

Package: MDP2 (via r-universe)

October 4, 2026

Type Package

Title Markov Decision Processes (MDPs)

Version 3.0.0.11

Author Lars Relund Nielsen [aut, cre]
(<https://orcid.org/0000-0002-4802-3071>)

Maintainer Lars Relund Nielsen <lars@relund.dk>

Description Create and optimize (semi) MDPs with discrete time steps and state space. Both hierarchical and ordinary-traditional MDPs can be modeled.

License GPL (>= 3)

URL <https://relund.github.io/mdp/>, <https://github.com/relund/mdp/>,
<http://relund.github.io/mdp/>

BugReports <https://github.com/relund/mdp/issues>

Depends R (>= 4.1.0)

Imports diagram, dplyr, shape, stringr, tidyr, magrittr, methods,
purrr, rlang, tibble, Rcpp (>= 0.11.5)

Suggests knitr, Matrix, rmarkdown, testthat (>= 3.0.0), readr, xml2,
covr, roxygen2

LinkingTo Rcpp, RcppArmadillo

VignetteBuilder knitr

Encoding UTF-8

Language en-US

Roxygen list(markdown = TRUE)

Config/roxygen2/version 8.0.0

Config/testthat/edition 3

Config/pak/sysreqs libicu-dev

Repository <https://relund.r-universe.dev>

Date/Publication 2026-08-05 10:38:10 UTC

RemoteUrl <https://github.com/relund/mdp>

RemoteRef HEAD

RemoteSha cc69a53dcf335464bc4e683adb21545447a2f7a7

Contents

binary_action_writer	2
binary_mdp_writer	6
convert_binary_to_hmp	11
convert_hmp_to_binary	12
get_bin_info_actions	13
get_bin_info_states	16
get_hypergraph	17
get_info	21
get_policy	25
get_rpo	30
get_steady_state_pr	31
get_w_idx	32
hmp_mdp_writer	32
load_mdp	36
memory_mdp_writer	40
plot.HMDP	43
plot_hypergraph	47
random_hmdp	52
run_calc_weights	53
run_policy_ite_ave	58
run_policy_ite_discount	59
run_value_ite	60
save_mdp	65
set_action_weight	65
set_policy	66
set_transition_weight	70

Index	71
--------------	-----------

binary_action_writer	<i>Function for writing actions of a HMDP model to binary files. The function defines sub-functions which can be used to define actions saved in a set of binary files. It is assumed that the states have been defined using binary_mdp_writer and that the id of the states is known (can be retrieved using e.g. state_idx_df).</i>
----------------------	--

Description

Binary files are efficient for storing large models. Compared to the HMP (XML) format the binary files use less storage space and loading the model is faster.

Usage

```
binary_action_writer(
  prefix = "",
  bin_names = c("actionIdx.bin", "actionIdxLbl.bin", "actionWeight.bin",
    "actionWeightLbl.bin", "transProb.bin", "transWeight.bin", "transWeightLbl.bin"),
  append = TRUE
)
```

Arguments

<code>prefix</code>	A character string with the prefix added to <code>bin_names</code> .
<code>bin_names</code>	A character vector of length 5 giving the names of the binary files storing the model.
<code>append</code>	Logical indicating whether should keep the currents actions (default - TRUE) defined or delete them and start over (FALSE).

Details

The returned writer exposes these functions:

- `set_weights(labels, ...)`: sets the labels of the weights used in the actions. `labels` is a vector of label names. `...` is currently ignored. Call this before building the model.
- `add_action(label = NULL, s_idx, weights, prob, ...)`: adds an action. `s_idx` is the id of the state defining the action. `weights` must be a vector of action weights. `prob` is a matrix (`s_idx`, `pr`) where the first column contains the id of the transition state; see the description of `actionIdx.bin` below, where scope is assumed to be 3. `...` is currently ignored.
- `end_action()`: ends an action.
- `close_writer()`: closes the writer. Call this when the model description is finished.

Five binary files are created:

- `actionIdx.bin`: integers defining all actions in the format `s_idx scope idx scope idx scope idx -1 s_idx scope`. `s_idx` corresponds to the index or line number in `stateIdx.bin`, starting from 0. The following (`scope`, `idx`) pairs indicate possible transitions. Scope can take four values:
 - 2: a transition to a child process, at stage zero in the child process.
 - 1: a transition to the next stage in the current process.
 - 0: a transition to the next stage in the father process.
 - 3: a transition to a state specified by its state `s_idx`.

For example, if `scope = 1` and `idx = 2`, the transition is to state number 3 at the next stage in the current process. If `scope = 3` and `idx = 5`, the transition is to the state specified at line 6 in `stateIdxLbl.bin`. This is useful when considering shared child processes.

- `actionIdxLbl.bin`: character data in the format `a_idx label a_idx label ....`. Here `a_idx` corresponds to the index or line number in `actionIdx.bin`, starting from 0. No delimiter is used.
- `actionWeight.bin`: doubles containing action weights in the format `"c1 c2 c3 c1 c2 c3 ..."`, assuming three weights for each action.

- `actionWeightLbl.bin`: character data containing the weight labels in the format `label1 label2 label3`, assuming three weights for each action.
- `transProb.bin`: doubles containing the transition probabilities defined in `actionIdx.bin`. The format is "`p1 p2 p3 -1 p1 -1 p1 p2 -1 ...`". Here `-1` indicates that a new action is considered.

Value

A list of functions.

Note

Note all indexes are starting from zero (C/C++ style).

Examples

```
## Use temp dir
wd <- setwd(tempdir())

# Create a small HMDP with two levels
w<-binary_mdp_writer()
w$set_weights(c("Duration","Net reward","Items"))
w$process()
  w$stage()
    w$state(label="M0")
      w$action(label="A0",weights=c(0,0,0),prob=c(2,0,1))
        w$process()
          w$stage()
            w$state(label="D")
              w$action(label="A0",weights=c(0,0,1),prob=c(1,0,0.5,1,1,0.5))
                w$end_action()
              w$end_state()
            w$end_stage()
          w$stage()
            w$state(label="C0")
              w$action(label="A0",weights=c(0,0,0),prob=c(1,0,1))
                w$end_action()
              w$action(label="A1",weights=c(1,2,1),prob=c(1,0,0.5,1,1,0.5))
                w$end_action()
              w$end_state()
            w$state(label="C1")
              w$action(label="A0",weights=c(0,0,0),prob=c(1,0,1))
                w$end_action()
              w$action(label="A1",weights=c(1,2,1),prob=c(1,0,0.5,1,1,0.5))
                w$end_action()
              w$end_state()
            w$end_stage()
          w$stage()
            w$state(label="C0")
              w$action(label="A0",weights=c(1,4,0),prob=c(0,0,1))
                w$end_action()
              w$end_state()
```

```

        w$state(label="C1")
        w$action(label="A0",weights=c(1,4,0),prob=c(0,0,1))
        w$end_action()
        w$end_state()
    w$end_stage()
w$end_process()
w$end_action()
w$action(label="A1",weights=c(0,0,0),prob=c(2,0,1))
w$process()
    w$stage()
        w$state(label="D")
        w$action(label="A0",weights=c(0,0,1),prob=c(1,0,1))
        w$end_action()
        w$end_state()
    w$end_stage()
w$stage()
    w$state(label="C0")
    w$action(label="A0",weights=c(0,0,0),prob=c(1,0,1))
    w$end_action()
    w$action(label="A1",weights=c(1,2,1),prob=c(1,0,0.5,1,1,0.5))
    w$end_action()
    w$end_state()
w$end_stage()
w$stage()
    w$state(label="C0")
    w$action(label="A0",weights=c(1,4,0),prob=c(0,0,1))
    w$end_action()
    w$end_state()
    w$state(label="C1")
    w$action(label="A0",weights=c(1,4,0),prob=c(0,0,1))
    w$end_action()
    w$action(label="A1",weights=c(0,10,5),prob=c(0,0,0.5,0,1,0.5))
    w$end_action()
    w$end_state()
w$end_stage()
w$end_process()
w$end_action()
w$end_state()
w$state(label="M1")
    w$action(label="A0",weights=c(0,0,0),prob=c(2,0,1))
    w$process()
        w$stage()
            w$state(label="D")
            w$action(label="A0",weights=c(0,0,1),prob=c(1,0,0.5,1,1,0.5))
            w$end_action()
            w$end_state()
        w$end_stage()
    w$stage()
        w$state(label="C0")
        w$action(label="A0",weights=c(0,0,0),prob=c(1,0,1))
        w$end_action()
        w$end_state()
    w$state(label="C1")

```

```

        w$action(label="A0",weights=c(0,0,0),prob=c(1,0,1))
        w$end_action()
    w$end_state()
w$end_stage()
w$stage()
    w$state(label="C0")
        w$action(label="A0",weights=c(1,4,0),prob=c(0,0,1))
        w$end_action()
    w$end_state()
    w$state(label="C1")
        w$action(label="A0",weights=c(1,4,0),prob=c(0,0,1))
        w$end_action()
    w$end_state()
w$end_stage()
w$end_process()
w$end_action()
w$end_state()
w$end_stage()
w$end_process()
w$close_writer()

## Info about the binary files (don't have to load the model first)
get_bin_info_states()
get_bin_info_actions()

## reset working dir
setwd(wd)

```

binary_mdp_writer	<i>Function for writing an HMDP model to binary files. The function defines sub-functions which can be used to define an HMDP model saved in a set of binary files.</i>
-------------------	---

Description

Binary files are efficient for storing large models. Compared to the HMP (XML) format the binary files use less storage space and loads the model faster.

Usage

```

binary_mdp_writer(
  prefix = "",
  bin_names = c("stateIdx.bin", "stateIdxLbl.bin", "actionIdx.bin", "actionIdxLbl.bin",
    "actionWeight.bin", "actionWeightLbl.bin", "transProb.bin", "externalProcesses.bin",
    "transWeight.bin", "transWeightLbl.bin"),
  get_log = TRUE
)

```

Arguments

prefix	A character string with the prefix added to bin_names.
bin_names	A character vector giving the names of the binary files storing the model.
get_log	Output log text.

Details

The returned writer exposes these functions:

- `set_weights(labels, ...)`: sets the labels of the weights used in the actions. `labels` is a vector of label names. `...` is currently ignored. Call this before building the model.
- `process()`: starts a (sub)process. It may also be used to specify a traditional MDP using matrices in MDPtoolbox style. In that style, `p` is a list of matrices, one per action, each of size $S \times S$ where S is the number of states. Each used row must sum to one, or all entries in a row must be zero if unused. `r` is a matrix of size $S \times A$, where A is the number of actions, and `d` is a matrix of size $S \times A$ with durations. If `d` is omitted, all durations are assumed to be 1.
- `end_process()`: ends a (sub)process.
- `stage(label = NULL)`: starts a stage. `label` is currently unused in the binary format.
- `end_stage()`: ends a stage.
- `state(label = NULL)`: starts a state and returns, invisibly, the state id. That id can later be referenced with scope 3.
- `end_state()`: ends a state.
- `action(scope = NULL, id = NULL, pr = NULL, prob = NULL, weights, trans_weights = NULL, label = NULL, end = FALSE, ...)`: starts an action. `weights` must be a vector of action weights. `trans_weights` must contain transition weights ordered by transition, with all transition weight labels for the first transition followed by all labels for the second transition, and so on. Transition probabilities can be entered in two ways:
 1. `prob` contains triples (`scope`, `id`, `pr`).
 2. `id` and `pr` are vectors of equal length. If `scope` is omitted, all scopes default to 1.
 See the description of `actionIdx.bin` below. If `end = TRUE`, calling `end_action()` is not necessary. `...` is currently ignored.
- `end_action()`: ends an action. Do not use this if `end = TRUE` was used when the action was specified.
- `include_process(prefix, label = NULL, weights, prob, term_states, trans_weights = NULL)`: includes an external process. External processes are loaded into memory only when needed, which helps with large models. `prefix` is the external process prefix. `weights` must be a vector of action weights, and `prob` must contain triples (`scope`, `idx`, `pr`); see the description of `actionIdx.bin` below. `term_states` must specify the number of states at the last stage in the external process. Inside an `include_process ... end_include_process` block, you must specify the father jump actions of the last stage in the external process. The external process is represented by its first and last stage together with its jump actions. The function returns, invisibly, the state ids of the first stage in the external process, which can later be referenced with scope 3.

- `end_include_process()`: ends an `include_process` block.
- `close_writer()`: closes the writer. Call this when the model description is finished.

Ten binary files are created:

- `stateIdx.bin`: integers defining all states in the format `"n0 s0 -1 n0 s0 a0 n1 s1 -1 n0 s0 a0 n1 s1 a1 n2 s2 -1 n0 s0 ..."`. Here `-1` indicates that a new state is considered.
- `stateIdxLbl.bin`: character data in the format `s_idx label s_idx label ....`. Here `s_idx` corresponds to the index or line number in `stateIdx.bin`, starting from 0. No delimiter is used.
- `actionIdx.bin`: integers defining all actions in the format `s_idx scope idx scope idx scope idx -1 s_idx scope s_idx scope ...`. Here `s_idx` corresponds to the index or line number in `stateIdx.bin`, starting from 0. The following `(scope, idx)` pairs indicate possible transitions. Scope can take four values:
 - 2: a transition to a child process, at stage zero in the child process.
 - 1: a transition to the next stage in the current process.
 - 0: a transition to the next stage in the father process.
 - 3: a transition to a state specified by its state `s_idx`.

For example, if `scope = 1` and `idx = 2`, the transition is to state number 3 at the next stage in the current process. If `scope = 3` and `idx = 5`, the transition is to the state specified at line 6 in `stateIdxLbl.bin`. This is useful when considering shared child processes.

- `actionIdxLbl.bin`: character data in the format `a_idx label a_idx label ....`. Here `a_idx` corresponds to the index or line number in `actionIdx.bin`, starting from 0. No delimiter is used.
- `actionWeight.bin`: doubles containing action weights in the format `"c1 c2 c3 c1 c2 c3 ..."`, assuming three weights for each action.
- `actionWeightLbl.bin`: character data containing the weight labels in the format `label1 label2 label3`, assuming three weights for each action.
- `transProb.bin`: doubles containing transition probabilities defined in `actionIdx.bin`. The format is `"p1 p2 p3 -1 p1 -1 p1 p2 -1 ..."`. Here `-1` indicates that a new action is considered.
- `externalProcesses.bin`: character data containing links to external processes in the format `stage_str prefix stage_str prefix ....`. Here `stage_str` corresponds to the stage index, for example `n0 s0 a0 n1`, of the stage corresponding to the first stage in the external process, and `prefix` is the external process prefix. No delimiter is used.
- `transWeight.bin`: doubles containing transition weights in the format `"t11 t12 t21 t22 -1 ..."`, assuming two transition weights for each transition and two transitions in the first action.
- `transWeightLbl.bin`: character data containing the transition weight labels.

Value

A list of functions.

Note

Note all indexes are starting from zero (C/C++ style).

Examples

```

## Use temp dir
wd <- setwd(tempdir())

# Create a small HMDP with two levels
w<-binary_mdp_writer()
w$set_weights(c("Duration","Net reward","Items"))
w$process()
  w$stage()
    w$state(label="M0")
      w$action(label="A0",weights=c(0,0,0),prob=c(2,0,1))
        w$process()
          w$stage()
            w$state(label="D")
              w$action(label="A0",weights=c(0,0,1),prob=c(1,0,0.5,1,1,0.5))
                w$end_action()
              w$end_state()
            w$end_stage()
          w$stage()
            w$state(label="C0")
              w$action(label="A0",weights=c(0,0,0),prob=c(1,0,1))
                w$end_action()
              w$action(label="A1",weights=c(1,2,1),prob=c(1,0,0.5,1,1,0.5))
                w$end_action()
              w$end_state()
            w$state(label="C1")
              w$action(label="A0",weights=c(0,0,0),prob=c(1,0,1))
                w$end_action()
              w$action(label="A1",weights=c(1,2,1),prob=c(1,0,0.5,1,1,0.5))
                w$end_action()
              w$end_state()
            w$end_stage()
          w$stage()
            w$state(label="C0")
              w$action(label="A0",weights=c(1,4,0),prob=c(0,0,1))
                w$end_action()
              w$end_state()
            w$state(label="C1")
              w$action(label="A0",weights=c(1,4,0),prob=c(0,0,1))
                w$end_action()
              w$end_state()
            w$end_stage()
          w$end_process()
        w$end_action()
      w$action(label="A1",weights=c(0,0,0),prob=c(2,0,1))
        w$process()
          w$stage()
            w$state(label="D")
              w$action(label="A0",weights=c(0,0,1),prob=c(1,0,1))
                w$end_action()
              w$end_state()
            w$end_stage()
          
```

```

w$stage()
  w$state(label="C0")
    w$action(label="A0",weights=c(0,0,0),prob=c(1,0,1))
    w$end_action()
    w$action(label="A1",weights=c(1,2,1),prob=c(1,0,0.5,1,1,0.5))
    w$end_action()
  w$end_state()
w$end_stage()
w$stage()
  w$state(label="C0")
    w$action(label="A0",weights=c(1,4,0),prob=c(0,0,1))
    w$end_action()
  w$end_state()
  w$state(label="C1")
    w$action(label="A0",weights=c(1,4,0),prob=c(0,0,1))
    w$end_action()
    w$action(label="A1",weights=c(0,10,5),prob=c(0,0,0.5,0,1,0.5))
    w$end_action()
  w$end_state()
w$end_stage()
w$end_process()
w$end_action()
w$end_state()
w$state(label="M1")
  w$action(label="A0",weights=c(0,0,0),prob=c(2,0,1))
  w$process()
    w$stage()
      w$state(label="D")
        w$action(label="A0",weights=c(0,0,1),prob=c(1,0,0.5,1,1,0.5))
        w$end_action()
      w$end_state()
    w$end_stage()
  w$stage()
    w$state(label="C0")
      w$action(label="A0",weights=c(0,0,0),prob=c(1,0,1))
      w$end_action()
    w$end_state()
    w$state(label="C1")
      w$action(label="A0",weights=c(0,0,0),prob=c(1,0,1))
      w$end_action()
    w$end_state()
  w$end_stage()
w$stage()
  w$state(label="C0")
    w$action(label="A0",weights=c(1,4,0),prob=c(0,0,1))
    w$end_action()
  w$end_state()
  w$state(label="C1")
    w$action(label="A0",weights=c(1,4,0),prob=c(0,0,1))
    w$end_action()
  w$end_state()
w$end_stage()
w$end_process()

```

```

        w$end_action()
        w$end_state()
        w$end_stage()
        w$end_process()
        w$close_writer()

## Info about the binary files (don't have to load the model first)
get_bin_info_states()
get_bin_info_actions()

## reset working dir
setwd(wd)

```

`convert_binary_to_hmp` *Convert a HMDP model stored in binary format to a hmp (XML) file. The function simply parse the binary files and create hmp files using the [hmp_mdp_writer\(\)](#).*

Description

Convert a HMDP model stored in binary format to a hmp (XML) file. The function simply parse the binary files and create hmp files using the [hmp_mdp_writer\(\)](#).

Usage

```

convert_binary_to_hmp(
  prefix = "",
  bin_names = c("stateIdx.bin", "stateIdxLbl.bin", "actionIdx.bin", "actionIdxLbl.bin",
    "actionWeight.bin", "actionWeightLbl.bin", "transProb.bin"),
  out = paste0(prefix, "converted.hmp"),
  duration = 1,
  get_log = TRUE
)

```

Arguments

<code>prefix</code>	A character string with the prefix which will be added to the binary files.
<code>bin_names</code>	A character vector of length 7 giving the names of the binary files storing the model.
<code>out</code>	The name of the HMP file (e.g. <code>r.hmp</code>).
<code>duration</code>	Weight number storing the duration (NULL if none).
<code>get_log</code>	Output log text.

Value

NULL (invisible).

Note

Note all indexes are starting from zero (C/C++ style).

See Also

[convert_hmp_to_binary\(\)](#).

Examples

```
## Set working dir
fDir <- system.file("models", package = "MDP2")
wd <- setwd(tempdir())
## Convert the machine example to a hmp file
prefix1 <- paste0(fDir, "/machine1_")
get_bin_info_states(prefix1)
convert_binary_to_hmp(prefix1, duration = NULL, out = "machine1_converted.hmp")
# have a look at the hmp file
cat(readr::read_file("machine1_converted.hmp"))

## Convert the machine example hmp file to binary files
convert_hmp_to_binary(file = paste0(fDir, "/machine1.hmp"), prefix = "machine_cov_")
get_bin_info_states(prefix = "machine_cov_")
## Convert the machine example with a single dummy node to a hmp file
#convert_binary_to_hmp("machine2_") # error since using scope = 3 not supported in hmp files

## Reset working dir
setwd(wd)
```

`convert_hmp_to_binary` *Convert a HMDP model stored in a hmp (xml) file to binary file format.*

Description

The function simply parse the hmp file and create binary files using the [binary_mdp_writer\(\)](#).

Usage

```
convert_hmp_to_binary(file, prefix = "", get_log = TRUE)
```

Arguments

<code>file</code>	The name of the HMP file (e.g. <code>r.hmp</code>).
<code>prefix</code>	A character string with the prefix which will be added to the binary files.
<code>get_log</code>	Output log text.

Value

NULL (invisible).

Note

Note all indexes are starting from zero (C/C++ style).

See Also

`binary_mdp_writer()`.

Examples

```
## Set working dir
fDir <- system.file("models", package = "MDP2")
wd <- setwd(tempdir())
## Convert the machine example to a hmp file
prefix1 <- paste0(fDir, "/machine1_")
get_bin_info_states(prefix1)
convert_binary_to_hmp(prefix1, duration = NULL, out = "machine1_converted.hmp")
# have a look at the hmp file
cat(readr::read_file("machine1_converted.hmp"))

## Convert the machine example hmp file to binary files
convert_hmp_to_binary(file = paste0(fDir, "/machine1.hmp"), prefix = "machine_cov_")
get_bin_info_states(prefix = "machine_cov_")
## Convert the machine example with a single dummy node to a hmp file
#convert_binary_to_hmp("machine2_") # error since using scope = 3 not supported in hmp files

## Reset working dir
setwd(wd)
```

get_bin_info_actions *Info about the actions in the HMDP model under consideration.*

Description

Info about the actions in the HMDP model under consideration.

Usage

```
get_bin_info_actions(
  prefix = "",
  labels = TRUE,
  file_a = "actionIdx.bin",
  file_pr = "transProb.bin",
  file_w = "actionWeight.bin",
  file_label_w = "actionWeightLbl.bin",
  file_label_a = "actionIdxLbl.bin"
)
```

Arguments

prefix	A character string with the prefix added to til binary files.
labels	Should labels be extracted.
file_a	The binary file containing the description of actions.
file_pr	The binary file containing the description of transition probabilities.
file_w	The binary file containing the description of weights.
file_label_w	The binary file containing the weight labels.
file_label_a	The binary file containing the action labels.

Value

A data frame with the information. Scope string contain the scope of the transitions and can be 4 values:

- 0: A transition to the next stage in the father process,
- 1: A transition to next stage in the current process,
- 2: A transition to a child process (stage zero in the child process),
- 3: A transition to the state with `s_id = idx` is considered.

The index string denote the index (id is scope = 3) of the state at the next stage.

Note

The model don't have to be loaded, i.e only read the binary files. The state id (`s_id`) will not be the same as in the loaded model!

Examples

```
## Use temp dir
wd <- setwd(tempdir())

# Create a small HMDP with two levels
w<-binary_mdp_writer()
w$set_weights(c("Duration","Net reward","Items"))
w$process()
  w$stage()
    w$state(label="M0")
      w$action(label="A0",weights=c(0,0,0),prob=c(2,0,1))
        w$process()
          w$stage()
            w$state(label="D")
              w$action(label="A0",weights=c(0,0,1),prob=c(1,0,0.5,1,1,0.5))
                w$end_action()
              w$end_state()
            w$end_stage()
          w$stage()
            w$state(label="C0")
              w$action(label="A0",weights=c(0,0,0),prob=c(1,0,1))
```

```

        w$end_action()
        w$action(label="A1",weights=c(1,2,1),prob=c(1,0,0.5,1,1,0.5))
        w$end_action()
    w$end_state()
    w$state(label="C1")
        w$action(label="A0",weights=c(0,0,0),prob=c(1,0,1))
        w$end_action()
        w$action(label="A1",weights=c(1,2,1),prob=c(1,0,0.5,1,1,0.5))
        w$end_action()
    w$end_state()
w$end_stage()
w$stage()
    w$state(label="C0")
        w$action(label="A0",weights=c(1,4,0),prob=c(0,0,1))
        w$end_action()
    w$end_state()
    w$state(label="C1")
        w$action(label="A0",weights=c(1,4,0),prob=c(0,0,1))
        w$end_action()
    w$end_state()
w$end_stage()
w$end_process()
w$end_action()
w$action(label="A1",weights=c(0,0,0),prob=c(2,0,1))
w$process()
    w$stage()
        w$state(label="D")
            w$action(label="A0",weights=c(0,0,1),prob=c(1,0,1))
            w$end_action()
        w$end_state()
    w$end_stage()
w$stage()
    w$state(label="C0")
        w$action(label="A0",weights=c(0,0,0),prob=c(1,0,1))
        w$end_action()
        w$action(label="A1",weights=c(1,2,1),prob=c(1,0,0.5,1,1,0.5))
        w$end_action()
    w$end_state()
w$end_stage()
w$stage()
    w$state(label="C0")
        w$action(label="A0",weights=c(1,4,0),prob=c(0,0,1))
        w$end_action()
    w$end_state()
    w$state(label="C1")
        w$action(label="A0",weights=c(1,4,0),prob=c(0,0,1))
        w$end_action()
        w$action(label="A1",weights=c(0,10,5),prob=c(0,0,0.5,0,1,0.5))
        w$end_action()
    w$end_state()
w$end_stage()
w$end_process()
w$end_action()

```

```

w$end_state()
w$state(label="M1")
  w$action(label="A0",weights=c(0,0,0),prob=c(2,0,1))
  w$process()
    w$stage()
      w$state(label="D")
        w$action(label="A0",weights=c(0,0,1),prob=c(1,0,0.5,1,1,0.5))
        w$end_action()
      w$end_state()
    w$end_stage()
  w$stage()
    w$state(label="C0")
      w$action(label="A0",weights=c(0,0,0),prob=c(1,0,1))
      w$end_action()
    w$end_state()
  w$state(label="C1")
    w$action(label="A0",weights=c(0,0,0),prob=c(1,0,1))
    w$end_action()
  w$end_state()
w$end_stage()
w$stage()
  w$state(label="C0")
    w$action(label="A0",weights=c(1,4,0),prob=c(0,0,1))
    w$end_action()
  w$end_state()
  w$state(label="C1")
    w$action(label="A0",weights=c(1,4,0),prob=c(0,0,1))
    w$end_action()
  w$end_state()
w$end_stage()
w$end_process()
w$end_action()
w$end_state()
w$end_stage()
w$end_process()
w$close_writer()

## Info about the binary files (don't have to load the model first)
get_bin_info_states()
get_bin_info_actions()

## reset working dir
setwd(wd)

```

get_bin_info_states	<i>Info about the states in the binary files of the HMDP model under consideration.</i>
---------------------	---

Description

Info about the states in the binary files of the HMDP model under consideration.

Usage

```
get_bin_info_states(
  prefix = "",
  labels = TRUE,
  state_str = TRUE,
  file_s = "stateIdx.bin",
  label_s = "stateIdxLbl.bin"
)
```

Arguments

prefix	A character string with the prefix added to til binary files.
labels	Should labels be extracted.
state_str	Should state strings be extracted. If false then add columns (n0, s0, a0, ...) where n0 the index of the stage at level 0, s0 the index of the state and a0 the index of the action. If the HMDP has more than one level columns index (d1, s1, a1, ...) are added.
file_s	The binary file containing the description of states.
label_s	The binary file containing the state labels.

Value

A data frame with the information.

Note

The model don't have to be loaded, i.e only read the binary files. The state id (s_id) will not be the same as in the loaded model!

get_hypergraph	<i>Return the (parts of) state-expanded hypergraph</i>
----------------	--

Description

The function is useful together with [plot_hypergraph\(\)](#).

Usage

```
get_hypergraph(mdp, ...)
```

Arguments

mdp	The MDP loaded using load_mdp() .
...	Arguments passed to get_info() .

Value

A list representing the hypergraph with two elements: a tibble nodes and a tibble hyperarcs. hyperarcs stores action_weights, trans, and pr as list-columns of vectors. trans_weights is a list-column of matrices with one row per transition and one column per transition-weight name-space.

See Also

[plot_hypergraph\(\)](#) and [plot.HMDP\(\)](#).

Examples

```
## Set working dir
wd <- setwd(system.file("models", package = "MDP2"))

#### A finite-horizon replacement problem ####
mdp<-load_mdp("machine1_")
plot(mdp)
plot(mdp, action_color = "label") # colors based on labels
plot(mdp, trans_labels = "state") # label transitions with target state labels
plot(mdp, trans_labels = "prob") # label transitions with transition probabilities
plot(mdp, action_color = "label", state_label = "s_id|label") # state labels are 's_id | label'
plot(mdp, state_label = "s_idx|label", radx = 0.01) # adjust radx in states
plot(mdp, state_label = "label", action_w_label = "none", action_label = "label",
      trans_labels = "s_id", radx = 0.01)

scrapValues <- c(30, 10, 5, 0) # scrap values (the values of the 4 states at stage 4)
run_value_ite(mdp, "Net reward", term_values = scrapValues)
plot(mdp, action_color = "policy") # highlight optimal policy
plot(mdp, actions_visible = "policy", state_label = "weight") # show only optimal policy

#### An infinite-horizon maintenance problem ####
mdp<-load_mdp("hct611-1_")
plot(mdp) # plot the first two stages
plot(mdp, action_color = "label") # colors based on labels
plot(mdp, action_color = "label", state_label = "s_id|label") # state labels are 's_id | label'
run_policy_ite_ave(mdp, "Net reward", "Duration")
plot(mdp, action_color = "policy") # highlight optimal policy
plot(mdp, actions_visible = "policy") # show only optimal policy

#### An infinite-horizon hierarchical replacement problem ####
library(magrittr)
mdp<-load_mdp("cow_")
hgf <- get_hypergraph(mdp)
# modify labels
dat <- hgf$nodes %>%
  dplyr::mutate(label = dplyr::case_when(
    label == "Low yield" ~ "L",
    label == "Avg yield" ~ "A",
    label == "High yield" ~ "H",
```

```

        label == "Dummy" ~ "D",
        label == "Bad genetic level" ~ "Bad",
        label == "Avg genetic level" ~ "Avg",
        label == "Good genetic level" ~ "Good",
        TRUE ~ "Error"
    ))
# assign nodes to grid ids
dat$g_id[1:3]<-85:87
dat$g_id[43:45]<-1:3
get_g_id<-function(process,stage,state) {
  if (process==0) start=18
  if (process==1) start=22
  if (process==2) start=26
  return(start + 14 * stage + state)
}
idx<-43
for (process in 0:2)
  for (stage in 0:4)
    for (state in 0:2) {
      if (stage==0 & state>0) break
      idx<-idx-1
      #cat(idx,process,stage,state,get_g_id(process,stage,state),"\n")
      dat$g_id[idx]<-get_g_id(process,stage,state)
    }
hgf$nodes <- dat
# modify labels
dat <- hgf$hyperarcs %>%
  dplyr::mutate(label = dplyr::case_when(
    label == "Replace" ~ "R",
    label == "Keep" ~ "K",
    label == "Dummy" ~ "D",
    TRUE ~ "Error"
  ),
  col = dplyr::case_when(
    label == "R" ~ "deepskyblue3",
    label == "K" ~ "darkorange1",
    label == "D" ~ "black",
    TRUE ~ "Error"
  ),
  lwd = 0.5,
  label = ""
)
hgf$hyperarcs <- dat
# plot hypergraph
oldpar <- par(mai = c(0, 0, 0, 0))
plot_hypergraph(grid_dim = c(14, 7), hgf, cex = 0.8, radx = 0.02, rady = 0.03)
par(oldpar)

## A simple finite-horizon MDP with action and transition weights
prefix <- file.path(tempdir(), "plot_transition_rewards_")
w <- binary_mdp_writer(prefix)
w$set_weights("Cost")

```

```

w$set_trans_weights(c("Reward", "Disease"))
w$process()
  w$stage()
    w$state(label = "S1")
      w$action(
        label = "A1", weights = 2, id = c(1), pr = c(1),
        trans_weights = c(20, 0.3), end = TRUE
      )
      w$action(
        label = "A2", weights = 1, id = c(0, 1), pr = c(0.3, 0.7),
        trans_weights = c(25, 0.4, 15, 0.2), end = TRUE
      )
    w$end_state()
  w$end_stage()
w$stage()
  w$state(label = "S2")
    w$action(
      label = "A3", weights = 3, id = c(0, 1, 2), pr = c(0.5, 0.3, 0.2),
      trans_weights = c(0, 0.05, 12, 0.2, 30, 0.8), end = TRUE
    )
    w$action(
      label = "A4", weights = 2, id = c(1, 2), pr = c(0.6, 0.4),
      trans_weights = c(22, 0.35, 27, 0.7), end = TRUE
    )
  w$end_state()
w$state(label = "S3")
  w$action(
    label = "A5", weights = 1, id = c(0, 1), pr = c(0.4, 0.6),
    trans_weights = c(5, 0, 16, 0.25), end = TRUE
  )
  w$action(
    label = "A6", weights = 4, id = c(0, 1, 2), pr = c(0.1, 0.3, 0.6),
    trans_weights = c(14, 0.15, 21, 0.45, 29, 1), end = TRUE
  )
w$end_state()
w$end_stage()
w$stage()
  w$state(label = "S4", end = TRUE)
  w$state(label = "S5", end = TRUE)
  w$state(label = "S6", end = TRUE)
w$end_stage()
w$end_process()
w$close_writer()

mdp <- load_mdp(prefix, get_log = FALSE)
plot(mdp, action_color = "label", trans_labels = "weights", action_w_label = "weight",
      radx = 0.005, rady = 0.01)

## Reset working dir
setwd(wd)

```

get_info

*Information about the MDP***Description**

Information about the MDP

Usage

```
get_info(
  mdp,
  s_id = 1:ifelse(mdp$time_horizon < Inf, mdp$states, mdp$states +
    mdp$founder_states_last) - 1,
  state_str = NULL,
  stage_str = NULL,
  with_list = TRUE,
  with_df = TRUE,
  df_level = "state",
  as_strings_state = TRUE,
  as_strings_actions = FALSE,
  with_harc = FALSE
)
```

Arguments

mdp	The MDP loaded using load_mdp() .
s_id	The id of the state(s) considered.
state_str	A character vector containing the index of the state(s) (e.g. "n0,s0,a0,n1,s1"). Parameter s_id are ignored if not NULL.
stage_str	A character vector containing the index of the stage(s) (e.g. "n0,s0,a0,n1"). Parameter s_id and idx_s are ignored if not NULL.
with_list	Output info as a list lst.
with_df	Output the info as a data frame.
df_level	If with_df and equal "state" the data frame contains a row for each state. If equal "action" the data frame contains a row for each action.
as_strings_state	Write state vector as a string; otherwise, output it as columns.
as_strings_actions	Write action vectors (weights, transitions and probabilities) as strings; otherwise, output it as columns.
with_harc	Output a hyperarcs data frame. Each row contains a hyperarc with the first column denoting the head (s_id), the tails (s_id) and the label.

Value

A list containing the list, data frame(s).

Examples

```
## Set working dir
wd <- setwd(tempdir())

# Create the small machine replacement problem used as an example in L.R. Nielsen and A.R.
# Kristensen. Finding the K best policies in a finite-horizon Markov decision process. European
# Journal of Operational Research, 175(2):1164-1179, 2006. doi:10.1016/j.ejor.2005.06.011.

## Create the MDP using a dummy replacement node
prefix<-"machine1_"
w <- binary_mdp_writer(prefix)
w$set_weights(c("Net reward"))
w$process()

  w$stage() # stage n=0
    w$state(label="Dummy") # v=(0,0)
      w$action(label="buy", weights=-100, prob=c(1,0,0.7, 1,1,0.3), end=TRUE)
    w$end_state()
  w$end_stage()
  w$stage() # stage n=1
    w$state(label="good") # v=(1,0)
      w$action(label="mt", weights=55, prob=c(1,0,1), end=TRUE)
      w$action(label="nmt", weights=70, prob=c(1,0,0.6, 1,1,0.4), end=TRUE)
    w$end_state()
    w$state(label="average") # v=(1,1)
      w$action(label="mt", weights=40, prob=c(1,0,1), end=TRUE)
      w$action(label="nmt", weights=50, prob=c(1,1,0.6, 1,2,0.4), end=TRUE)
    w$end_state()
  w$end_stage()
  w$stage() # stage n=2
    w$state(label="good") # v=(2,0)
      w$action(label="mt", weights=55, prob=c(1,0,1), end=TRUE)
      w$action(label="nmt", weights=70, prob=c(1,0,0.5, 1,1,0.5), end=TRUE)
    w$end_state()
    w$state(label="average") # v=(2,1)
      w$action(label="mt", weights=40, prob=c(1,0,1), end=TRUE)
      w$action(label="nmt", weights=50, prob=c(1,1,0.5, 1,2,0.5), end=TRUE)
    w$end_state()
    w$state(label="not working") # v=(2,2)
      w$action(label="mt", weights=30, prob=c(1,0,1), end=TRUE)
      w$action(label="rep", weights=5, prob=c(1,3,1), end=TRUE)
    w$end_state()
  w$end_stage()
  w$stage() # stage n=3
    w$state(label="good") # v=(3,0)
      w$action(label="mt", weights=55, prob=c(1,0,1), end=TRUE)
      w$action(label="nmt", weights=70, prob=c(1,0,0.2, 1,1,0.8), end=TRUE)
    w$end_state()
    w$state(label="average") # v=(3,1)
      w$action(label="mt", weights=40, prob=c(1,0,1), end=TRUE)
      w$action(label="nmt", weights=50, prob=c(1,1,0.2, 1,2,0.8), end=TRUE)
    w$end_state()
    w$state(label="not working") # v=(3,2)
```

```

        w$action(label="mt", weights=30, prob=c(1,0,1), end=TRUE)
        w$action(label="rep", weights=5, prob=c(1,3,1), end=TRUE)
    w$end_state()
    w$state(label="replaced")      # v=(3,3)
        w$action(label="Dummy", weights=0, prob=c(1,3,1), end=TRUE)
    w$end_state()
w$end_stage()
w$stage() # stage n=4
    w$state(label="good", end=TRUE)      # v=(4,0)
    w$state(label="average", end=TRUE)   # v=(4,1)
    w$state(label="not working", end=TRUE) # v=(4,2)
    w$state(label="replaced", end=TRUE)  # v=(4,3)
w$end_stage()
w$end_process()
w$close_writer()

## Load the model into memory
mdp<-load_mdp(prefix)
mdp
plot(mdp)

get_info(mdp, with_list = FALSE)
get_info(mdp, with_list = FALSE, df_level = "action", as_strings_actions = TRUE)
get_info(mdp, with_list = FALSE, df_level = "action", as_strings_actions = FALSE)

## Perform value iteration
w<-"Net reward"          # label of the weight we want to optimize
scrapValues<-c(30,10,5,0) # scrap values (the values of the 4 states at stage 4)
run_value_ite(mdp, w, term_values=scrapValues)
get_policy(mdp)          # optimal policy

## Calculate the weights of the policy always to maintain
library(magrittr)
policy <- get_info(mdp, with_list = FALSE, df_level = "action")$df %>%
  dplyr::filter(label_action == "mt") %>%
  dplyr::select(s_id, a_idx)
set_policy(mdp, policy)
run_calc_weights(mdp, w, term_values=scrapValues)
get_policy(mdp)

# The example given in L.R. Nielsen and A.R. Kristensen. Finding the K best
# policies in a finite-horizon Markov decision process. European Journal of
# Operational Research, 175(2):1164-1179, 2006. doi:10.1016/j.ejor.2005.06.011,
# does actually not have any dummy replacement node as in the MDP above. The same
# model can be created using a single dummy node at the end of the process.

## Create the MDP using a single dummy node
prefix<-"machine2_"
w <- binary_mdp_writer(prefix)
w$set_weights(c("Net reward"))
w$process()

```

```

w$stage() # stage n=0
  w$state(label="Dummy") # v=(0,0)
    w$action(label="buy", weights=-100, prob=c(1,0,0.7, 1,1,0.3), end=TRUE)
  w$end_state()
w$end_stage()
w$stage() # stage n=1
  w$state(label="good") # v=(1,0)
    w$action(label="mt", weights=55, prob=c(1,0,1), end=TRUE)
    w$action(label="nmt", weights=70, prob=c(1,0,0.6, 1,1,0.4), end=TRUE)
  w$end_state()
  w$state(label="average") # v=(1,1)
    w$action(label="mt", weights=40, prob=c(1,0,1), end=TRUE)
    w$action(label="nmt", weights=50, prob=c(1,1,0.6, 1,2,0.4), end=TRUE)
  w$end_state()
w$end_stage()
w$stage() # stage n=2
  w$state(label="good") # v=(2,0)
    w$action(label="mt", weights=55, prob=c(1,0,1), end=TRUE)
    w$action(label="nmt", weights=70, prob=c(1,0,0.5, 1,1,0.5), end=TRUE)
  w$end_state()
  w$state(label="average") # v=(2,1)
    w$action(label="mt", weights=40, prob=c(1,0,1), end=TRUE)
    w$action(label="nmt", weights=50, prob=c(1,1,0.5, 1,2,0.5), end=TRUE)
  w$end_state()
  w$state(label="not working") # v=(2,2)
    w$action(label="mt", weights=30, prob=c(1,0,1), end=TRUE)
  w$action(label="rep", weights=5, prob=c(3,12,1), end=TRUE) # transition to s_id=12 (Dummy)
  w$end_state()
w$end_stage()
w$stage() # stage n=3
  w$state(label="good") # v=(3,0)
    w$action(label="mt", weights=55, prob=c(1,0,1), end=TRUE)
    w$action(label="nmt", weights=70, prob=c(1,0,0.2, 1,1,0.8), end=TRUE)
  w$end_state()
  w$state(label="average") # v=(3,1)
    w$action(label="mt", weights=40, prob=c(1,0,1), end=TRUE)
    w$action(label="nmt", weights=50, prob=c(1,1,0.2, 1,2,0.8), end=TRUE)
  w$end_state()
  w$state(label="not working") # v=(3,2)
    w$action(label="mt", weights=30, prob=c(1,0,1), end=TRUE)
    w$action(label="rep", weights=5, prob=c(3,12,1), end=TRUE)
  w$end_state()
w$end_stage()
w$stage() # stage n=4
  w$state(label="good") # v=(4,0)
    w$action(label="rep", weights=30, prob=c(1,0,1), end=TRUE)
  w$end_state()
  w$state(label="average") # v=(4,1)
    w$action(label="rep", weights=10, prob=c(1,0,1), end=TRUE)
  w$end_state()
  w$state(label="not working") # v=(4,2)
    w$action(label="rep", weights=5, prob=c(1,0,1), end=TRUE)
  w$end_state()

```

```

    w$end_stage()
    w$stage() # stage n=5
    w$state(label="Dummy", end=TRUE) # v=(5,0)
    w$end_stage()
  w$end_process()
  w$close_writer()

## Have a look at the state-expanded hypergraph
mdp<-load_mdp(prefix)
mdp
plot(mdp)

get_info(mdp, with_list = FALSE)
get_info(mdp, with_list = FALSE, df_level = "action", as_strings_actions = TRUE)
get_info(mdp, with_list = FALSE, df_level = "action", as_strings_actions = FALSE)

## Perform value iteration
w<-"Net reward" # label of the weight we want to optimize
run_value_ite(mdp, w, term_values = 0)
get_policy(mdp) # optimal policy

## Calculate the weights of the policy always to maintain
library(magrittr)
policy <- get_info(mdp, with_list = FALSE, df_level = "action")$df %>%
  dplyr::filter(label_action == "mt") %>%
  dplyr::select(s_id, a_idx)
set_policy(mdp, policy)
run_calc_weights(mdp, w, term_values=scrapValues)
get_policy(mdp)

## Reset working dir
setwd(wd)

```

get_policy

*Get parts of the optimal policy.***Description**

Get parts of the optimal policy.

Usage

```

get_policy(
  mdp,
  s_id = ifelse(mdp$time_horizon >= Inf, mdp$founder_states_last + 1,
    1):ifelse(mdp$time_horizon >= Inf, mdp$states + mdp$founder_states_last, mdp$states)
    - 1,
  stage_str = NULL,
  state_labels = TRUE,

```

```

    action_labels = TRUE,
    action_idx = TRUE,
    rewards = TRUE,
    state_str = TRUE,
    external = NULL,
    ...
)

```

Arguments

mdp	The MDP loaded using load_mdp() .
s_id	Vector of id's of the states we want to retrieve.
stage_str	Stage string. If specified then find s_id based on the stage string.
state_labels	Add state labels.
action_labels	Add action labels of policy.
action_idx	Add action index.
rewards	Add weights calculated for each state.
state_str	Add the state string for each state.
external	A vector of stage strings corresponding to external processes we want the optimal policy of.
...	Parameters passed on when find the optimal policy of the external processes. Note if external is specified then it must contain stage strings from mdp\$external. Moreover you must specify further arguments passed on to run_value_ite() used for recreating the optimal policy e.g. the g value and the label for weight and duration. See the vignette about external processes.

Value

The policy (data frame).

Examples

```

## Set working dir
wd <- setwd(tempdir())

# Create the small machine replacement problem used as an example in L.R. Nielsen and A.R.
# Kristensen. Finding the K best policies in a finite-horizon Markov decision process. European
# Journal of Operational Research, 175(2):1164-1179, 2006. doi:10.1016/j.ejor.2005.06.011.

## Create the MDP using a dummy replacement node
prefix<-"machine1_"
w <- binary_mdp_writer(prefix)
w$set_weights(c("Net reward"))
w$process()
  w$stage()    # stage n=0
    w$state(label="Dummy")      # v=(0,0)
      w$action(label="buy", weights=-100, prob=c(1,0,0.7, 1,1,0.3), end=TRUE)

```

```

        w$end_state()
w$end_stage()
w$stage() # stage n=1
    w$state(label="good") # v=(1,0)
        w$action(label="mt", weights=55, prob=c(1,0,1), end=TRUE)
        w$action(label="nmt", weights=70, prob=c(1,0,0.6, 1,1,0.4), end=TRUE)
    w$end_state()
    w$state(label="average") # v=(1,1)
        w$action(label="mt", weights=40, prob=c(1,0,1), end=TRUE)
        w$action(label="nmt", weights=50, prob=c(1,1,0.6, 1,2,0.4), end=TRUE)
    w$end_state()
w$end_stage()
w$stage() # stage n=2
    w$state(label="good") # v=(2,0)
        w$action(label="mt", weights=55, prob=c(1,0,1), end=TRUE)
        w$action(label="nmt", weights=70, prob=c(1,0,0.5, 1,1,0.5), end=TRUE)
    w$end_state()
    w$state(label="average") # v=(2,1)
        w$action(label="mt", weights=40, prob=c(1,0,1), end=TRUE)
        w$action(label="nmt", weights=50, prob=c(1,1,0.5, 1,2,0.5), end=TRUE)
    w$end_state()
    w$state(label="not working") # v=(2,2)
        w$action(label="mt", weights=30, prob=c(1,0,1), end=TRUE)
        w$action(label="rep", weights=5, prob=c(1,3,1), end=TRUE)
    w$end_state()
w$end_stage()
w$stage() # stage n=3
    w$state(label="good") # v=(3,0)
        w$action(label="mt", weights=55, prob=c(1,0,1), end=TRUE)
        w$action(label="nmt", weights=70, prob=c(1,0,0.2, 1,1,0.8), end=TRUE)
    w$end_state()
    w$state(label="average") # v=(3,1)
        w$action(label="mt", weights=40, prob=c(1,0,1), end=TRUE)
        w$action(label="nmt", weights=50, prob=c(1,1,0.2, 1,2,0.8), end=TRUE)
    w$end_state()
    w$state(label="not working") # v=(3,2)
        w$action(label="mt", weights=30, prob=c(1,0,1), end=TRUE)
        w$action(label="rep", weights=5, prob=c(1,3,1), end=TRUE)
    w$end_state()
    w$state(label="replaced") # v=(3,3)
        w$action(label="Dummy", weights=0, prob=c(1,3,1), end=TRUE)
    w$end_state()
w$end_stage()
w$stage() # stage n=4
    w$state(label="good", end=TRUE) # v=(4,0)
    w$state(label="average", end=TRUE) # v=(4,1)
    w$state(label="not working", end=TRUE) # v=(4,2)
    w$state(label="replaced", end=TRUE) # v=(4,3)
w$end_stage()
w$end_process()
w$close_writer()

## Load the model into memory

```

```

mdp<-load_mdp(prefix)
mdp
plot(mdp)

get_info(mdp, with_list = FALSE)
get_info(mdp, with_list = FALSE, df_level = "action", as_strings_actions = TRUE)
get_info(mdp, with_list = FALSE, df_level = "action", as_strings_actions = FALSE)

## Perform value iteration
w<-"Net reward"          # label of the weight we want to optimize
scrapValues<-c(30,10,5,0) # scrap values (the values of the 4 states at stage 4)
run_value_ite(mdp, w, term_values=scrapValues)
get_policy(mdp)          # optimal policy

## Calculate the weights of the policy always to maintain
library(magrittr)
policy <- get_info(mdp, with_list = FALSE, df_level = "action")$df %>%
  dplyr::filter(label_action == "mt") %>%
  dplyr::select(s_id, a_idx)
set_policy(mdp, policy)
run_calc_weights(mdp, w, term_values=scrapValues)
get_policy(mdp)

# The example given in L.R. Nielsen and A.R. Kristensen. Finding the K best
# policies in a finite-horizon Markov decision process. European Journal of
# Operational Research, 175(2):1164-1179, 2006. doi:10.1016/j.ejor.2005.06.011,
# does actually not have any dummy replacement node as in the MDP above. The same
# model can be created using a single dummy node at the end of the process.

## Create the MDP using a single dummy node
prefix<-"machine2_"
w <- binary_mdp_writer(prefix)
w$set_weights(c("Net reward"))
w$process()
  w$stage() # stage n=0
    w$state(label="Dummy") # v=(0,0)
      w$action(label="buy", weights=-100, prob=c(1,0,0.7, 1,1,0.3), end=TRUE)
    w$end_state()
  w$end_stage()
  w$stage() # stage n=1
    w$state(label="good") # v=(1,0)
      w$action(label="mt", weights=55, prob=c(1,0,1), end=TRUE)
      w$action(label="nmt", weights=70, prob=c(1,0,0.6, 1,1,0.4), end=TRUE)
    w$end_state()
    w$state(label="average") # v=(1,1)
      w$action(label="mt", weights=40, prob=c(1,0,1), end=TRUE)
      w$action(label="nmt", weights=50, prob=c(1,1,0.6, 1,2,0.4), end=TRUE)
    w$end_state()
  w$end_stage()
  w$stage() # stage n=2
    w$state(label="good") # v=(2,0)

```

```

        w$action(label="mt", weights=55, prob=c(1,0,1), end=TRUE)
        w$action(label="nmt", weights=70, prob=c(1,0,0.5, 1,1,0.5), end=TRUE)
    w$end_state()
    w$state(label="average")          # v=(2,1)
        w$action(label="mt", weights=40, prob=c(1,0,1), end=TRUE)
        w$action(label="nmt", weights=50, prob=c(1,1,0.5, 1,2,0.5), end=TRUE)
    w$end_state()
    w$state(label="not working")      # v=(2,2)
        w$action(label="mt", weights=30, prob=c(1,0,1), end=TRUE)
    w$action(label="rep", weights=5, prob=c(3,12,1), end=TRUE) # transition to s_id=12 (Dummy)
    w$end_state()
w$end_stage()
w$stage() # stage n=3
    w$state(label="good")            # v=(3,0)
        w$action(label="mt", weights=55, prob=c(1,0,1), end=TRUE)
        w$action(label="nmt", weights=70, prob=c(1,0,0.2, 1,1,0.8), end=TRUE)
    w$end_state()
    w$state(label="average")          # v=(3,1)
        w$action(label="mt", weights=40, prob=c(1,0,1), end=TRUE)
        w$action(label="nmt", weights=50, prob=c(1,1,0.2, 1,2,0.8), end=TRUE)
    w$end_state()
    w$state(label="not working")      # v=(3,2)
        w$action(label="mt", weights=30, prob=c(1,0,1), end=TRUE)
        w$action(label="rep", weights=5, prob=c(3,12,1), end=TRUE)
    w$end_state()
w$end_stage()
w$stage() # stage n=4
    w$state(label="good")            # v=(4,0)
        w$action(label="rep", weights=30, prob=c(1,0,1), end=TRUE)
    w$end_state()
    w$state(label="average")          # v=(4,1)
        w$action(label="rep", weights=10, prob=c(1,0,1), end=TRUE)
    w$end_state()
    w$state(label="not working")      # v=(4,2)
        w$action(label="rep", weights=5, prob=c(1,0,1), end=TRUE)
    w$end_state()
w$end_stage()
w$stage() # stage n=5
    w$state(label="Dummy", end=TRUE) # v=(5,0)
w$end_stage()
w$end_process()
w$close_writer()

## Have a look at the state-expanded hypergraph
mdp<-load_mdp(prefix)
mdp
plot(mdp)

get_info(mdp, with_list = FALSE)
get_info(mdp, with_list = FALSE, df_level = "action", as_strings_actions = TRUE)
get_info(mdp, with_list = FALSE, df_level = "action", as_strings_actions = FALSE)

## Perform value iteration

```

```

w<-"Net reward"          # label of the weight we want to optimize
run_value_ite(mdp, w, term_values = 0)
get_policy(mdp)          # optimal policy

## Calculate the weights of the policy always to maintain
library(magrittr)
policy <- get_info(mdp, with_list = FALSE, df_level = "action")$df %>%
  dplyr::filter(label_action == "mt") %>%
  dplyr::select(s_id, a_idx)
set_policy(mdp, policy)
run_calc_weights(mdp, w, term_values=scrapValues)
get_policy(mdp)

## Reset working dir
setwd(wd)

```

get_rpo	<i>Calculate the retention pay-off (RPO) or opportunity cost for some states.</i>
---------	---

Description

The RPO is defined as the difference between the weight of the state when using action `i_a` and the maximum weight of the node when using another predecessor different from `i_a`.

Usage

```

get_rpo(
  mdp,
  w,
  i_a,
  s_id = ifelse(mdp$time_horizon >= Inf, mdp$founder_states_last + 1,
    1):ifelse(mdp$time_horizon >= Inf, mdp$states + mdp$founder_states_last, mdp$states)
    - 1,
  criterion = "expected",
  dur = "",
  rate = 0,
  rate_base = 1,
  discount_factor = NULL,
  g = 0,
  objective = c("max", "min"),
  discount_method = "continuous",
  state_str = TRUE
)

```

Arguments

mdp	The MDP loaded using load_mdp() .
w	The label of the weight we calculate RPO for.
i_a	The action index we calculate the RPO with respect to (same size as s_id).
s_id	Vector of id's of the states we want to retrieve.
criterion	The Bellman operator shortcut. If expected use expected weights, if discount use discounted expected weights, if average use average expected weights.
dur	The label of the duration/time such that discount rates can be calculated.
rate	The interest rate.
rate_base	The time-horizon the rate is valid over.
discount_factor	The discount rate for one time unit. If specified rate and rate_base are not used to calculate the discount rate.
g	The optimal gain (g) calculated (used if criterion = "average").
objective	Optimize by maximizing ("max") or minimizing ("min") the Bellman value.
discount_method	Either 'continuous' or 'discrete', corresponding to discount factor $\exp(-rate/rate_base)$ or $1/(1 + rate/rate_base)$, respectively. Only used if discount_factor is NULL.
state_str	Output the state string.

Value

The RPO (matrix/data frame).

get_steady_state_pr	<i>Calculate the steady state transition probabilities for the founder process (level 0).</i>
---------------------	---

Description

Assume that we consider an ergodic/irreducible time-homogeneous Markov chain specified using a policy in the MDP.

Usage

```
get_steady_state_pr(mdp, get_log = FALSE)
```

Arguments

mdp	The MDP loaded using load_mdp() .
get_log	Output log text.

Value

A vector with steady state probabilities for all the states at the founder level.

<code>get_w_idx</code>	<i>Return the index of a weight in the model. Note that index always start from zero (C++ style), i.e. the first weight, the first state at a stage etc has index 0.</i>
------------------------	--

Description

Return the index of a weight in the model. Note that index always start from zero (C++ style), i.e. the first weight, the first state at a stage etc has index 0.

Usage

```
get_w_idx(mdp, w_lbl)
```

Arguments

<code>mdp</code>	The MDP loaded using load_mdp() .
<code>w_lbl</code>	The label/string of the weight.

Value

The index (integer).

<code>hmp_mdp_writer</code>	<i>Function for writing an HMDP model to a hmp file (XML). The function define sub-functions which can be used to define an HMDP model stored in a hmp file.</i>
-----------------------------	--

Description

HMP files are in XML format and human readable using e.g. a text editor. HMP files are not suitable for storing large HMDP models since text files are very verbose. Moreover, approximation of the weights and probabilities may occur since the parser writing the hmp file may no output all digits. If you consider large models then use the binary file format instead.

Usage

```
hmp_mdp_writer(
  file = "r.hmp",
  rate = 0.1,
  rate_base = 1,
  precision = 1e-05,
  desc = "HMP file created using hmp_mdp_writer in R",
  get_log = TRUE
)
```

Arguments

file	The name of the file storing the model (e.g. r.hmp).
rate	The interest rate (used if consider discounting).
rate_base	The time where the rate is taken over, e.g. if the rate is 0.1 and rate_base is 365 days then we have an interest rate of 10 percent over the year.
precision	The precision used when checking if probabilities sum to one.
desc	Description of the model.
get_log	Output log text.

Details

The returned writer exposes these functions:

- `set_weights(labels, duration)`: sets the labels of the weights used in the actions. `labels` is a vector of label names. `duration` identifies which label corresponds to duration or time. For example, if the first entry in `labels` is time, then `duration = 1`. Call this before building the model.
- `set_trans_weights(labels)`: sets the labels of transition-level weights.
- `process()`: starts a (sub)process.
- `end_process()`: ends a (sub)process.
- `stage(label = NULL)`: starts a stage.
- `end_stage()`: ends a stage.
- `state(label = NULL)`: starts a state and returns the state index `s_idx`.
- `end_state()`: ends a state.
- `action(label = NULL, weights, prob, states_next = NULL, trans_weights = NULL)`: starts an action. `weights` must be a vector of action weights, and `prob` must contain triples (`scope`, `idx`, `pr`). `scope` can take three values:
 - 0: a transition to the next stage in the father process.
 - 1: a transition to the next stage in the current process.
 - 2: a transition to a child process, at stage zero in the child process.

The `idx` value denotes the index of the state at the stage considered. For example, if `scope = 1` and `idx = 2`, the transition is to state number 3 at the next stage in the current process, counting from zero. `scope = 3` is not supported in the hmp file format. `states_next` is the number of states in the next stage of the process and is only needed when there is a transition to the father.

- `end_action()`: ends an action.
- `close_writer()`: closes the writer. Call this when the model description is finished.

Value

A list of functions.

Note

Note all indexes are starting from zero (C/C++ style).

Examples

```
## Use temp dir
wd <- setwd(tempdir())

## Create a small HMDP with two levels
w<-hmp_mdp_writer()
w$set_weights(c("Duration","Net reward","Items"), duration=1)
w$process()
w$stage()
  w$state(label="M0")
  w$action(label="A0",weights=c(0,0,0),prob=c(2,0,1))
  w$process()
  w$stage()
  w$state(label="D")
    w$action(label="A0",weights=c(0,0,1),prob=c(1,0,0.5,1,1,0.5))
    w$end_action()
  w$end_state()
w$end_stage()
w$stage()
  w$state(label="C0")
    w$action(label="A0",weights=c(0,0,0),prob=c(1,0,1))
    w$end_action()
    w$action(label="A1",weights=c(1,2,1),prob=c(1,0,0.5,1,1,0.5))
    w$end_action()
  w$end_state()
w$state(label="C1")
  w$action(label="A0",weights=c(0,0,0),prob=c(1,0,1))
  w$end_action()
  w$action(label="A1",weights=c(1,2,1),prob=c(1,0,0.5,1,1,0.5))
  w$end_action()
w$end_state()
w$end_stage()
w$stage()
  w$state(label="C0")
    w$action(label="A0",weights=c(1,4,0),prob=c(0,0,1), states_next=0)
    w$end_action()
  w$end_state()
w$state(label="C1")
  w$action(label="A0",weights=c(1,4,0),prob=c(0,0,1), states_next=0)
  w$end_action()
w$end_state()
w$end_stage()
w$end_process()
w$end_action()
w$action(label="A1",weights=c(0,0,0),prob=c(2,0,1))
w$process()
  w$stage()
  w$state(label="D")
```

```

        w$action(label="A0",weights=c(0,0,1),prob=c(1,0,1))
        w$end_action()
    w$end_state()
w$end_stage()
w$stage()
w$state(label="C0")
    w$action(label="A0",weights=c(0,0,0),prob=c(1,0,1))
    w$end_action()
    w$action(label="A1",weights=c(1,2,1),prob=c(1,0,0.5,1,1,0.5))
    w$end_action()
w$end_state()
w$end_stage()
w$stage()
w$state(label="C0")
    w$action(label="A0",weights=c(1,4,0),prob=c(0,0,1), states_next=0)
    w$end_action()
w$end_state()
w$state(label="C1")
    w$action(label="A0",weights=c(1,4,0),prob=c(0,0,1), states_next=0)
    w$end_action()
    w$action(label="A1",weights=c(0,10,5),prob=c(0,0,0.5,0,1,0.5), states_next=0)
    w$end_action()
w$end_state()
w$end_stage()
w$end_process()
w$end_action()
w$end_state()
w$state(label="M1")
w$action(label="A0",weights=c(0,0,0),prob=c(2,0,1))
w$process()
    w$stage()
    w$state(label="D")
        w$action(label="A0",weights=c(0,0,1),prob=c(1,0,0.5,1,1,0.5))
        w$end_action()
    w$end_state()
w$end_stage()
w$stage()
w$state(label="C0")
    w$action(label="A0",weights=c(0,0,0),prob=c(1,0,1))
    w$end_action()
w$end_state()
w$state(label="C1")
    w$action(label="A0",weights=c(0,0,0),prob=c(1,0,1))
    w$end_action()
w$end_state()
w$end_stage()
w$stage()
w$state(label="C0")
    w$action(label="A0",weights=c(1,4,0),prob=c(0,0,1), states_next=0)
    w$end_action()
w$end_state()
w$state(label="C1")
    w$action(label="A0",weights=c(1,4,0),prob=c(0,0,1), states_next=0)

```

```

        w$end_action()
        w$end_state()
        w$end_stage()
        w$end_process()
        w$end_action()
        w$end_state()
        w$end_stage()
    w$end_process()
    w$close_writer()

    ## Have a look at the hmp file
    cat(readr::read_file("r.hmp"))

    ## Reset working dir
    setwd(wd)

```

load_mdp	<i>Load the HMDP model defined in the binary files. The model are created in memory using the external C++ library.</i>
----------	---

Description

Load the HMDP model defined in the binary files. The model are created in memory using the external C++ library.

Usage

```

load_mdp(
  prefix = "",
  bin_names = c("stateIdx.bin", "stateIdxLbl.bin", "actionIdx.bin", "actionIdxLbl.bin",
    "actionWeight.bin", "actionWeightLbl.bin", "transProb.bin", "externalProcesses.bin",
    "transWeight.bin", "transWeightLbl.bin"),
  eps = 1e-05,
  check = TRUE,
  verbose = FALSE,
  get_log = TRUE
)

```

Arguments

prefix	A character string with the prefix added to bin_names. Used to identify a specific model.
bin_names	A character vector of length 7 giving the names of the binary files storing the model.
eps	The sum of the transition probabilities must at most differ eps from one.
check	Check if the MDP seems correct.
verbose	More output when running algorithms.
get_log	Output the log messages.

Value

A list containing relevant information about the model such as model file names (bin_names), time horizon (time_horizon), number of states (states), number of states at last stage of the founder process (founder_states_last), number of actions (actions), number of levels (levels), names of the weights associated to each action (weight_names) and a pointer ptr to the model object in memory. Note for models with an infinite time-horizon the states at the founder level is repeated at stage two so have something aka a double array representation.

Examples

```
## Set working dir
wd <- setwd(tempdir())

# Create the small machine replacement problem used as an example in L.R. Nielsen and A.R.
# Kristensen. Finding the K best policies in a finite-horizon Markov decision process. European
# Journal of Operational Research, 175(2):1164-1179, 2006. doi:10.1016/j.ejor.2005.06.011.

## Create the MDP using a dummy replacement node
prefix<-"machine1_"
w <- binary_mdp_writer(prefix)
w$set_weights(c("Net reward"))
w$process()
  w$stage() # stage n=0
    w$state(label="Dummy") # v=(0,0)
      w$action(label="buy", weights=-100, prob=c(1,0,0.7, 1,1,0.3), end=TRUE)
    w$end_state()
  w$end_stage()
  w$stage() # stage n=1
    w$state(label="good") # v=(1,0)
      w$action(label="mt", weights=55, prob=c(1,0,1), end=TRUE)
      w$action(label="nmt", weights=70, prob=c(1,0,0.6, 1,1,0.4), end=TRUE)
    w$end_state()
    w$state(label="average") # v=(1,1)
      w$action(label="mt", weights=40, prob=c(1,0,1), end=TRUE)
      w$action(label="nmt", weights=50, prob=c(1,1,0.6, 1,2,0.4), end=TRUE)
    w$end_state()
  w$end_stage()
  w$stage() # stage n=2
    w$state(label="good") # v=(2,0)
      w$action(label="mt", weights=55, prob=c(1,0,1), end=TRUE)
      w$action(label="nmt", weights=70, prob=c(1,0,0.5, 1,1,0.5), end=TRUE)
    w$end_state()
    w$state(label="average") # v=(2,1)
      w$action(label="mt", weights=40, prob=c(1,0,1), end=TRUE)
      w$action(label="nmt", weights=50, prob=c(1,1,0.5, 1,2,0.5), end=TRUE)
    w$end_state()
    w$state(label="not working") # v=(2,2)
      w$action(label="mt", weights=30, prob=c(1,0,1), end=TRUE)
      w$action(label="rep", weights=5, prob=c(1,3,1), end=TRUE)
    w$end_state()
  w$end_stage()
  w$stage() # stage n=3
```

```

w$state(label="good")          # v=(3,0)
  w$action(label="mt", weights=55, prob=c(1,0,1), end=TRUE)
  w$action(label="nmt", weights=70, prob=c(1,0,0.2, 1,1,0.8), end=TRUE)
w$end_state()
w$state(label="average")       # v=(3,1)
  w$action(label="mt", weights=40, prob=c(1,0,1), end=TRUE)
  w$action(label="nmt", weights=50, prob=c(1,1,0.2, 1,2,0.8), end=TRUE)
w$end_state()
w$state(label="not working")   # v=(3,2)
  w$action(label="mt", weights=30, prob=c(1,0,1), end=TRUE)
  w$action(label="rep", weights=5, prob=c(1,3,1), end=TRUE)
w$end_state()
w$state(label="replaced")      # v=(3,3)
  w$action(label="Dummy", weights=0, prob=c(1,3,1), end=TRUE)
w$end_state()
w$end_stage()
w$stage() # stage n=4
  w$state(label="good", end=TRUE)      # v=(4,0)
  w$state(label="average", end=TRUE)   # v=(4,1)
  w$state(label="not working", end=TRUE) # v=(4,2)
  w$state(label="replaced", end=TRUE)  # v=(4,3)
w$end_stage()
w$end_process()
w$close_writer()

## Load the model into memory
mdp<-load_mdp(prefix)
mdp
plot(mdp)

get_info(mdp, with_list = FALSE)
get_info(mdp, with_list = FALSE, df_level = "action", as_strings_actions = TRUE)
get_info(mdp, with_list = FALSE, df_level = "action", as_strings_actions = FALSE)

## Perform value iteration
w<-"Net reward"                # label of the weight we want to optimize
scrapValues<-c(30,10,5,0)      # scrap values (the values of the 4 states at stage 4)
run_value_ite(mdp, w, term_values=scrapValues)
get_policy(mdp)                # optimal policy

## Calculate the weights of the policy always to maintain
library(magrittr)
policy <- get_info(mdp, with_list = FALSE, df_level = "action")$df %>%
  dplyr::filter(label_action == "mt") %>%
  dplyr::select(s_id, a_idx)
set_policy(mdp, policy)
run_calc_weights(mdp, w, term_values=scrapValues)
get_policy(mdp)

# The example given in L.R. Nielsen and A.R. Kristensen. Finding the K best
# policies in a finite-horizon Markov decision process. European Journal of

```

Operational Research, 175(2):1164-1179, 2006. doi:10.1016/j.ejor.2005.06.011,
 # does actually not have any dummy replacement node as in the MDP above. The same
 # model can be created using a single dummy node at the end of the process.

```
## Create the MDP using a single dummy node
prefix<-"machine2_"
w <- binary_mdp_writer(prefix)
w$set_weights(c("Net reward"))
w$process()
  w$stage() # stage n=0
    w$state(label="Dummy") # v=(0,0)
      w$action(label="buy", weights=-100, prob=c(1,0,0.7, 1,1,0.3), end=TRUE)
    w$end_state()
  w$end_stage()
  w$stage() # stage n=1
    w$state(label="good") # v=(1,0)
      w$action(label="mt", weights=55, prob=c(1,0,1), end=TRUE)
      w$action(label="nmt", weights=70, prob=c(1,0,0.6, 1,1,0.4), end=TRUE)
    w$end_state()
    w$state(label="average") # v=(1,1)
      w$action(label="mt", weights=40, prob=c(1,0,1), end=TRUE)
      w$action(label="nmt", weights=50, prob=c(1,1,0.6, 1,2,0.4), end=TRUE)
    w$end_state()
  w$end_stage()
  w$stage() # stage n=2
    w$state(label="good") # v=(2,0)
      w$action(label="mt", weights=55, prob=c(1,0,1), end=TRUE)
      w$action(label="nmt", weights=70, prob=c(1,0,0.5, 1,1,0.5), end=TRUE)
    w$end_state()
    w$state(label="average") # v=(2,1)
      w$action(label="mt", weights=40, prob=c(1,0,1), end=TRUE)
      w$action(label="nmt", weights=50, prob=c(1,1,0.5, 1,2,0.5), end=TRUE)
    w$end_state()
    w$state(label="not working") # v=(2,2)
      w$action(label="mt", weights=30, prob=c(1,0,1), end=TRUE)
      w$action(label="rep", weights=5, prob=c(3,12,1), end=TRUE) # transition to s_id=12 (Dummy)
    w$end_state()
  w$end_stage()
  w$stage() # stage n=3
    w$state(label="good") # v=(3,0)
      w$action(label="mt", weights=55, prob=c(1,0,1), end=TRUE)
      w$action(label="nmt", weights=70, prob=c(1,0,0.2, 1,1,0.8), end=TRUE)
    w$end_state()
    w$state(label="average") # v=(3,1)
      w$action(label="mt", weights=40, prob=c(1,0,1), end=TRUE)
      w$action(label="nmt", weights=50, prob=c(1,1,0.2, 1,2,0.8), end=TRUE)
    w$end_state()
    w$state(label="not working") # v=(3,2)
      w$action(label="mt", weights=30, prob=c(1,0,1), end=TRUE)
      w$action(label="rep", weights=5, prob=c(3,12,1), end=TRUE)
    w$end_state()
  w$end_stage()
  w$stage() # stage n=4
```

```

w$state(label="good")          # v=(4,0)
  w$action(label="rep", weights=30, prob=c(1,0,1), end=TRUE)
w$end_state()
w$state(label="average")       # v=(4,1)
  w$action(label="rep", weights=10, prob=c(1,0,1), end=TRUE)
w$end_state()
w$state(label="not working")   # v=(4,2)
  w$action(label="rep", weights=5, prob=c(1,0,1), end=TRUE)
w$end_state()
w$end_stage()
w$stage() # stage n=5
  w$state(label="Dummy", end=TRUE)          # v=(5,0)
w$end_stage()
w$end_process()
w$close_writer()

## Have a look at the state-expanded hypergraph
mdp<-load_mdp(prefix)
mdp
plot(mdp)

get_info(mdp, with_list = FALSE)
get_info(mdp, with_list = FALSE, df_level = "action", as_strings_actions = TRUE)
get_info(mdp, with_list = FALSE, df_level = "action", as_strings_actions = FALSE)

## Perform value iteration
w<-"Net reward"          # label of the weight we want to optimize
run_value_ite(mdp, w, term_values = 0)
get_policy(mdp)          # optimal policy

## Calculate the weights of the policy always to maintain
library(magrittr)
policy <- get_info(mdp, with_list = FALSE, df_level = "action")$df %>%
  dplyr::filter(label_action == "mt") %>%
  dplyr::select(s_id, a_idx)
set_policy(mdp, policy)
run_calc_weights(mdp, w, term_values=scrapValues)
get_policy(mdp)

## Reset working dir
setwd(wd)

```

memory_mdp_writer

Function for building an HMDP model directly in memory.

Description

memory_mdp_writer() defines the same main sub-functions as [binary_mdp_writer\(\)](#), but stores states and actions directly in C++ memory instead of writing intermediate binary files. close_writer() compiles the model and returns the loaded "HMDP" object.

Usage

```
memory_mdp_writer(
  prefix = "",
  eps = 1e-05,
  check = TRUE,
  verbose = FALSE,
  get_log = TRUE
)
```

Arguments

prefix	A character string kept for compatibility and stored in the returned object meta-data.
eps	The sum of transition probabilities must at most differ eps from one when check = TRUE.
check	Check if the MDP seems correct before returning it.
verbose	More output when compiling and running algorithms.
get_log	Output the log messages.

Details

External or included processes are not supported by `memory_mdp_writer()`.

Value

A list of functions. Calling `close_writer()` returns an "HMDP" object.

Note

Note all indexes are starting from zero (C/C++ style).

Examples

```
## Use temp dir
wd <- setwd(tempdir())

# Create a small HMDP with two levels
w<-memory_mdp_writer()
w$set_weights(c("Duration","Net reward","Items"))
w$process()
  w$stage()
    w$state(label="M0")
      w$action(label="A0",weights=c(0,0,0),prob=c(2,0,1))
        w$process()
          w$stage()
            w$state(label="D")
              w$action(label="A0",weights=c(0,0,1),prob=c(1,0,0.5,1,1,0.5))
                w$end_action()
              w$end_state()
```

```

w$end_stage()
w$stage()
  w$state(label="C0")
    w$action(label="A0",weights=c(0,0,0),prob=c(1,0,1))
    w$end_action()
    w$action(label="A1",weights=c(1,2,1),prob=c(1,0,0.5,1,1,0.5))
    w$end_action()
  w$end_state()
  w$state(label="C1")
    w$action(label="A0",weights=c(0,0,0),prob=c(1,0,1))
    w$end_action()
    w$action(label="A1",weights=c(1,2,1),prob=c(1,0,0.5,1,1,0.5))
    w$end_action()
  w$end_state()
w$end_stage()
w$stage()
  w$state(label="C0")
    w$action(label="A0",weights=c(1,4,0),prob=c(0,0,1))
    w$end_action()
  w$end_state()
  w$state(label="C1")
    w$action(label="A0",weights=c(1,4,0),prob=c(0,0,1))
    w$end_action()
  w$end_state()
w$end_stage()
w$end_process()
w$end_action()
w$action(label="A1",weights=c(0,0,0),prob=c(2,0,1))
w$process()
  w$stage()
    w$state(label="D")
      w$action(label="A0",weights=c(0,0,1),prob=c(1,0,1))
      w$end_action()
    w$end_state()
  w$end_stage()
w$stage()
  w$state(label="C0")
    w$action(label="A0",weights=c(0,0,0),prob=c(1,0,1))
    w$end_action()
    w$action(label="A1",weights=c(1,2,1),prob=c(1,0,0.5,1,1,0.5))
    w$end_action()
  w$end_state()
w$end_stage()
w$stage()
  w$state(label="C0")
    w$action(label="A0",weights=c(1,4,0),prob=c(0,0,1))
    w$end_action()
  w$end_state()
  w$state(label="C1")
    w$action(label="A0",weights=c(1,4,0),prob=c(0,0,1))
    w$end_action()
    w$action(label="A1",weights=c(0,10,5),prob=c(0,0,0.5,0,1,0.5))
    w$end_action()

```

```

        w$end_state()
        w$end_stage()
        w$end_process()
        w$end_action()
    w$end_state()
    w$state(label="M1")
    w$action(label="A0",weights=c(0,0,0),prob=c(2,0,1))
    w$process()
    w$stage()
    w$state(label="D")
    w$action(label="A0",weights=c(0,0,1),prob=c(1,0,0.5,1,1,0.5))
    w$end_action()
    w$end_state()
    w$end_stage()
    w$stage()
    w$state(label="C0")
    w$action(label="A0",weights=c(0,0,0),prob=c(1,0,1))
    w$end_action()
    w$end_state()
    w$state(label="C1")
    w$action(label="A0",weights=c(0,0,0),prob=c(1,0,1))
    w$end_action()
    w$end_state()
    w$end_stage()
    w$stage()
    w$state(label="C0")
    w$action(label="A0",weights=c(1,4,0),prob=c(0,0,1))
    w$end_action()
    w$end_state()
    w$state(label="C1")
    w$action(label="A0",weights=c(1,4,0),prob=c(0,0,1))
    w$end_action()
    w$end_state()
    w$end_stage()
    w$end_process()
    w$end_action()
    w$end_state()
    w$end_stage()
w$end_process()
w$close_writer()

## Info about the binary files (don't have to load the model first)
if (FALSE) {
    get_bin_info_states()
    get_bin_info_actions()
}

## reset working dir
setwd(wd)

```

Description

Plot the state-expanded hypergraph of the MDP.

Usage

```
## S3 method for class 'HMDP'
plot(x, ...)
```

Arguments

```
x                The MDP model.
...              Arguments passed to plot\_hypergraph\(\).
```

Value

No return value (NULL invisible), called for side effects (plotting).

See Also

[get_hypergraph\(\)](#) and [plot_hypergraph\(\)](#) for possible arguments.

Examples

```
## Set working dir
wd <- setwd(system.file("models", package = "MDP2"))

#### A finite-horizon replacement problem ####
mdp<-load_mdp("machine1_")
plot(mdp)
plot(mdp, action_color = "label") # colors based on labels
plot(mdp, trans_labels = "state") # label transitions with target state labels
plot(mdp, trans_labels = "prob") # label transitions with transition probabilities
plot(mdp, action_color = "label", state_label = "s_id|label") # state labels are 's_id | label'
plot(mdp, state_label = "s_idx|label", radx = 0.01) # adjust radx in states
plot(mdp, state_label = "label", action_w_label = "none", action_label = "label",
      trans_labels = "s_id", radx = 0.01)

scrapValues <- c(30, 10, 5, 0) # scrap values (the values of the 4 states at stage 4)
run_value_ite(mdp, "Net reward", term_values = scrapValues)
plot(mdp, action_color = "policy") # highlight optimal policy
plot(mdp, actions_visible = "policy", state_label = "weight") # show only optimal policy

#### An infinite-horizon maintenance problem ####
mdp<-load_mdp("hct611-1_")
plot(mdp) # plot the first two stages
plot(mdp, action_color = "label") # colors based on labels
plot(mdp, action_color = "label", state_label = "s_id|label") # state labels are 's_id | label'
run_policy_ite_ave(mdp, "Net reward", "Duration")
plot(mdp, action_color = "policy") # highlight optimal policy
plot(mdp, actions_visible = "policy") # show only optimal policy
```

```
#### An infinite-horizon hierarchical replacement problem ####
library(magrittr)
mdp<-load_mdp("cow_")
hgf <- get_hypergraph(mdp)
# modify labels
dat <- hgf$nodes %>%
  dplyr::mutate(label = dplyr::case_when(
    label == "Low yield" ~ "L",
    label == "Avg yield" ~ "A",
    label == "High yield" ~ "H",
    label == "Dummy" ~ "D",
    label == "Bad genetic level" ~ "Bad",
    label == "Avg genetic level" ~ "Avg",
    label == "Good genetic level" ~ "Good",
    TRUE ~ "Error"
  ))
# assign nodes to grid ids
dat$g_id[1:3]<-85:87
dat$g_id[43:45]<-1:3
get_g_id<-function(process,stage,state) {
  if (process==0) start=18
  if (process==1) start=22
  if (process==2) start=26
  return(start + 14 * stage + state)
}
idx<-43
for (process in 0:2)
  for (stage in 0:4)
    for (state in 0:2) {
      if (stage==0 & state>0) break
      idx<-idx-1
      #cat(idx,process,stage,state,get_g_id(process,stage,state),"\n")
      dat$g_id[idx]<-get_g_id(process,stage,state)
    }
hgf$nodes <- dat
# modify labels
dat <- hgf$hyperarcs %>%
  dplyr::mutate(label = dplyr::case_when(
    label == "Replace" ~ "R",
    label == "Keep" ~ "K",
    label == "Dummy" ~ "D",
    TRUE ~ "Error"
  ),
  col = dplyr::case_when(
    label == "R" ~ "deepskyblue3",
    label == "K" ~ "darkorange1",
    label == "D" ~ "black",
    TRUE ~ "Error"
  ),
  lwd = 0.5,
  label = "")
```

```

)
hgf$hyperarcs <- dat
# plot hypergraph
oldpar <- par(mai = c(0, 0, 0, 0))
plot_hypergraph(grid_dim = c(14, 7), hgf, cex = 0.8, radx = 0.02, rady = 0.03)
par(oldpar)

## A simple finite-horizon MDP with action and transition weights
prefix <- file.path(tempdir(), "plot_transition_rewards_")
w <- binary_mdp_writer(prefix)
w$set_weights("Cost")
w$set_trans_weights(c("Reward", "Disease"))
w$process()
  w$stage()
    w$state(label = "S1")
      w$action(
        label = "A1", weights = 2, id = c(1), pr = c(1),
        trans_weights = c(20, 0.3), end = TRUE
      )
      w$action(
        label = "A2", weights = 1, id = c(0, 1), pr = c(0.3, 0.7),
        trans_weights = c(25, 0.4, 15, 0.2), end = TRUE
      )
    w$end_state()
  w$end_stage()
  w$stage()
    w$state(label = "S2")
      w$action(
        label = "A3", weights = 3, id = c(0, 1, 2), pr = c(0.5, 0.3, 0.2),
        trans_weights = c(0, 0.05, 12, 0.2, 30, 0.8), end = TRUE
      )
      w$action(
        label = "A4", weights = 2, id = c(1, 2), pr = c(0.6, 0.4),
        trans_weights = c(22, 0.35, 27, 0.7), end = TRUE
      )
    w$end_state()
  w$state(label = "S3")
    w$action(
      label = "A5", weights = 1, id = c(0, 1), pr = c(0.4, 0.6),
      trans_weights = c(5, 0, 16, 0.25), end = TRUE
    )
    w$action(
      label = "A6", weights = 4, id = c(0, 1, 2), pr = c(0.1, 0.3, 0.6),
      trans_weights = c(14, 0.15, 21, 0.45, 29, 1), end = TRUE
    )
  w$end_state()
w$end_stage()
w$stage()
  w$state(label = "S4", end = TRUE)
  w$state(label = "S5", end = TRUE)
  w$state(label = "S6", end = TRUE)
w$end_stage()

```

```

w$send_process()
w$close_writer()

mdp <- load_mdp(prefix, get_log = FALSE)
plot(mdp, action_color = "label", trans_labels = "weights", action_w_label = "weight",
      radx = 0.005, rady = 0.01)

## Reset working dir
setwd(wd)

```

plot_hypergraph

Plot parts of the state expanded hypergraph.

Description

The plot is created based on a grid (rows, cols). Each grid point is numbered from bottom to top and left to right (starting from 1), i.e. given grid point with coordinates (r, c) (where (1, 1) is the top left corner and (rows, cols) is the bottom right corner) the grid id is 'c

- 1. - rows + r'. You must assign a node to the hypergraph to a grid point (see below).

Usage

```

plot_hypergraph(
  hgf,
  grid_dim,
  show_grid = FALSE,
  radx = 0.03,
  rady = 0.05,
  state_lwd = 0.5,
  state_lty = 1,
  cex = 1,
  mar_x = 0.035,
  mar_y = 0.15,
  draw_border = FALSE,
  action_offset = 0.025,
  trans_labels = "none",
  trans_label_cex = 0.8 * cex,
  trans_label_adj = c(0.5, -0.6),
  state_label = "label",
  action_label = "label",
  action_w_label = "none",
  action_color = c("", "label", "policy"),
  actions_visible = c("all", "policy"),
  arrow = c("backward", "forward", "none"),
  connected_to = NULL,
  recalc_grid = FALSE,
  draw_time_index = c("none", "top", "bottom"),

```

```

    mdp = NULL,
    ...
)

```

Arguments

hgf	A list with the hypergraph containing two data frames, normally found using <code>get_hypergraph()</code> . The data frame nodes must have columns: <code>s_id</code> (state id), <code>g_id</code> (grid id) and <code>label</code> (node label). The data frame hyperarcs must have columns <code>s_id</code> (head node), <code>trans</code> (a list-column of tail node ids), <code>pr</code> (a list-column of transition probabilities), <code>action_weights</code> (a list-column of action weights), <code>trans_weights</code> (a list-column of transition-by-weight matrices), <code>a_idx</code> (action index), <code>label</code> (action label), <code>lwd</code> (hyperarc line width), <code>lty</code> (hyperarc line type) and <code>col</code> (hyperarc color).
grid_dim	A 2-dim vector (rows, cols) representing the size of the grid.
show_grid	If true show the grid points (good for debugging).
radx	Horizontal radius of the box.
radx	Vertical radius of the box.
state_lwd	Line width of the state/node border, or "custom" to read per-node values from <code>hgf\$nodes\$state_lwd</code> .
state_lty	Line type of the state/node border, using base R line type conventions, or "custom" to read per-node values from <code>hgf\$nodes\$state_lty</code> .
cex	Relative size of text.
mar_x	Horizontal margin.
mar_y	Vertical margin.
draw_border	If TRUE, draw a border around the plot region and report the outside and inside padding (good for debugging).
action_offset	Distance used to separate actions with the same start and trans states. Set to 0 to draw overlapping actions.
trans_labels	Transition-label mode. "none" draws no transition labels (the default); "custom" draws values from an optional <code>trans_labels</code> list-column in <code>hgf\$hyperarcs</code> ; otherwise use a -separated combination of "label", "s_id", "prob", and "weights", for example "prob weights". The older "state" spelling is treated as "label".
trans_label_cex	Relative size of transition-label text.
trans_label_adj	Position adjustment passed to <code>textempty()</code> for transition labels, drawn at the middle of each split-to-transition branch.
state_label	What to plot in states. "custom" uses a <code>state_label</code> column in <code>hgf\$nodes</code> ; otherwise use a -separated combination of "label" (state label, default), "s_id" (state id), "s_idx" (stage-based state index), and "weight" (optimal weight of the state).

action_label	What to plot near the split. One of "none", "custom" (uses an action_label column in hgf\$hyperarcs), or a -separated combination of "label" (action label, default) and "a_idx".
action_w_label	What to plot from the start state to the split. One of "none" (default), "weight", or "custom" (uses an action_w_label column in hgf\$hyperarcs).
action_color	Action coloring scheme. Default "" uses black lines. "label" uses different colors based on the action labels. "policy" highlights the current policy.
actions_visible	Action visibility mode. "all" (default) shows all actions. "policy" only shows actions in the current policy.
arrow	Arrow direction in time. "backward" (default) keeps the current arrow direction, "forward" reverses arrowheads, and "none" draws hyperarc lines without arrowheads.
connected_to	Optional vector of state ids. If supplied, plot only states reachable from these states by following visible hyperarcs forward, and trim hyperarcs and transition-level data to the remaining states.
recalc_grid	If TRUE and connected_to is supplied, recalculate the grid for the visible nodes. Nodes keep their original columns, but visible nodes within each column are placed consecutively from the top and the number of grid rows is reduced to the maximum number of visible nodes in any column.
draw_time_index	Where to draw time/stage labels. "none" (default) draws no labels, "top" draws labels above each column, and "bottom" draws labels below each column. Finite-horizon MDPs are labelled $t=0, 1, 2, \dots$; infinite-horizon MDPs are labelled $t, t+1, \dots$
mdp	The MDP model. Required if state_label contains "weight", action_color = "policy", or actions_visible = "policy".
...	Graphical parameters passed to textempty.

Value

No return value (NULL invisible), called for side effects (plotting).

See Also

[get_hypergraph\(\)](#) and [plot.HMDP\(\)](#).

Examples

```
## Set working dir
wd <- setwd(system.file("models", package = "MDP2"))

#### A finite-horizon replacement problem ####
mdp<-load_mdp("machine1_")
plot(mdp)
plot(mdp, action_color = "label") # colors based on labels
plot(mdp, trans_labels = "state") # label transitions with target state labels
```

```

plot(mdp, trans_labels = "prob") # label transitions with transition probabilities
plot(mdp, action_color = "label", state_label = "s_id|label") # state labels are 's_id | label'
plot(mdp, state_label = "s_idx|label", radx = 0.01) # adjust radx in states
plot(mdp, state_label = "label", action_w_label = "none", action_label = "label",
      trans_labels = "s_id", radx = 0.01)

scrapValues <- c(30, 10, 5, 0) # scrap values (the values of the 4 states at stage 4)
run_value_ite(mdp, "Net reward", term_values = scrapValues)
plot(mdp, action_color = "policy") # highlight optimal policy
plot(mdp, actions_visible = "policy", state_label = "weight") # show only optimal policy

#### An infinite-horizon maintenance problem ####
mdp<-load_mdp("hct611-1-")
plot(mdp) # plot the first two stages
plot(mdp, action_color = "label") # colors based on labels
plot(mdp, action_color = "label", state_label = "s_id|label") # state labels are 's_id | label'
run_policy_ite_ave(mdp, "Net reward", "Duration")
plot(mdp, action_color = "policy") # highlight optimal policy
plot(mdp, actions_visible = "policy") # show only optimal policy

#### An infinite-horizon hierarchical replacement problem ####
library(magrittr)
mdp<-load_mdp("cow-")
hgf <- get_hypergraph(mdp)
# modify labels
dat <- hgf$nodes %>%
  dplyr::mutate(label = dplyr::case_when(
    label == "Low yield" ~ "L",
    label == "Avg yield" ~ "A",
    label == "High yield" ~ "H",
    label == "Dummy" ~ "D",
    label == "Bad genetic level" ~ "Bad",
    label == "Avg genetic level" ~ "Avg",
    label == "Good genetic level" ~ "Good",
    TRUE ~ "Error"
  ))
# assign nodes to grid ids
dat$g_id[1:3]<-85:87
dat$g_id[43:45]<-1:3
get_g_id<-function(process,stage,state) {
  if (process==0) start=18
  if (process==1) start=22
  if (process==2) start=26
  return(start + 14 * stage + state)
}
idx<-43
for (process in 0:2)
  for (stage in 0:4)
    for (state in 0:2) {
      if (stage==0 & state>0) break
      idx<-idx-1
    }

```

```

        #cat(idx,process,stage,state,get_g_id(process,stage,state),"\n")
        dat$g_id[idx]<-get_g_id(process,stage,state)
    }
    hgf$nodes <- dat
    # modify labels
    dat <- hgf$hyperarcs %>%
      dplyr::mutate(label = dplyr::case_when(
        label == "Replace" ~ "R",
        label == "Keep" ~ "K",
        label == "Dummy" ~ "D",
        TRUE ~ "Error"
      ),
      col = dplyr::case_when(
        label == "R" ~ "deepskyblue3",
        label == "K" ~ "darkorange1",
        label == "D" ~ "black",
        TRUE ~ "Error"
      ),
      lwd = 0.5,
      label = ""
    )
    hgf$hyperarcs <- dat
    # plot hypergraph
    oldpar <- par(mai = c(0, 0, 0, 0))
    plot_hypergraph(grid_dim = c(14, 7), hgf, cex = 0.8, radx = 0.02, rady = 0.03)
    par(oldpar)

```

```

## A simple finite-horizon MDP with action and transition weights
prefix <- file.path(tempdir(), "plot_transition_rewards_")
w <- binary_mdp_writer(prefix)
w$set_weights("Cost")
w$set_trans_weights(c("Reward", "Disease"))
w$process()
  w$stage()
    w$state(label = "S1")
      w$action(
        label = "A1", weights = 2, id = c(1), pr = c(1),
        trans_weights = c(20, 0.3), end = TRUE
      )
      w$action(
        label = "A2", weights = 1, id = c(0, 1), pr = c(0.3, 0.7),
        trans_weights = c(25, 0.4, 15, 0.2), end = TRUE
      )
    w$end_state()
  w$end_stage()
w$stage()
  w$state(label = "S2")
    w$action(
      label = "A3", weights = 3, id = c(0, 1, 2), pr = c(0.5, 0.3, 0.2),
      trans_weights = c(0, 0.05, 12, 0.2, 30, 0.8), end = TRUE
    )
    w$action(

```

```

        label = "A4", weights = 2, id = c(1, 2), pr = c(0.6, 0.4),
        trans_weights = c(22, 0.35, 27, 0.7), end = TRUE
    )
    w$end_state()
    w$state(label = "S3")
    w$action(
        label = "A5", weights = 1, id = c(0, 1), pr = c(0.4, 0.6),
        trans_weights = c(5, 0, 16, 0.25), end = TRUE
    )
    w$action(
        label = "A6", weights = 4, id = c(0, 1, 2), pr = c(0.1, 0.3, 0.6),
        trans_weights = c(14, 0.15, 21, 0.45, 29, 1), end = TRUE
    )
    w$end_state()
w$end_stage()
w$stage()
    w$state(label = "S4", end = TRUE)
    w$state(label = "S5", end = TRUE)
    w$state(label = "S6", end = TRUE)
    w$end_stage()
w$end_process()
w$close_writer()

mdp <- load_mdp(prefix, get_log = FALSE)
plot(mdp, action_color = "label", trans_labels = "weights", action_w_label = "weight",
     radx = 0.005, rady = 0.01)

## Reset working dir
setwd(wd)

```

random_hmdp

Generate a "random" HMDP stored in a set of binary files.

Description

Generate a "random" HMDP stored in a set of binary files.

Usage

```

random_hmdp(
  prefix = "",
  levels = 3,
  time_horizon = c(Inf, 3, 4),
  states = c(2, 4, 5),
  actions = c(1, 2),
  child_process_pr = 0.5,
  external_process_pr = 0,
  rewards = c(0, 100),
  durations = c(1, 10),

```

```

    reward_name = "Reward",
    duration_name = "Duration"
)

```

Arguments

prefix	A character string with the prefix added to the file(s).
levels	Maximum number of levels. Set child_process_pr = 1 if want exact this number of levels.
time_horizon	The time horizon for each level (vector). For the founder the time-horizon can be Inf.
states	Number of states at each stage at a given level (vector of length levels)
actions	Min and max number of actions at a state.
child_process_pr	Probability of creating a child process when define action.
external_process_pr	Probability of creating an external process given that we create a child process. Only works if levels>2 and and currently does not generate external processes which include external processes.
rewards	Min and max reward used.
durations	Min and max duration used.
reward_name	Weight name used for reward.
duration_name	Weight name used for duration.

Value

The file prefix (character).

run_calc_weights	<i>Calculate weights based on current policy. Normally run after an optimal policy has been found.</i>
------------------	--

Description

Calculate weights based on current policy. Normally run after an optimal policy has been found.

Usage

```

run_calc_weights(
  mdp,
  w_lbl,
  criterion = "expected",
  dur_lbl = NULL,
  rate = 0,
  rate_base = 1,

```

```

discount_factor = NULL,
term_values = NULL,
discount_method = "continuous"
)

```

Arguments

mdp	The MDP loaded using <code>load_mdp()</code> .
w_lbl	The label of the weight we consider.
criterion	The Bellman operator shortcut. If expected use expected weights, if discount use discounted expected weights, if average use average expected weights, if min use minimum-successor weights, if max use maximum-successor weights, if second_moment use the second moment of total accumulated weight, and if variance use the law-of-total-variance recursion under the current policy.
dur_lbl	The label of the duration/time such that discount rates can be calculated.
rate	The interest rate.
rate_base	The time-horizon the rate is valid over.
discount_factor	The discount rate for one time unit. If specified rate and rate_base are not used to calculate the discount rate.
term_values	The terminal values used (values of the last stage in the MDP).
discount_method	Either 'continuous' or 'discrete', corresponding to discount factor $\exp(-\text{rate}/\text{rate_base})$ or $1/(1 + \text{rate}/\text{rate_base})$, respectively. Only used if discount_factor is NULL.

Value

Nothing.

Examples

```

## Set working dir
wd <- setwd(tempdir())

# Create the small machine replacement problem used as an example in L.R. Nielsen and A.R.
# Kristensen. Finding the K best policies in a finite-horizon Markov decision process. European
# Journal of Operational Research, 175(2):1164-1179, 2006. doi:10.1016/j.ejor.2005.06.011.

## Create the MDP using a dummy replacement node
prefix<-"machine1_"
w <- binary_mdp_writer(prefix)
w$set_weights(c("Net reward"))
w$process()
w$stage() # stage n=0
w$state(label="Dummy") # v=(0,0)
w$action(label="buy", weights=-100, prob=c(1,0,0.7, 1,1,0.3), end=TRUE)
w$end_state()

```

```

w$end_stage()
w$stage() # stage n=1
  w$state(label="good") # v=(1,0)
    w$action(label="mt", weights=55, prob=c(1,0,1), end=TRUE)
    w$action(label="nmt", weights=70, prob=c(1,0,0.6, 1,1,0.4), end=TRUE)
  w$end_state()
  w$state(label="average") # v=(1,1)
    w$action(label="mt", weights=40, prob=c(1,0,1), end=TRUE)
    w$action(label="nmt", weights=50, prob=c(1,1,0.6, 1,2,0.4), end=TRUE)
  w$end_state()
w$end_stage()
w$stage() # stage n=2
  w$state(label="good") # v=(2,0)
    w$action(label="mt", weights=55, prob=c(1,0,1), end=TRUE)
    w$action(label="nmt", weights=70, prob=c(1,0,0.5, 1,1,0.5), end=TRUE)
  w$end_state()
  w$state(label="average") # v=(2,1)
    w$action(label="mt", weights=40, prob=c(1,0,1), end=TRUE)
    w$action(label="nmt", weights=50, prob=c(1,1,0.5, 1,2,0.5), end=TRUE)
  w$end_state()
  w$state(label="not working") # v=(2,2)
    w$action(label="mt", weights=30, prob=c(1,0,1), end=TRUE)
    w$action(label="rep", weights=5, prob=c(1,3,1), end=TRUE)
  w$end_state()
w$end_stage()
w$stage() # stage n=3
  w$state(label="good") # v=(3,0)
    w$action(label="mt", weights=55, prob=c(1,0,1), end=TRUE)
    w$action(label="nmt", weights=70, prob=c(1,0,0.2, 1,1,0.8), end=TRUE)
  w$end_state()
  w$state(label="average") # v=(3,1)
    w$action(label="mt", weights=40, prob=c(1,0,1), end=TRUE)
    w$action(label="nmt", weights=50, prob=c(1,1,0.2, 1,2,0.8), end=TRUE)
  w$end_state()
  w$state(label="not working") # v=(3,2)
    w$action(label="mt", weights=30, prob=c(1,0,1), end=TRUE)
    w$action(label="rep", weights=5, prob=c(1,3,1), end=TRUE)
  w$end_state()
  w$state(label="replaced") # v=(3,3)
    w$action(label="Dummy", weights=0, prob=c(1,3,1), end=TRUE)
  w$end_state()
w$end_stage()
w$stage() # stage n=4
  w$state(label="good", end=TRUE) # v=(4,0)
  w$state(label="average", end=TRUE) # v=(4,1)
  w$state(label="not working", end=TRUE) # v=(4,2)
  w$state(label="replaced", end=TRUE) # v=(4,3)
w$end_stage()
w$end_process()
w$close_writer()

## Load the model into memory
mdp<-load_mdp(prefix)

```

```

mdp
plot(mdp)

get_info(mdp, with_list = FALSE)
get_info(mdp, with_list = FALSE, df_level = "action", as_strings_actions = TRUE)
get_info(mdp, with_list = FALSE, df_level = "action", as_strings_actions = FALSE)

## Perform value iteration
w<-"Net reward"          # label of the weight we want to optimize
scrapValues<-c(30,10,5,0) # scrap values (the values of the 4 states at stage 4)
run_value_ite(mdp, w, term_values=scrapValues)
get_policy(mdp)          # optimal policy

## Calculate the weights of the policy always to maintain
library(magrittr)
policy <- get_info(mdp, with_list = FALSE, df_level = "action")$df %>%
  dplyr::filter(label_action == "mt") %>%
  dplyr::select(s_id, a_idx)
set_policy(mdp, policy)
run_calc_weights(mdp, w, term_values=scrapValues)
get_policy(mdp)

# The example given in L.R. Nielsen and A.R. Kristensen. Finding the K best
# policies in a finite-horizon Markov decision process. European Journal of
# Operational Research, 175(2):1164-1179, 2006. doi:10.1016/j.ejor.2005.06.011,
# does actually not have any dummy replacement node as in the MDP above. The same
# model can be created using a single dummy node at the end of the process.

## Create the MDP using a single dummy node
prefix<-"machine2_"
w <- binary_mdp_writer(prefix)
w$set_weights(c("Net reward"))
w$process()
  w$stage() # stage n=0
    w$state(label="Dummy") # v=(0,0)
      w$action(label="buy", weights=-100, prob=c(1,0,0.7, 1,1,0.3), end=TRUE)
    w$end_state()
  w$end_stage()
  w$stage() # stage n=1
    w$state(label="good") # v=(1,0)
      w$action(label="mt", weights=55, prob=c(1,0,1), end=TRUE)
      w$action(label="nmt", weights=70, prob=c(1,0,0.6, 1,1,0.4), end=TRUE)
    w$end_state()
    w$state(label="average") # v=(1,1)
      w$action(label="mt", weights=40, prob=c(1,0,1), end=TRUE)
      w$action(label="nmt", weights=50, prob=c(1,1,0.6, 1,2,0.4), end=TRUE)
    w$end_state()
  w$end_stage()
  w$stage() # stage n=2
    w$state(label="good") # v=(2,0)
      w$action(label="mt", weights=55, prob=c(1,0,1), end=TRUE)

```

```

        w$action(label="nmt", weights=70, prob=c(1,0,0.5, 1,1,0.5), end=TRUE)
    w$end_state()
    w$state(label="average")      # v=(2,1)
        w$action(label="mt", weights=40, prob=c(1,0,1), end=TRUE)
        w$action(label="nmt", weights=50, prob=c(1,1,0.5, 1,2,0.5), end=TRUE)
    w$end_state()
    w$state(label="not working")  # v=(2,2)
        w$action(label="mt", weights=30, prob=c(1,0,1), end=TRUE)
    w$action(label="rep", weights=5, prob=c(3,12,1), end=TRUE) # transition to s_id=12 (Dummy)
    w$end_state()
w$end_stage()
w$stage() # stage n=3
    w$state(label="good")        # v=(3,0)
        w$action(label="mt", weights=55, prob=c(1,0,1), end=TRUE)
        w$action(label="nmt", weights=70, prob=c(1,0,0.2, 1,1,0.8), end=TRUE)
    w$end_state()
    w$state(label="average")      # v=(3,1)
        w$action(label="mt", weights=40, prob=c(1,0,1), end=TRUE)
        w$action(label="nmt", weights=50, prob=c(1,1,0.2, 1,2,0.8), end=TRUE)
    w$end_state()
    w$state(label="not working")  # v=(3,2)
        w$action(label="mt", weights=30, prob=c(1,0,1), end=TRUE)
        w$action(label="rep", weights=5, prob=c(3,12,1), end=TRUE)
    w$end_state()
w$end_stage()
w$stage() # stage n=4
    w$state(label="good")        # v=(4,0)
        w$action(label="rep", weights=30, prob=c(1,0,1), end=TRUE)
    w$end_state()
    w$state(label="average")      # v=(4,1)
        w$action(label="rep", weights=10, prob=c(1,0,1), end=TRUE)
    w$end_state()
    w$state(label="not working")  # v=(4,2)
        w$action(label="rep", weights=5, prob=c(1,0,1), end=TRUE)
    w$end_state()
w$end_stage()
w$stage() # stage n=5
    w$state(label="Dummy", end=TRUE)      # v=(5,0)
w$end_stage()
w$end_process()
w$close_writer()

## Have a look at the state-expanded hypergraph
mdp<-load_mdp(prefix)
mdp
plot(mdp)

get_info(mdp, with_list = FALSE)
get_info(mdp, with_list = FALSE, df_level = "action", as_strings_actions = TRUE)
get_info(mdp, with_list = FALSE, df_level = "action", as_strings_actions = FALSE)

## Perform value iteration
w<-"Net reward"          # label of the weight we want to optimize

```

```

run_value_ite(mdp, w, term_values = 0)
get_policy(mdp)      # optimal policy

## Calculate the weights of the policy always to maintain
library(magrittr)
policy <- get_info(mdp, with_list = FALSE, df_level = "action")$df %>%
  dplyr::filter(label_action == "mt") %>%
  dplyr::select(s_id, a_idx)
set_policy(mdp, policy)
run_calc_weights(mdp, w, term_values=scrapValues)
get_policy(mdp)

## Reset working dir
setwd(wd)

```

run_policy_ite_ave	<i>Perform policy iteration using the average expected-weight Bellman operator on the MDP.</i>
--------------------	--

Description

The policy can afterwards be received using functions `get_policy` and `get_policy_w`.

Usage

```

run_policy_ite_ave(
  mdp,
  w,
  dur,
  max_ite = 100,
  objective = c("max", "min"),
  get_log = TRUE
)

```

Arguments

mdp	The MDP loaded using <code>load_mdp()</code> .
w	The label of the weight we optimize.
dur	The label of the duration/time such that discount rates can be calculated.
max_ite	Max number of iterations. If the model does not satisfy the unichain assumption the algorithm may loop.
objective	Optimize by maximizing ("max") or minimizing ("min") the Bellman value.
get_log	Output the log messages.

Value

The optimal gain (g) calculated.

See Also

[get_policy\(\)](#).

run_policy_ite_discount

Perform policy iteration using the discounted expected-weight Bellman operator on the MDP.

Description

The policy can afterwards be received using functions `get_policy` and `get_policy_w`.

Usage

```
run_policy_ite_discount(
    mdp,
    w,
    dur,
    rate = 0,
    rate_base = 1,
    discount_factor = NULL,
    max_ite = 100,
    discount_method = "continuous",
    objective = c("max", "min"),
    get_log = TRUE
)
```

Arguments

<code>mdp</code>	The MDP loaded using load_mdp() .
<code>w</code>	The label of the weight we optimize.
<code>dur</code>	The label of the duration/time such that discount rates can be calculated.
<code>rate</code>	The interest rate.
<code>rate_base</code>	The time-horizon the rate is valid over.
<code>discount_factor</code>	The discount rate for one time unit. If specified <code>rate</code> and <code>rate_base</code> are not used to calculate the discount rate.
<code>max_ite</code>	Max number of iterations. If the model does not satisfy the unichain assumption the algorithm may loop.
<code>discount_method</code>	Either 'continuous' or 'discrete', corresponding to discount factor $\exp(-\text{rate}/\text{rate_base})$ or $1/(1 + \text{rate}/\text{rate_base})$, respectively. Only used if <code>discount_factor</code> is NULL.
<code>objective</code>	Optimize by maximizing ("max") or minimizing ("min") the Bellman value.
<code>get_log</code>	Output the log messages.

Value

Nothing.

See Also

[get_policy\(\)](#).

run_value_ite

Perform value iteration on the MDP.

Description

If the MDP has a finite time-horizon then arguments times and eps are ignored.

Usage

```
run_value_ite(
  mdp,
  w,
  dur = NULL,
  rate = 0,
  rate_base = 1,
  discount_factor = NULL,
  max_ite = 100,
  eps = 1e-05,
  term_values = NULL,
  g = NULL,
  objective = c("max", "min"),
  bellman_op = c("auto", "expected", "discount", "average", "min", "max",
    "second_moment"),
  get_log = TRUE,
  discount_method = "continuous"
)
```

Arguments

mdp	The MDP loaded using load_mdp() .
w	The label of the weight we optimize.
dur	The label of the duration/time such that discount rates can be calculated.
rate	Interest rate.
rate_base	The time-horizon the rate is valid over.
discount_factor	The discount rate for one time unit. If specified rate and rate_base are not used to calculate the discount rate.
max_ite	The max number of iterations value iteration is performed.

eps	Stopping tolerance. If $\max(w(t)-w(t+1)) < \text{eps}$ then stop the algorithm, i.e the policy becomes epsilon optimal (see Puterman p161).
term_values	The terminal values used (values of the last stage in the MDP).
g	Average weight. If specified then do a single iteration using the update equations under the average expected-weight Bellman operator with the specified g value.
objective	Optimize by maximizing ("max") or minimizing ("min") the Bellman value.
bellman_op	Bellman operator. Use "auto" for existing behavior, "min" for the minimum-successor operator, "max" for the maximum-successor operator, or "second_moment" for the second moment of total accumulated weight.
get_log	Output the log messages.
discount_method	Either 'continuous' or 'discrete', corresponding to discount factor $\exp(-\text{rate}/\text{rate_base})$ or $1/(1 + \text{rate}/\text{rate_base})$, respectively. Only used if discount_factor is NULL.

Value

NULL (invisible)

References

Puterman, M. Markov Decision Processes, Wiley-Interscience, 1994.

Examples

```
## Set working dir
wd <- setwd(tempdir())

# Create the small machine replacement problem used as an example in L.R. Nielsen and A.R.
# Kristensen. Finding the K best policies in a finite-horizon Markov decision process. European
# Journal of Operational Research, 175(2):1164-1179, 2006. doi:10.1016/j.ejor.2005.06.011.

## Create the MDP using a dummy replacement node
prefix<-"machine1_"
w <- binary_mdp_writer(prefix)
w$set_weights(c("Net reward"))
w$process()
  w$stage() # stage n=0
    w$state(label="Dummy") # v=(0,0)
      w$action(label="buy", weights=-100, prob=c(1,0,0.7, 1,1,0.3), end=TRUE)
    w$end_state()
  w$end_stage()
  w$stage() # stage n=1
    w$state(label="good") # v=(1,0)
      w$action(label="mt", weights=55, prob=c(1,0,1), end=TRUE)
      w$action(label="nmt", weights=70, prob=c(1,0,0.6, 1,1,0.4), end=TRUE)
    w$end_state()
  w$state(label="average") # v=(1,1)
    w$action(label="mt", weights=40, prob=c(1,0,1), end=TRUE)
```

```

        w$action(label="nmt", weights=50, prob=c(1,1,0.6, 1,2,0.4), end=TRUE)
    w$end_state()
w$end_stage()
w$stage() # stage n=2
    w$state(label="good") # v=(2,0)
        w$action(label="mt", weights=55, prob=c(1,0,1), end=TRUE)
        w$action(label="nmt", weights=70, prob=c(1,0,0.5, 1,1,0.5), end=TRUE)
    w$end_state()
    w$state(label="average") # v=(2,1)
        w$action(label="mt", weights=40, prob=c(1,0,1), end=TRUE)
        w$action(label="nmt", weights=50, prob=c(1,1,0.5, 1,2,0.5), end=TRUE)
    w$end_state()
    w$state(label="not working") # v=(2,2)
        w$action(label="mt", weights=30, prob=c(1,0,1), end=TRUE)
        w$action(label="rep", weights=5, prob=c(1,3,1), end=TRUE)
    w$end_state()
w$end_stage()
w$stage() # stage n=3
    w$state(label="good") # v=(3,0)
        w$action(label="mt", weights=55, prob=c(1,0,1), end=TRUE)
        w$action(label="nmt", weights=70, prob=c(1,0,0.2, 1,1,0.8), end=TRUE)
    w$end_state()
    w$state(label="average") # v=(3,1)
        w$action(label="mt", weights=40, prob=c(1,0,1), end=TRUE)
        w$action(label="nmt", weights=50, prob=c(1,1,0.2, 1,2,0.8), end=TRUE)
    w$end_state()
    w$state(label="not working") # v=(3,2)
        w$action(label="mt", weights=30, prob=c(1,0,1), end=TRUE)
        w$action(label="rep", weights=5, prob=c(1,3,1), end=TRUE)
    w$end_state()
    w$state(label="replaced") # v=(3,3)
        w$action(label="Dummy", weights=0, prob=c(1,3,1), end=TRUE)
    w$end_state()
w$end_stage()
w$stage() # stage n=4
    w$state(label="good", end=TRUE) # v=(4,0)
    w$state(label="average", end=TRUE) # v=(4,1)
    w$state(label="not working", end=TRUE) # v=(4,2)
    w$state(label="replaced", end=TRUE) # v=(4,3)
w$end_stage()
w$end_process()
w$close_writer()

## Load the model into memory
mdp<-load_mdp(prefix)
mdp
plot(mdp)

get_info(mdp, with_list = FALSE)
get_info(mdp, with_list = FALSE, df_level = "action", as_strings_actions = TRUE)
get_info(mdp, with_list = FALSE, df_level = "action", as_strings_actions = FALSE)

## Perform value iteration

```

```

w<-"Net reward"          # label of the weight we want to optimize
scrapValues<-c(30,10,5,0) # scrap values (the values of the 4 states at stage 4)
run_value_ite(mdp, w, term_values=scrapValues)
get_policy(mdp)          # optimal policy

## Calculate the weights of the policy always to maintain
library(magrittr)
policy <- get_info(mdp, with_list = FALSE, df_level = "action")$df %>%
  dplyr::filter(label_action == "mt") %>%
  dplyr::select(s_id, a_idx)
set_policy(mdp, policy)
run_calc_weights(mdp, w, term_values=scrapValues)
get_policy(mdp)

# The example given in L.R. Nielsen and A.R. Kristensen. Finding the K best
# policies in a finite-horizon Markov decision process. European Journal of
# Operational Research, 175(2):1164-1179, 2006. doi:10.1016/j.ejor.2005.06.011,
# does actually not have any dummy replacement node as in the MDP above. The same
# model can be created using a single dummy node at the end of the process.

## Create the MDP using a single dummy node
prefix<-"machine2_"
w <- binary_mdp_writer(prefix)
w$set_weights(c("Net reward"))
w$process()
  w$stage() # stage n=0
    w$state(label="Dummy") # v=(0,0)
      w$action(label="buy", weights=-100, prob=c(1,0,0.7, 1,1,0.3), end=TRUE)
    w$end_state()
  w$end_stage()
  w$stage() # stage n=1
    w$state(label="good") # v=(1,0)
      w$action(label="mt", weights=55, prob=c(1,0,1), end=TRUE)
      w$action(label="nmt", weights=70, prob=c(1,0,0.6, 1,1,0.4), end=TRUE)
    w$end_state()
    w$state(label="average") # v=(1,1)
      w$action(label="mt", weights=40, prob=c(1,0,1), end=TRUE)
      w$action(label="nmt", weights=50, prob=c(1,1,0.6, 1,2,0.4), end=TRUE)
    w$end_state()
  w$end_stage()
  w$stage() # stage n=2
    w$state(label="good") # v=(2,0)
      w$action(label="mt", weights=55, prob=c(1,0,1), end=TRUE)
      w$action(label="nmt", weights=70, prob=c(1,0,0.5, 1,1,0.5), end=TRUE)
    w$end_state()
    w$state(label="average") # v=(2,1)
      w$action(label="mt", weights=40, prob=c(1,0,1), end=TRUE)
      w$action(label="nmt", weights=50, prob=c(1,1,0.5, 1,2,0.5), end=TRUE)
    w$end_state()
    w$state(label="not working") # v=(2,2)
      w$action(label="mt", weights=30, prob=c(1,0,1), end=TRUE)

```

```

    w$action(label="rep", weights=5, prob=c(3,12,1), end=TRUE) # transition to s_id=12 (Dummy)
  w$end_state()
w$end_stage()
w$stage() # stage n=3
  w$state(label="good") # v=(3,0)
    w$action(label="mt", weights=55, prob=c(1,0,1), end=TRUE)
    w$action(label="nmt", weights=70, prob=c(1,0,0.2, 1,1,0.8), end=TRUE)
  w$end_state()
  w$state(label="average") # v=(3,1)
    w$action(label="mt", weights=40, prob=c(1,0,1), end=TRUE)
    w$action(label="nmt", weights=50, prob=c(1,1,0.2, 1,2,0.8), end=TRUE)
  w$end_state()
  w$state(label="not working") # v=(3,2)
    w$action(label="mt", weights=30, prob=c(1,0,1), end=TRUE)
    w$action(label="rep", weights=5, prob=c(3,12,1), end=TRUE)
  w$end_state()
w$end_stage()
w$stage() # stage n=4
  w$state(label="good") # v=(4,0)
    w$action(label="rep", weights=30, prob=c(1,0,1), end=TRUE)
  w$end_state()
  w$state(label="average") # v=(4,1)
    w$action(label="rep", weights=10, prob=c(1,0,1), end=TRUE)
  w$end_state()
  w$state(label="not working") # v=(4,2)
    w$action(label="rep", weights=5, prob=c(1,0,1), end=TRUE)
  w$end_state()
w$end_stage()
w$stage() # stage n=5
  w$state(label="Dummy", end=TRUE) # v=(5,0)
w$end_stage()
w$end_process()
w$close_writer()

## Have a look at the state-expanded hypergraph
mdp<-load_mdp(prefix)
mdp
plot(mdp)

get_info(mdp, with_list = FALSE)
get_info(mdp, with_list = FALSE, df_level = "action", as_strings_actions = TRUE)
get_info(mdp, with_list = FALSE, df_level = "action", as_strings_actions = FALSE)

## Perform value iteration
w<-"Net reward" # label of the weight we want to optimize
run_value_ite(mdp, w, term_values = 0)
get_policy(mdp) # optimal policy

## Calculate the weights of the policy always to maintain
library(magrittr)
policy <- get_info(mdp, with_list = FALSE, df_level = "action")$df %>%
  dplyr::filter(label_action == "mt") %>%
  dplyr::select(s_id, a_idx)

```

```
set_policy(mdp, policy)
run_calc_weights(mdp, w, term_values=scrapValues)
get_policy(mdp)

## Reset working dir
setwd(wd)
```

save_mdp	<i>Save the MDP to binary files</i>
----------	-------------------------------------

Description

Currently do not save external files.

Usage

```
save_mdp(mdp, prefix = "", get_log = TRUE)
```

Arguments

mdp	The MDP loaded using load_mdp() .
prefix	A character string with the prefix added to bin_names. Used to identify a specific model.
get_log	Output the log as a message.

Value

???

set_action_weight	<i>Set an action weight in a loaded MDP</i>
-------------------	---

Description

Modifies one action-level weight in the in-memory model. Use [save_mdp\(\)](#) if the modified model should be written to binary files.

Usage

```
set_action_weight(mdp, weight, s_id, a_idx, weight_name)
```

Arguments

mdp	The MDP loaded using <code>load_mdp()</code> or created using <code>memory_mdp_writer()</code> .
weight	A finite numeric scalar with the new action weight.
s_id	Zero-based state index.
a_idx	Zero-based action index within the state.
weight_name	The name of an action-level weight.

Value

NULL, invisibly.

set_policy	<i>Modify the current policy by setting policy action of states.</i>
------------	--

Description

If the policy does not contain all states then the actions from the previous optimal policy are used.

Usage

```
set_policy(mdp, policy)
```

Arguments

mdp	The MDP loaded using <code>load_mdp()</code> .
policy	A data frame with two columns state id <code>s_id</code> and action index <code>a_idx</code> .

Value

NULL (invisible)

Examples

```
## Set working dir
wd <- setwd(tempdir())

# Create the small machine replacement problem used as an example in L.R. Nielsen and A.R.
# Kristensen. Finding the K best policies in a finite-horizon Markov decision process. European
# Journal of Operational Research, 175(2):1164-1179, 2006. doi:10.1016/j.ejor.2005.06.011.

## Create the MDP using a dummy replacement node
prefix<-"machine1_"
w <- binary_mdp_writer(prefix)
w$set_weights(c("Net reward"))
w$process()
  w$stage() # stage n=0
    w$state(label="Dummy") # v=(0,0)
```

```

        w$action(label="buy", weights=-100, prob=c(1,0,0.7, 1,1,0.3), end=TRUE)
    w$end_state()
w$end_stage()
w$stage() # stage n=1
    w$state(label="good") # v=(1,0)
        w$action(label="mt", weights=55, prob=c(1,0,1), end=TRUE)
        w$action(label="nmt", weights=70, prob=c(1,0,0.6, 1,1,0.4), end=TRUE)
    w$end_state()
    w$state(label="average") # v=(1,1)
        w$action(label="mt", weights=40, prob=c(1,0,1), end=TRUE)
        w$action(label="nmt", weights=50, prob=c(1,1,0.6, 1,2,0.4), end=TRUE)
    w$end_state()
w$end_stage()
w$stage() # stage n=2
    w$state(label="good") # v=(2,0)
        w$action(label="mt", weights=55, prob=c(1,0,1), end=TRUE)
        w$action(label="nmt", weights=70, prob=c(1,0,0.5, 1,1,0.5), end=TRUE)
    w$end_state()
    w$state(label="average") # v=(2,1)
        w$action(label="mt", weights=40, prob=c(1,0,1), end=TRUE)
        w$action(label="nmt", weights=50, prob=c(1,1,0.5, 1,2,0.5), end=TRUE)
    w$end_state()
    w$state(label="not working") # v=(2,2)
        w$action(label="mt", weights=30, prob=c(1,0,1), end=TRUE)
        w$action(label="rep", weights=5, prob=c(1,3,1), end=TRUE)
    w$end_state()
w$end_stage()
w$stage() # stage n=3
    w$state(label="good") # v=(3,0)
        w$action(label="mt", weights=55, prob=c(1,0,1), end=TRUE)
        w$action(label="nmt", weights=70, prob=c(1,0,0.2, 1,1,0.8), end=TRUE)
    w$end_state()
    w$state(label="average") # v=(3,1)
        w$action(label="mt", weights=40, prob=c(1,0,1), end=TRUE)
        w$action(label="nmt", weights=50, prob=c(1,1,0.2, 1,2,0.8), end=TRUE)
    w$end_state()
    w$state(label="not working") # v=(3,2)
        w$action(label="mt", weights=30, prob=c(1,0,1), end=TRUE)
        w$action(label="rep", weights=5, prob=c(1,3,1), end=TRUE)
    w$end_state()
    w$state(label="replaced") # v=(3,3)
        w$action(label="Dummy", weights=0, prob=c(1,3,1), end=TRUE)
    w$end_state()
w$end_stage()
w$stage() # stage n=4
    w$state(label="good", end=TRUE) # v=(4,0)
    w$state(label="average", end=TRUE) # v=(4,1)
    w$state(label="not working", end=TRUE) # v=(4,2)
    w$state(label="replaced", end=TRUE) # v=(4,3)
w$end_stage()
w$end_process()
w$close_writer()

```

```

## Load the model into memory
mdp<-load_mdp(prefix)
mdp
plot(mdp)

get_info(mdp, with_list = FALSE)
get_info(mdp, with_list = FALSE, df_level = "action", as_strings_actions = TRUE)
get_info(mdp, with_list = FALSE, df_level = "action", as_strings_actions = FALSE)

## Perform value iteration
w<-"Net reward"          # label of the weight we want to optimize
scrapValues<-c(30,10,5,0) # scrap values (the values of the 4 states at stage 4)
run_value_ite(mdp, w, term_values=scrapValues)
get_policy(mdp)          # optimal policy

## Calculate the weights of the policy always to maintain
library(magrittr)
policy <- get_info(mdp, with_list = FALSE, df_level = "action")$df %>%
  dplyr::filter(label_action == "mt") %>%
  dplyr::select(s_id, a_idx)
set_policy(mdp, policy)
run_calc_weights(mdp, w, term_values=scrapValues)
get_policy(mdp)

# The example given in L.R. Nielsen and A.R. Kristensen. Finding the K best
# policies in a finite-horizon Markov decision process. European Journal of
# Operational Research, 175(2):1164-1179, 2006. doi:10.1016/j.ejor.2005.06.011,
# does actually not have any dummy replacement node as in the MDP above. The same
# model can be created using a single dummy node at the end of the process.

## Create the MDP using a single dummy node
prefix<-"machine2_"
w <- binary_mdp_writer(prefix)
w$set_weights(c("Net reward"))
w$process()
  w$stage() # stage n=0
    w$state(label="Dummy") # v=(0,0)
      w$action(label="buy", weights=-100, prob=c(1,0,0.7, 1,1,0.3), end=TRUE)
    w$end_state()
  w$end_stage()
  w$stage() # stage n=1
    w$state(label="good") # v=(1,0)
      w$action(label="mt", weights=55, prob=c(1,0,1), end=TRUE)
      w$action(label="nmt", weights=70, prob=c(1,0,0.6, 1,1,0.4), end=TRUE)
    w$end_state()
    w$state(label="average") # v=(1,1)
      w$action(label="mt", weights=40, prob=c(1,0,1), end=TRUE)
      w$action(label="nmt", weights=50, prob=c(1,1,0.6, 1,2,0.4), end=TRUE)
    w$end_state()
  w$end_stage()
  w$stage() # stage n=2

```

```

w$state(label="good")          # v=(2,0)
  w$action(label="mt", weights=55, prob=c(1,0,1), end=TRUE)
  w$action(label="nmt", weights=70, prob=c(1,0,0.5, 1,1,0.5), end=TRUE)
w$end_state()
w$state(label="average")        # v=(2,1)
  w$action(label="mt", weights=40, prob=c(1,0,1), end=TRUE)
  w$action(label="nmt", weights=50, prob=c(1,1,0.5, 1,2,0.5), end=TRUE)
w$end_state()
w$state(label="not working")    # v=(2,2)
  w$action(label="mt", weights=30, prob=c(1,0,1), end=TRUE)
w$action(label="rep", weights=5, prob=c(3,12,1), end=TRUE) # transition to s_id=12 (Dummy)
w$end_state()
w$end_stage()
w$stage() # stage n=3
  w$state(label="good")          # v=(3,0)
    w$action(label="mt", weights=55, prob=c(1,0,1), end=TRUE)
    w$action(label="nmt", weights=70, prob=c(1,0,0.2, 1,1,0.8), end=TRUE)
  w$end_state()
  w$state(label="average")        # v=(3,1)
    w$action(label="mt", weights=40, prob=c(1,0,1), end=TRUE)
    w$action(label="nmt", weights=50, prob=c(1,1,0.2, 1,2,0.8), end=TRUE)
  w$end_state()
  w$state(label="not working")    # v=(3,2)
    w$action(label="mt", weights=30, prob=c(1,0,1), end=TRUE)
    w$action(label="rep", weights=5, prob=c(3,12,1), end=TRUE)
  w$end_state()
w$end_stage()
w$stage() # stage n=4
  w$state(label="good")          # v=(4,0)
    w$action(label="rep", weights=30, prob=c(1,0,1), end=TRUE)
  w$end_state()
  w$state(label="average")        # v=(4,1)
    w$action(label="rep", weights=10, prob=c(1,0,1), end=TRUE)
  w$end_state()
  w$state(label="not working")    # v=(4,2)
    w$action(label="rep", weights=5, prob=c(1,0,1), end=TRUE)
  w$end_state()
w$end_stage()
w$stage() # stage n=5
  w$state(label="Dummy", end=TRUE) # v=(5,0)
w$end_stage()
w$end_process()
w$close_writer()

## Have a look at the state-expanded hypergraph
mdp<-load_mdp(prefix)
mdp
plot(mdp)

get_info(mdp, with_list = FALSE)
get_info(mdp, with_list = FALSE, df_level = "action", as_strings_actions = TRUE)
get_info(mdp, with_list = FALSE, df_level = "action", as_strings_actions = FALSE)

```

```
## Perform value iteration
w<-"Net reward"          # label of the weight we want to optimize
run_value_ite(mdp, w, term_values = 0)
get_policy(mdp)          # optimal policy

## Calculate the weights of the policy always to maintain
library(magrittr)
policy <- get_info(mdp, with_list = FALSE, df_level = "action")$df %>%
  dplyr::filter(label_action == "mt") %>%
  dplyr::select(s_id, a_idx)
set_policy(mdp, policy)
run_calc_weights(mdp, w, term_values=scrapValues)
get_policy(mdp)

## Reset working dir
setwd(wd)
```

set_transition_weight *Set a transition weight in a loaded MDP*

Description

Modifies one transition-level weight in the in-memory model. Use [get_info\(\)](#) to inspect the transition ordering within an action, and use [save_mdp\(\)](#) if the modified model should be written to binary files.

Usage

```
set_transition_weight(mdp, weight, s_id, a_idx, transition_idx, weight_name)
```

Arguments

mdp	The MDP loaded using load_mdp() or created using memory_mdp_writer() .
weight	A finite numeric scalar with the new transition weight.
s_id	Zero-based state index.
a_idx	Zero-based action index within the state.
transition_idx	Zero-based transition index within the action.
weight_name	The name of a transition-level weight.

Value

NULL, invisibly.

Index

binary_action_writer, 2
binary_mdp_writer, 6
binary_mdp_writer(), 12, 13, 40

convert_binary_to_hmp, 11
convert_hmp_to_binary, 12
convert_hmp_to_binary(), 12

get_bin_info_actions, 13
get_bin_info_states, 16
get_hypergraph, 17
get_hypergraph(), 44, 48, 49
get_info, 21
get_info(), 17, 70
get_policy, 25
get_policy(), 59, 60
get_rpo, 30
get_steady_state_pr, 31
get_w_idx, 32

hmp_mdp_writer, 32
hmp_mdp_writer(), 11

load_mdp, 36
load_mdp(), 17, 21, 26, 31, 32, 54, 58–60, 65, 66, 70

memory_mdp_writer, 40
memory_mdp_writer(), 66, 70

plot.HMDP, 43
plot.HMDP(), 18, 49
plot_hypergraph, 47
plot_hypergraph(), 17, 18, 44

random_hmdp, 52
run_calc_weights, 53
run_policy_ite_ave, 58
run_policy_ite_discount, 59
run_value_ite, 60
run_value_ite(), 26

save_mdp, 65
save_mdp(), 65, 70
set_action_weight, 65
set_policy, 66
set_transition_weight, 70