Where `A` is an action number and the `V1` through `VN` are real values representing the components of a particular vector that has the associated action. The action number is the 0-based index of the action as specified in the input POMDP file. The vector represents the coefficients of a hyperplane representing one facet of the piecewise linear and convex (PWLC) value function. Note that the length of the lists needs to be equal to the number of states in the POMDP.

Policy Graph

The policy graph is returned as a file with the extension `.pg`. Each line of the file represents one node of the policy graph and its contents are:

```
N A Z1 Z2 Z3 ...
...
```

Here `N` is a node ID giving the node a unique name, numbered sequentially and lining up sequentially with the value function vectors in the corresponding output `.alpha` file above.

The `A` is the action number defined for this node; it is an integer referring to the the POMDP file actions by its 0-based index number. These are followed by a list of node IDs, one for each observation. Thus the list will have a length equal to the number of observations in the POMDP. This list specifies the transitions in the policy graph. The `n`th number in the list will be the index of the node that follows this one when the observation received is `n`.

Details about the file formats and pomdp-solve can be found in the References section.

Value

`find_pomdp_solve()` returns the path to the 'pomdp-solve' executable as a string or stops with an error.

References

- Kaelbling, L.P., Littman, M.L., Cassandra, A.R. (1998). Planning and acting in partially observable stochastic domains. *Artificial Intelligence*. **101** (1–2): 99-134. doi:10.1016/S00043702(98)00023X
- Anthony R. Cassandra, pomdp-solve documentation, <https://www.pomdp.org/code/index.html>

Examples

```
# find the location of the pomdp-solve executable
find_pomdp_solve()

# get pomdp-solve options
system2(find_pomdp_solve(), args = "-h")

# solve a POMDP file that ships with this package in a temporary directory
tmp <- tempdir()
pomdp_file <- file.path(tmp, "tiger.aaai.POMDP")
file.copy(system.file("tiger.aaai.POMDP", package = "pomdpSolve"), pomdp_file)

system2(find_pomdp_solve(), args = paste("-pomdp", pomdp_file, "-method grid"))
dir(tmp)

# read the raw policy graph
read.table(file.path(tmp, "tiger.aaai-0.pg"))

# read the raw value function
readLines(file.path(tmp, "tiger.aaai-0.alpha"))
```

Index

[find_pomdp_solve](#), [2](#)
[find_pomdp_solve\(\)](#), [2](#)

[pomdp-solve \(find_pomdp_solve\)](#), [2](#)
[pomdp::POMDP](#), [2](#)
[pomdp::pomdp-package](#), [2](#)
[pomdp::solve_POMDP\(\)](#), [2](#), [3](#)
[pomdpsolve \(find_pomdp_solve\)](#), [2](#)
[pomdpSolve-package](#), [2](#)

[system2\(\)](#), [3](#)