Proactive Planning and Execution Strategies with Multiple Hypotheses

Jorrit T'Hooft

Context

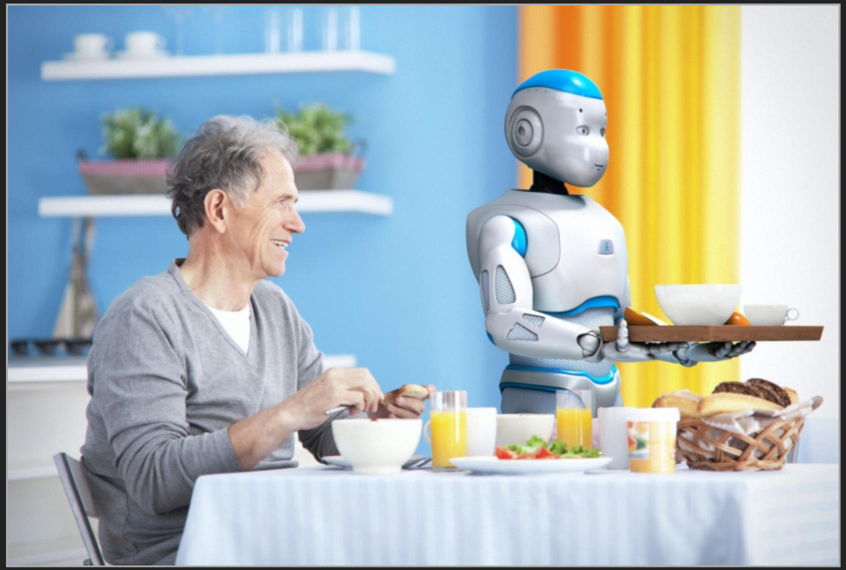

Context

Complex Missions

- Human Assistance
- Search & Rescue
- Surveillance
- Dangerous Work

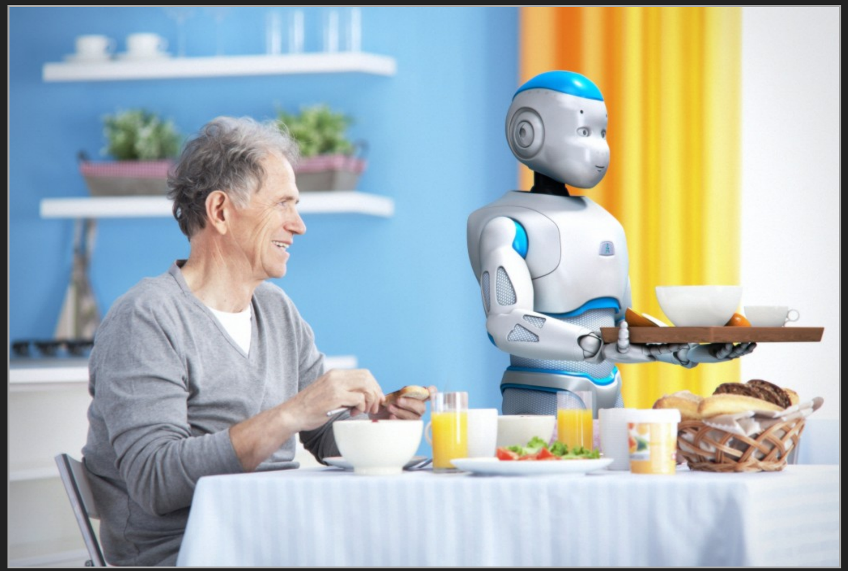

Context

Complex Missions

- Human Assistance
- Search & Rescue
- Surveillance
- Dangerous Work

Complex Environments

- Dynamic
- Unstructured
- Open
- Partially Observable

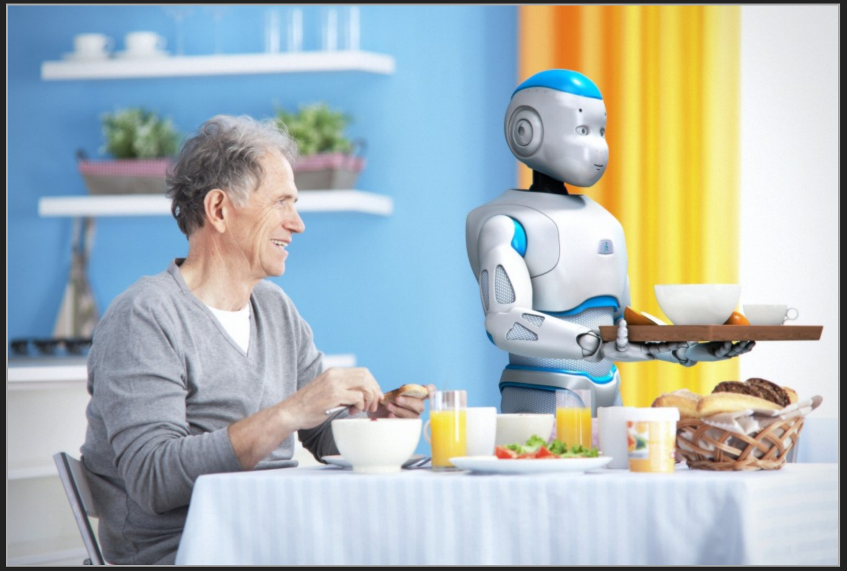

Context

Complex Missions

- Human Assistance
- Search & Rescue
- Surveillance
- Dangerous Work

Complex Environments

- Dynamic
- Unstructured
- Open
- Partially Observable

- ▸ Impossible to precompute a plan for each possible situation.
- ▸ Service robots need to plan and manage their actions online!

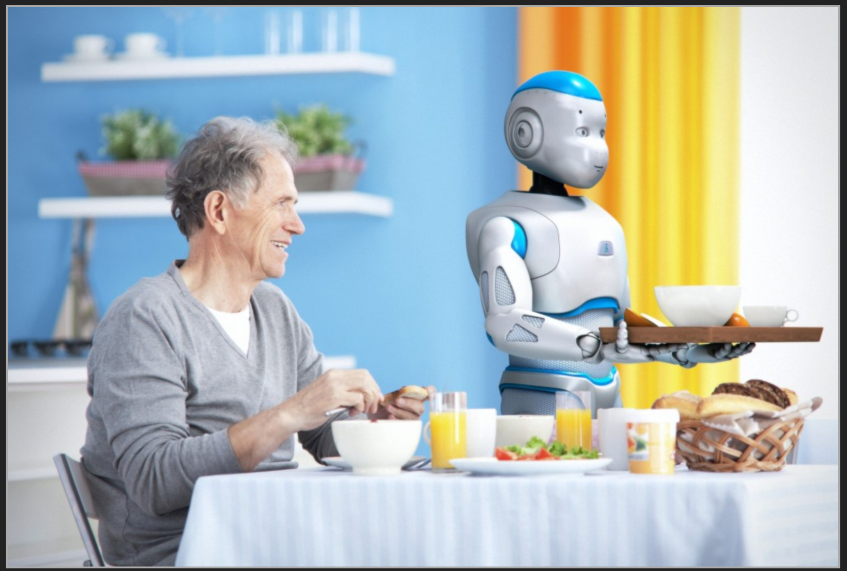

Online Planning & Execution

Triggers replanning to adapt the plan during execution.

1. M. Ghallab, D. Nau, and P. Traverso, “The actor’s view of automated planning and acting: A position paper”, Artificial Intelligence, vol. 208, pp. 1–17, 2014.
2. S. Arora, S. Choudhury, D. Althoff, and S. Scherer, “Emergency maneuver library-ensuring safe navigation in partially known environments”, in ICRA, 2015.
3. D. Kortenkamp and R. Simmons, “Robotic systems architectures and programming”, in Springer Handbook of Robotics, 2008.

Online Planning & Execution

Triggers replanning to adapt the plan during execution.

Planning & Execution Strategy

The way of interleaving the planning, selection and execution of actions.

- When to call which deliberative function, with which parameters.[1]

1. M. Ghallab, D. Nau, and P. Traverso, “The actor’s view of automated planning and acting: A position paper”, Artificial Intelligence, vol. 208, pp. 1–17, 2014.
2. S. Arora, S. Choudhury, D. Althoff, and S. Scherer, “Emergency maneuver library-ensuring safe navigation in partially known environments”, in ICRA, 2015.
3. D. Kortenkamp and R. Simmons, “Robotic systems architectures and programming”, in Springer Handbook of Robotics, 2008.

Online Planning & Execution

Triggers replanning to adapt the plan during execution.

Planning & Execution Strategy

The way of interleaving the planning, selection and execution of actions.

- When to call which deliberative function, with which parameters.[1]

Default Behavior

To ensure the integrity of the system and/or its environment when no valid plan is available.

- Such default behavior can be precomputed or obtained by fast (sub-optimal) planning.[2]

1. M. Ghallab, D. Nau, and P. Traverso, “The actor’s view of automated planning and acting: A position paper”, Artificial Intelligence, vol. 208, pp. 1–17, 2014.
2. S. Arora, S. Choudhury, D. Althoff, and S. Scherer, “Emergency maneuver library-ensuring safe navigation in partially known environments”, in ICRA, 2015.
3. D. Kortenkamp and R. Simmons, “Robotic systems architectures and programming”, in Springer Handbook of Robotics, 2008.

Online Planning & Execution

Triggers replanning to adapt the plan during execution.

Planning & Execution Strategy

The way of interleaving the planning, selection and execution of actions.

- When to call which deliberative function, with which parameters.[1]

Default Behavior

To ensure the integrity of the system and/or its environment when no valid plan is available.

- Such default behavior can be precomputed or obtained by fast (sub-optimal) planning.[2]

Reactivity

The robot must be able to adapt its behavior fast enough when unexpected events occur.

- Constraints on the software architecture.[3]

---

1. M. Ghallab, D. Nau, and P. Traverso, "The actor's view of automated planning and acting: A position paper", Artificial Intelligence, vol. 208, pp. 1–17, 2014.
2. S. Arora, S. Choudhury, D. Althoff, and S. Scherer, "Emergency maneuver library-ensuring safe navigation in partially known environments", in ICRA, 2015.
3. D. Kortenkamp and R. Simmons, "Robotic systems architectures and programming", in Springer Handbook of Robotics, 2008.

Two Existing Strategies

1. E. Marder-Eppstein, E. Berger, T. Foote, B. P. Gerkey, and K. Konolige, "The Office Marathon: Robust navigation in an indoor office environment", in ICRA 2010.
2. S. W. Yoon, A. Fern, and R. Givan, "FF-Replan: A baseline for probabilistic planning", in ICAPS 2007.
3. J. Van Den Berg, D. Ferguson, and J. Kuffner, "Anytime path planning and replanning in dynamic environments", in ICRA 2006.
4. B. Marthi, "Robust navigation execution by planning in belief space", in RSS 2012.

Two Existing Strategies

Plan-Replan[1,2]

Only replans when strictly necessary.

- Triggers replanning only when no valid plan is available (only when default behavior is executed).

1. E. Marder-Eppstein, E. Berger, T. Foote, B. P. Gerkey, and K. Konolige, “The Office Marathon: Robust navigation in an indoor office environment”, in ICRA 2010.
2. S. W. Yoon, A. Fern, and R. Givan, “FF-Replan: A baseline for probabilistic planning”, in ICAPS 2007.
3. J. Van Den Berg, D. Ferguson, and J. Kuffner, “Anytime path planning and replanning in dynamic environments”, in ICRA 2006.
4. B. Marthi, “Robust navigation execution by planning in belief space”, in RSS 2012.

Two Existing Strategies

Plan-Replan[1,2]

Only replans when strictly necessary.

- Triggers replanning only when no valid plan is available
  (only when default behavior is executed).

Continuous Planning[3,4]

Tries to avoid the situation where no valid plan is available, by integrating observed changes.

- Triggers replanning continuously during execution,
  with the freshest information available,
  to continuously update the plan.

1. E. Marder-Eppstein, E. Berger, T. Foote, B. P. Gerkey, and K. Konolige, “The Office Marathon: Robust navigation in an indoor office environment”, in ICRA 2010.
2. S. W. Yoon, A. Fern, and R. Givan, “FF-Replan: A baseline for probabilistic planning”, in ICAPS 2007.
3. J. Van Den Berg, D. Ferguson, and J. Kuffner, “Anytime path planning and replanning in dynamic environments”, in ICRA 2006.
4. B. Marthi, “Robust navigation execution by planning in belief space”, in RSS 2012.

Continuous Proactive Planning with Multiple Hypotheses

Extends *Continuous Planning* with the following concepts:

Continuous Proactive Planning with Multiple Hypotheses

Proactive Planning[1]

Anticipating situations by searching for solutions in a proactive manner.

Generating multiple solution-plans by proactive planning for multiple hypotheses (which can correspond to predictable future situations).

1. E. Burns, J. Benton, W. Ruml, S. W. Yoon, and M. B. Do, "Anticipatory on-line planning," in ICAPS 2012.

Continuous Proactive Planning with Multiple Hypotheses

Proactive Planning[1]

Anticipating situations by searching for solutions in a proactive manner.

Generating multiple solution-plans by proactive planning for multiple hypotheses (which can correspond to predictable future situations).

Hypotheses

A collection of information in order to generate a corresponding type of solution-plan.

Such as:

- Constraints on the search space.
- A particular configuration or initialization ...

1. E. Burns, J. Benton, W. Ruml, S. W. Yoon, and M. B. Do, “Anticipatory on-line planning,” in ICAPS 2012.

Continuous Proactive Planning with Multiple Hypotheses

Proactive Planning[1]

Anticipating situations by searching for solutions in a proactive manner.

Generating multiple solution-plans by proactive planning for multiple hypotheses (which can correspond to predictable future situations).

Hypotheses

A collection of information in order to generate a corresponding type of solution-plan.

Such as:

- Constraints on the search space.
- A particular configuration or initialization ...

Using One or Multiple Planners

Select the appropriate planner(s) to solve for each hypothesis.

A planner can be more suited than another:

- Ability to deal with uncertainties.
- Solving duration.
- Optimality guarentees ...

Solving a hypothesis with different planners to obtain different plans for the same constraints.

1. E. Burns, J. Benton, W. Ruml, S. W. Yoon, and M. B. Do, "Anticipatory on-line planning," in ICAPS 2012.

Continuous Proactive Planning with Multiple Hypotheses

Proactive Planning[1]

Anticipating situations by searching for solutions in a proactive manner.

Generating multiple solution-plans by proactive planning for multiple hypotheses (which can correspond to predictable future situations).

Hypotheses

A collection of information in order to generate a corresponding type of solution-plan.

Such as:

- Constraints on the search space.
- A particular configuration or initialization ...

Using One or Multiple Planners

Select the appropriate planner(s) to solve for each hypothesis.

A planner can be more suited than another:

- Ability to deal with uncertainties.
- Solving duration.
- Optimality guarentees ...

Solving a hypothesis with different planners to obtain different plans for the same constraints.

Tackling (some) Uncertainties with Multiple Hypotheses

Generate hypotheses on-the-fly taking into account only appropriate uncertainties.

To speed up the search:
take into account only a subset of the uncertainties.

To discard the need of modeling the uncertainties:
solve for possible future situations separately.

1. E. Burns, J. Benton, W. Ruml, S. W. Yoon, and M. B. Do, “Anticipatory on-line planning,” in ICAPS 2012.

Continuous Proactive Planning with Multiple Hypotheses

High-level algorithm:

loop

Continuous Proactive Planning with Multiple Hypotheses

High-level algorithm:

update context

loop

update current goal, current state, planning results (if any)

Continuous Proactive Planning with Multiple Hypotheses

High-level algorithm:

	loop
update context	update current goal, current state, planning results (if any)

manage actions
execute most appropriate action
	if no action executing **then**
	select most appropriate action to execute; launch selected action
	else
	verify if the executing action is still the most appropriate
	if a more appropriate action is available **then**
	stop the executing of current action; launch new most appropriate action;

Continuous Proactive Planning with Multiple Hypotheses

High-level algorithm:

	loop
update context	update current goal, current state, planning results (if any)

manage actions
execute most appropriate action
	if no action executing **then**
	select most appropriate action to execute; launch selected action
	else
	verify if the executing action is still the most appropriate
	if a more appropriate action is available **then**
	stop the executing of current action; launch new most appropriate action;

manage planning
plan for appropriate hypothesis with appropriate planner
	if currently planning **then**
	if the hypothesis for which we are planning is still appropriate **then**
	continue planning
	else
	stop the corresponding algorithm
	if currently not planning **then**
	if hypothesis left to explore **then**
	$h \leftarrow$ select hypothesis to plan for; select planning algorithm for h
	create planning algorithm input for h; launch the planner for h

Tests – Autonomous Navigation in Dynamic Environments

Environments from grid-based path planning benchmarks[1]

1. N. Sturtevant, “Benchmarks for Grid-Based Pathfinding,” Transactions on Computational Intelligence and AI in Games, 2012.

Tests – Autonomous Navigation in Dynamic Environments

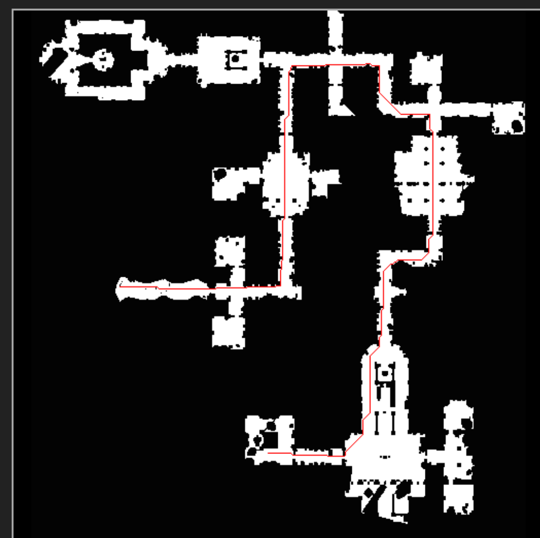
dao


wc3maps

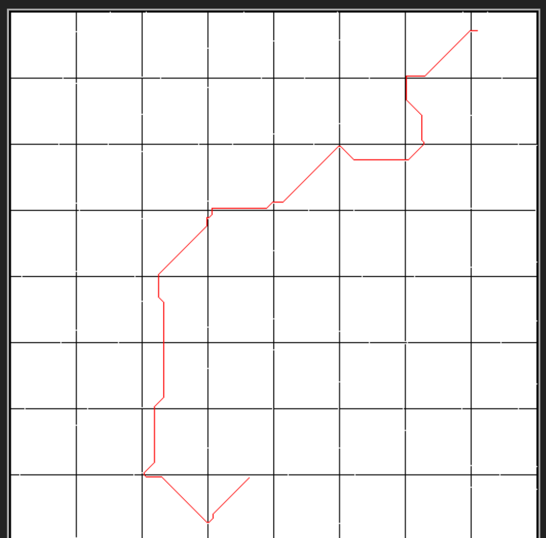
rooms

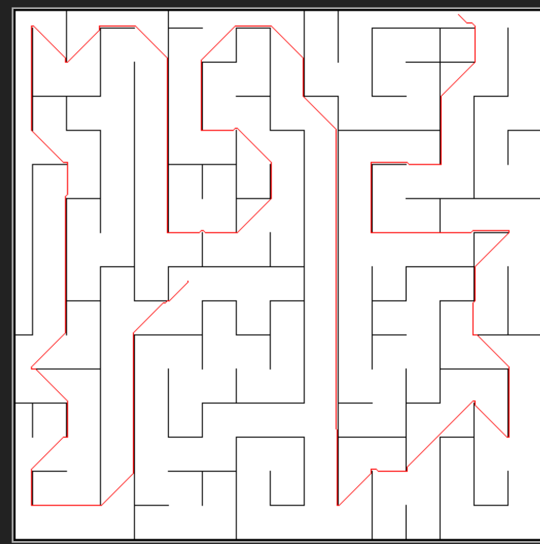
mazes

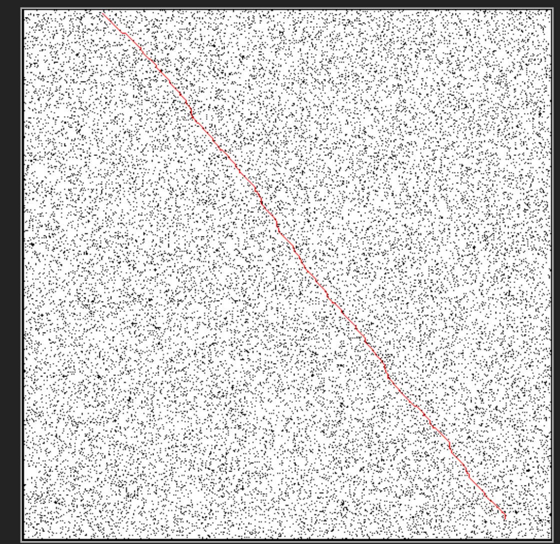
random (10% filled)

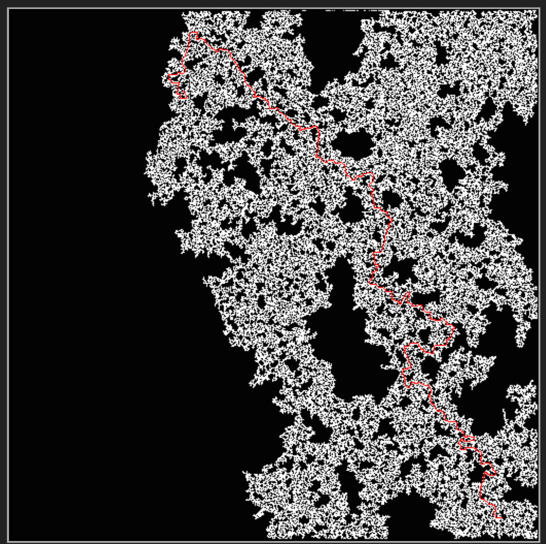
random (40% filled)

Environments from grid-based path planning benchmarks[1]

1. N. Sturtevant, "Benchmarks for Grid-Based Pathfinding," Transactions on Computational Intelligence and AI in Games, 2012.

Test Conditions

1. Peter E Hart, Nils J Nilsson, and Bertram Raphael, "A formal basis for the heuristic determination of minimum cost paths", Systems Science and Cybernetics 1968 .
2. S. Koenig and M. Likhachev, “Improved fast replanning for robot navigation in unknown terrain”, in ICRA 2002.

Test Conditions

Actions

8 "normal" actions
1 "default" action

1. Peter E Hart, Nils J Nilsson, and Bertram Raphael, "A formal basis for the heuristic determination of minimum cost paths", Systems Science and Cybernetics 1968 .
2. S. Koenig and M. Likhachev, “Improved fast replanning for robot navigation in unknown terrain”, in ICRA 2002.

Test Conditions

Actions

8 "normal" actions
1 "default" action

Plans

Path of "normal" actions

1. Peter E Hart, Nils J Nilsson, and Bertram Raphael, "A formal basis for the heuristic determination of minimum cost paths", Systems Science and Cybernetics 1968
.
2. S. Koenig and M. Likhachev, “Improved fast replanning for robot navigation in unknown terrain”, in ICRA 2002.

Test Conditions

Actions

8 "normal" actions
1 "default" action

Plans

Path of "normal" actions

2 Planners

A*[1]

D* Lite[2]

1. Peter E Hart, Nils J Nilsson, and Bertram Raphael, "A formal basis for the heuristic determination of minimum cost paths", Systems Science and Cybernetics 1968 .
2. S. Koenig and M. Likhachev, “Improved fast replanning for robot navigation in unknown terrain”, in ICRA 2002.

Test Conditions

Actions

8 "normal" actions

1 "default" action

Plans

Path of "normal" actions

2 Planners

A*[1]

D* Lite[2]

Dynamic obstacles

Invalidating the followed path with probability $P_{obstacle} \in \{0.2, 0.5, 0.8\}$

1. Peter E Hart, Nils J Nilsson, and Bertram Raphael, "A formal basis for the heuristic determination of minimum cost paths", Systems Science and Cybernetics 1968 .
2. S. Koenig and M. Likhachev, “Improved fast replanning for robot navigation in unknown terrain”, in ICRA 2002.

Test Conditions

Actions

8 "normal" actions
1 "default" action

Plans

Path of "normal" actions

2 Planners

A*[1]

D* Lite[2]

Dynamic obstacles

Invalidating the followed path with probability $P_{obstacle} \in \{0.2, 0.5, 0.8\}$

Minimum duration of the "default" actions: 2 cases

A: No minimum duration ⇒ The duration of a default action equals the (re)planning duration
B: Minimum duration of 0.5s

1. Peter E Hart, Nils J Nilsson, and Bertram Raphael, "A formal basis for the heuristic determination of minimum cost paths", Systems Science and Cybernetics 1968 .
2. S. Koenig and M. Likhachev, “Improved fast replanning for robot navigation in unknown terrain”, in ICRA 2002.

Types of Hypotheses

Types of Hypotheses

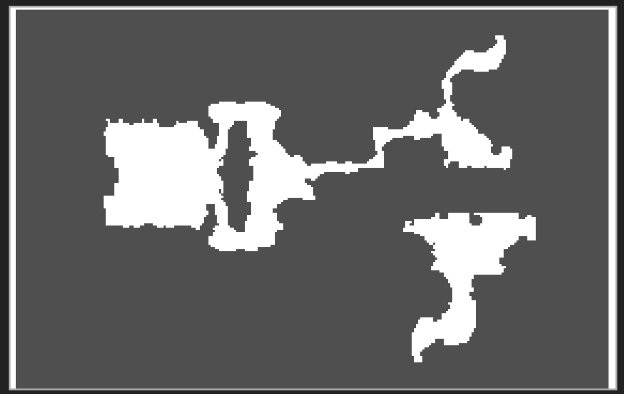

Global

Takes into account the complete search space

Types of Hypotheses

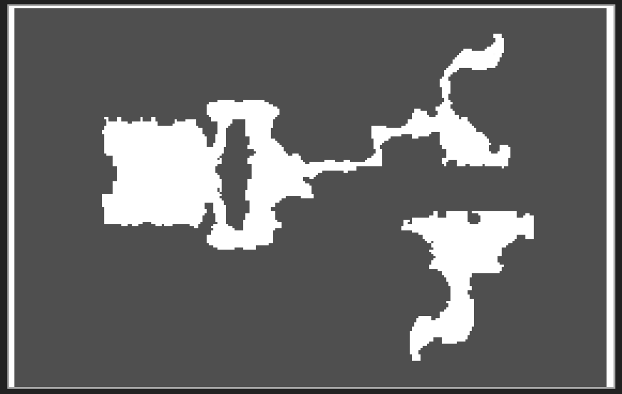

Global

Takes into account the complete search space

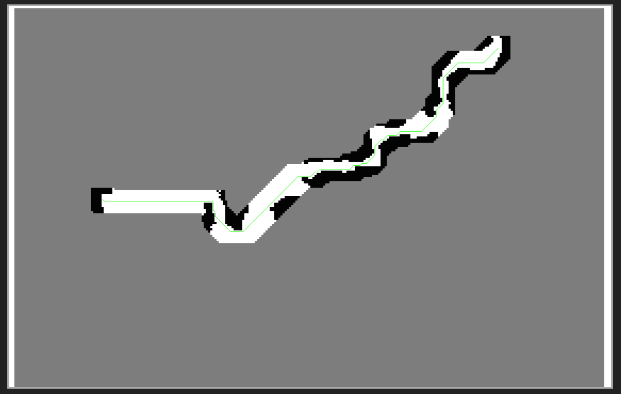

Path Region

Limits the search space to a region around a previous path

Types of Hypotheses

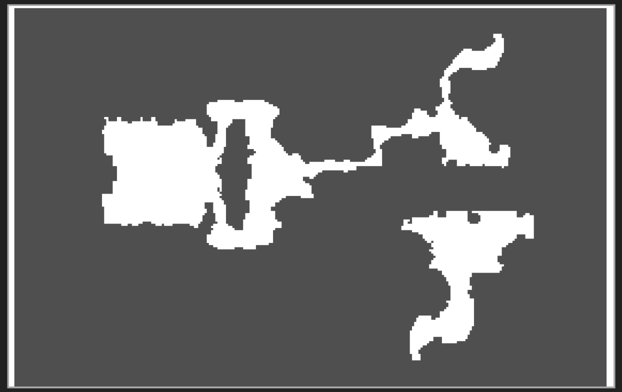

Global

Takes into account the complete search space

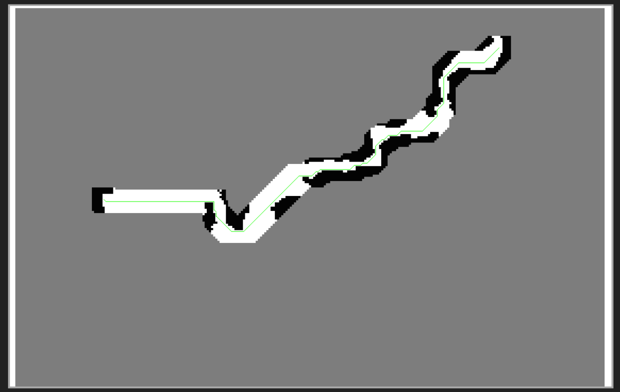

Path Region

Limits the search space to a region around a previous path

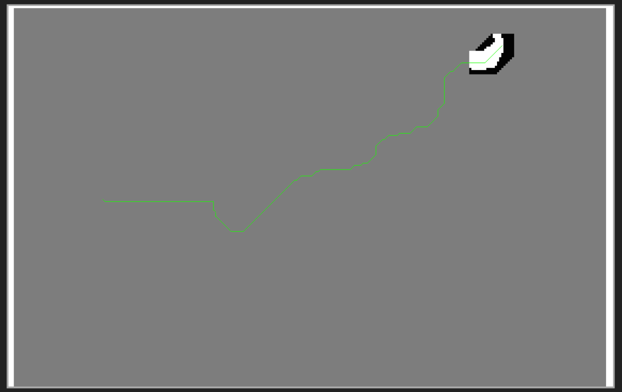

Sub-Path Region

Applies the *path region* hypothesis to the first part of a previous path.

Combines the new solution for the first part, with the previous path for the remaining part.

Tested Instantiations

Tested Instantiations

Plan-Replan

- **PR-A**
  - 1 global hypothesis / A* planner
- **PR-D**
  - 1 global hypothesis / D* Lite planner

Tested Instantiations

Plan-Replan

- **PR-A**
  - 1 global hypothesis / A* planner
- **PR-D**
  - 1 global hypothesis / D* Lite planner

Continuous Planning

- **CP-D**
  - 1 global hypothesis / D* Lite planner

Tested Instantiations

Plan-Replan

PR-A

1 global hypothesis / A* planner

PR-D

1 global hypothesis / D* Lite planner

Continuous Planning

CP-D

1 global hypothesis / D* Lite planner

Continuous Proactive Planning with Multiple (10) Hypotheses

CPP-1

1 global hypothesis / D* Lite planner
9 sub-path region hypotheses with different sub-goals / A* planner

CPP-2

1 global hypothesis / D* Lite planner
9 sub-path region hypotheses with different sub-goals + obstacle prediction / A* planner

CPP-3

1 global hypothesis / D* Lite planner
9 global hypotheses with obstacle prediction / D* Lite planner

Test Results

Results for dao (similar for the first five environments)

	$P_{obstacle} = 0.2$	$P_{obstacle} = 0.5$	$P_{obstacle} = 0.8$
A			
B			

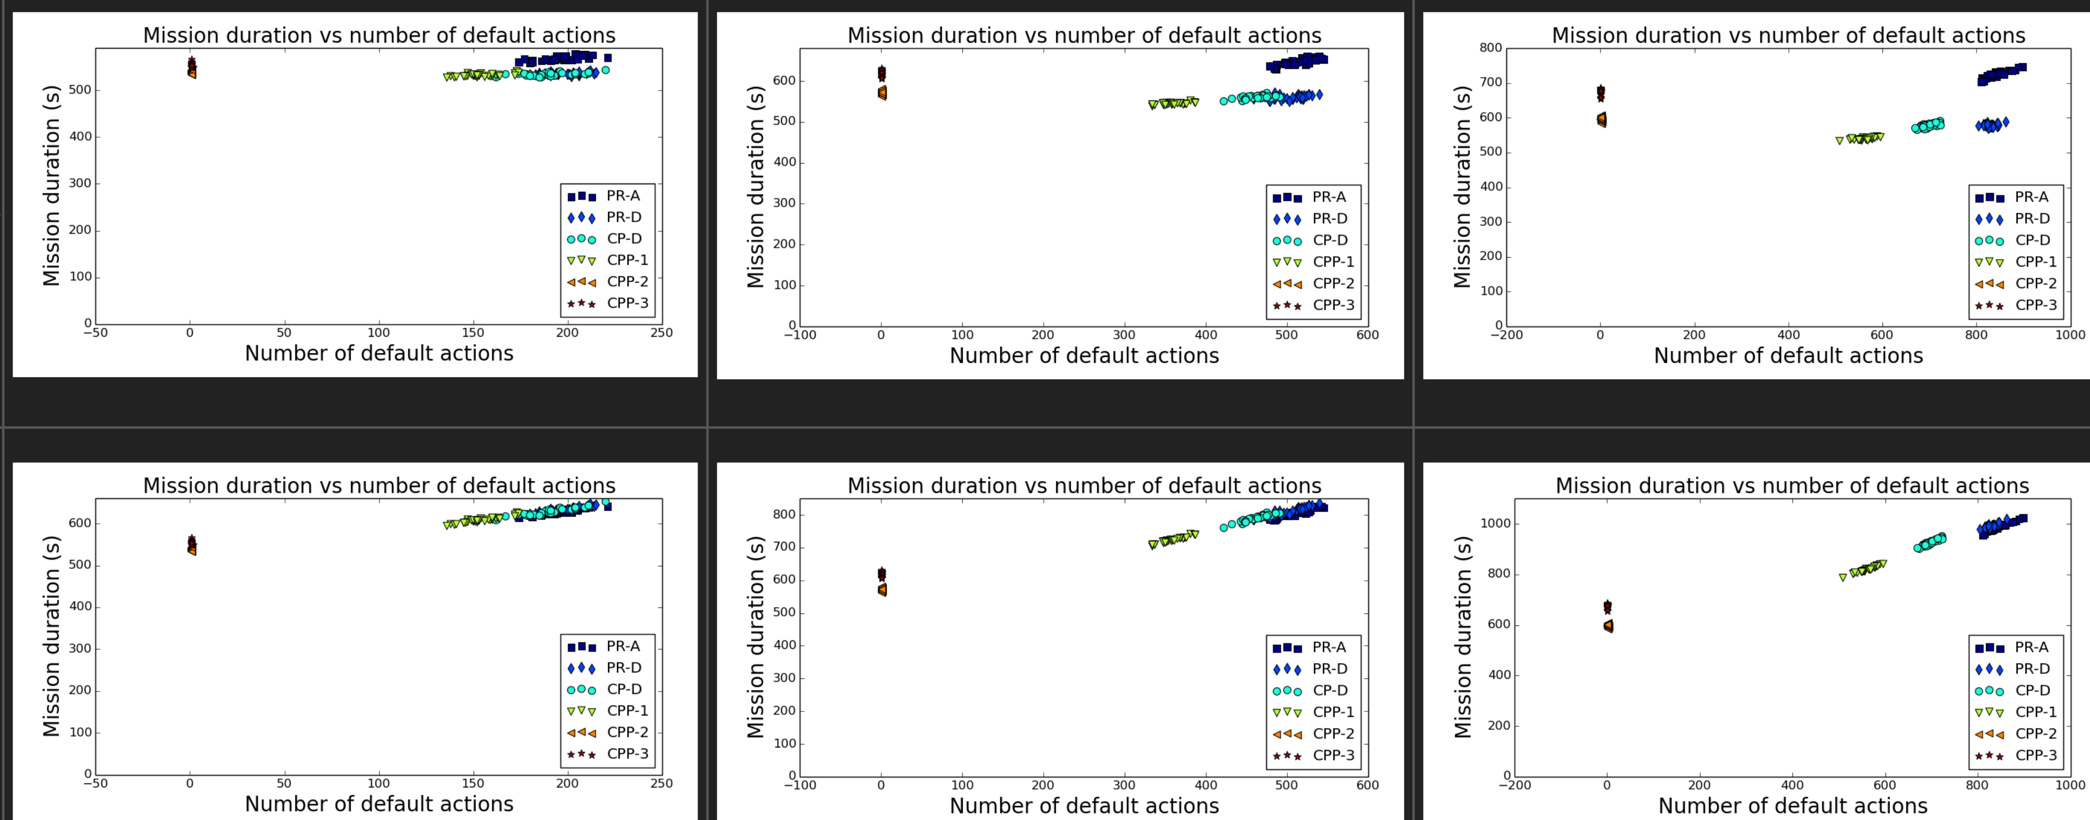

Dynamic obstacles

Invalidating the followed path with probability $P_{obstacle} \in \{0.2, 0.5, 0.8\}$

Minimum duration of the "default" actions: 2 cases

A: No minimum duration $\Rightarrow$ The duration of a default action equals the (re)planning duration

B: Minimum duration of 0.5s

Test Results

Results for dao (similar for the first five environments)

	$P_{obstacle} = 0.5$	$P_{obstacle} = 0.8$
A	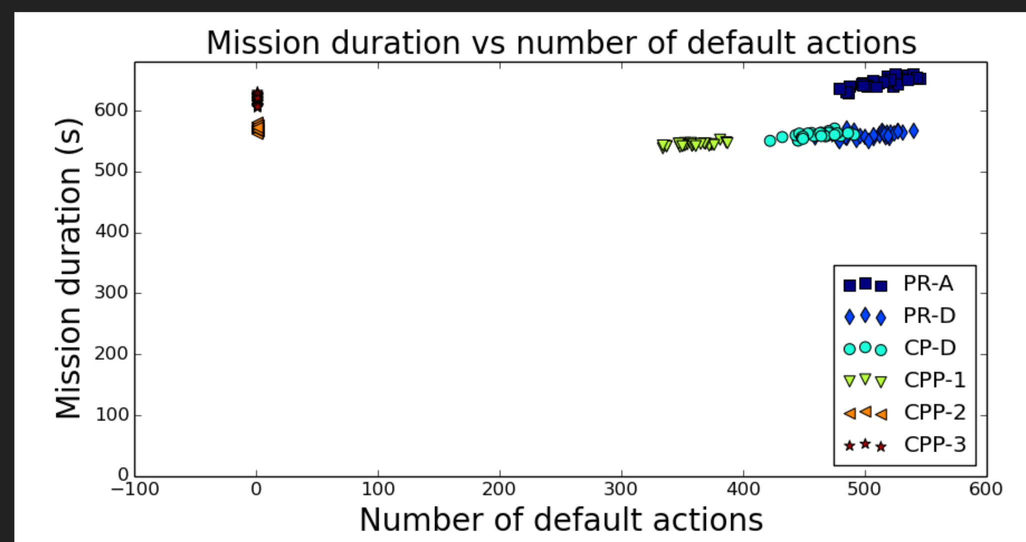	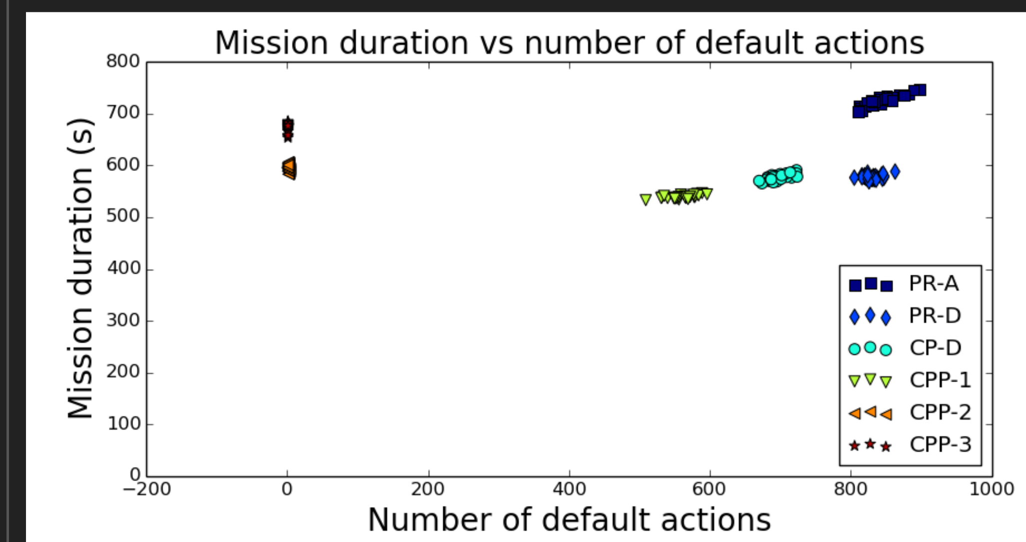
B	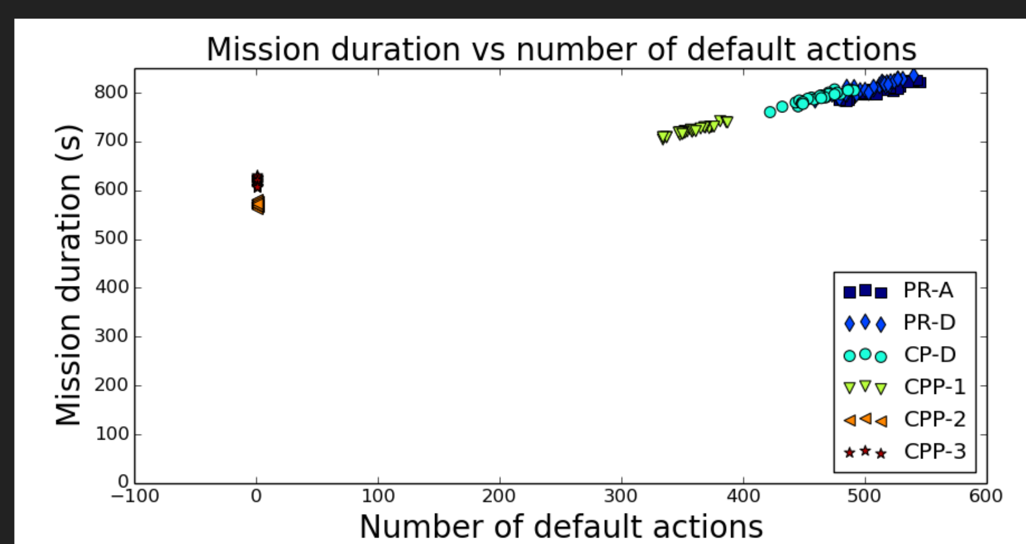	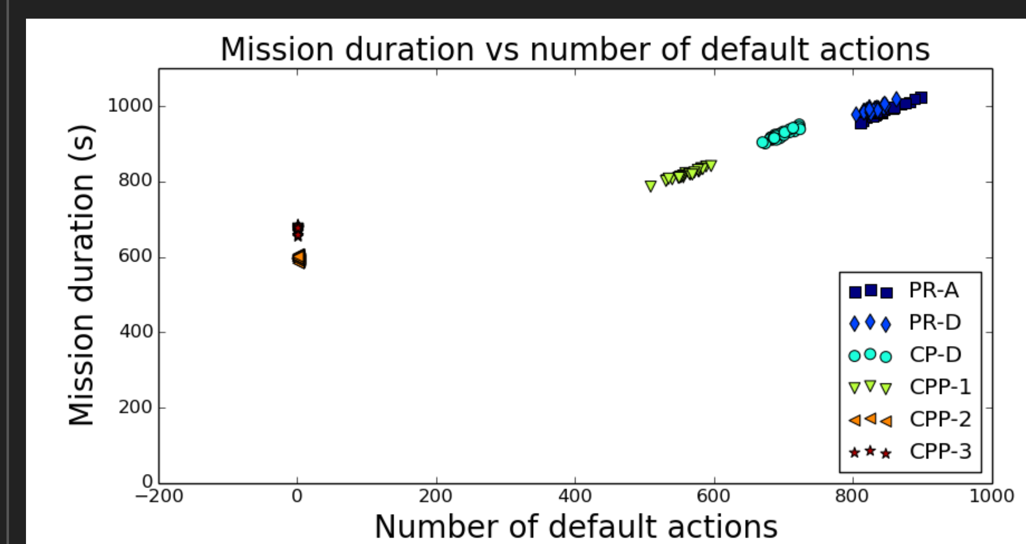

Test Results

Results for random (40% filled) (sixth environment)

$P_{obstacle} = 0.2$ | $P_{obstacle} = 0.5$ | $P_{obstacle} = 0.8$

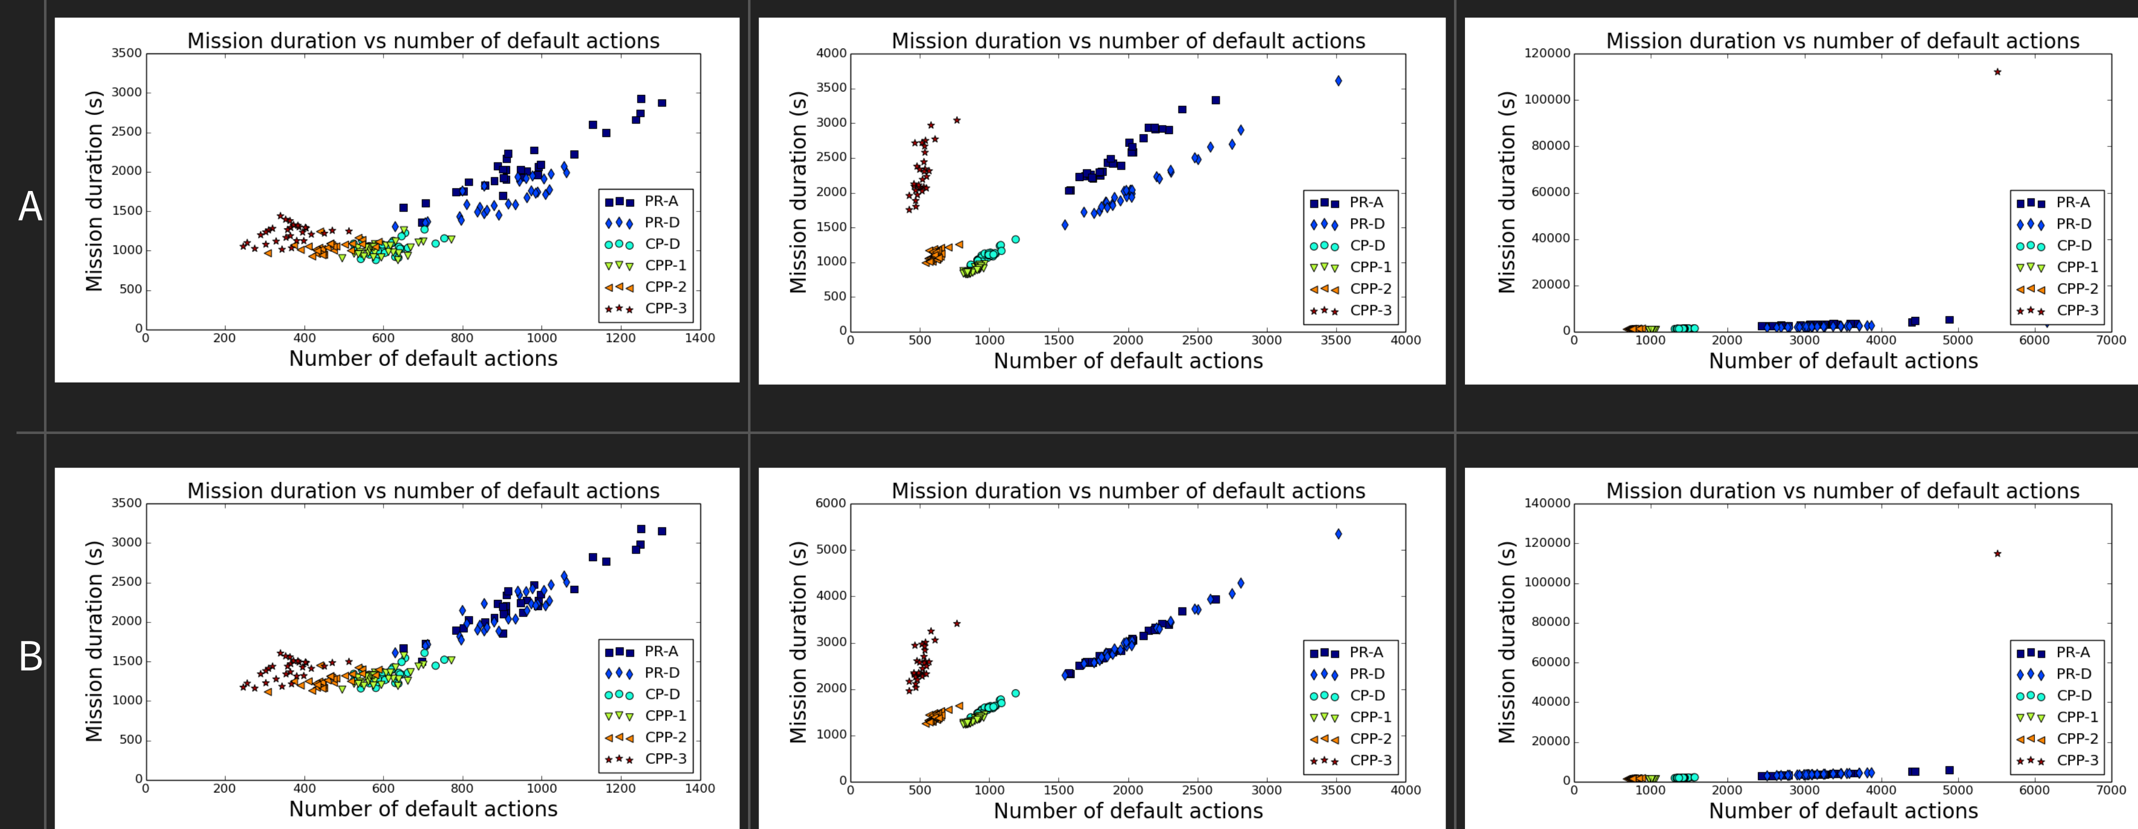

Dynamic obstacles

Invalidating the followed path with probability $P_{obstacle} \in \{0.2, 0.5, 0.8\}$

Minimum duration of the "default" actions: 2 cases

A: No minimum duration $\Rightarrow$ The duration of a default action equals the (re)planning duration

B: Minimum duration of 0.5s

Test Results

Results for random (40% filled) (sixth environment)

	$P_{obstacle} = 0.5$	$P_{obstacle} = 0.8$
A	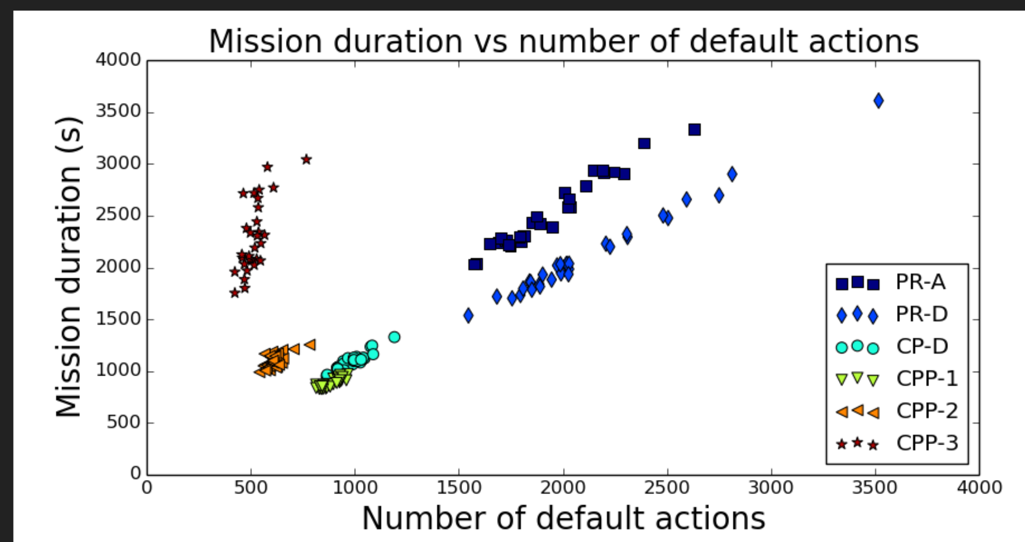	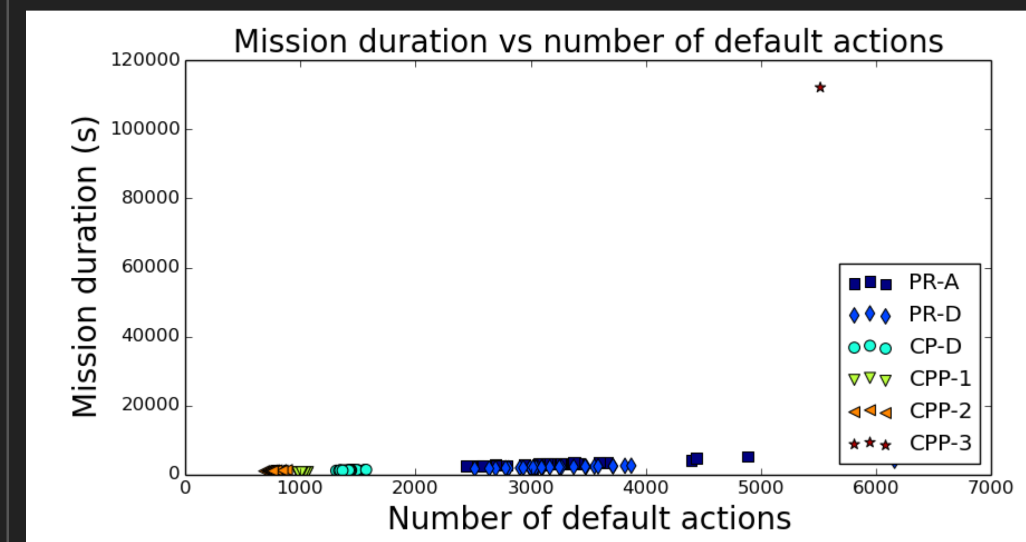
B	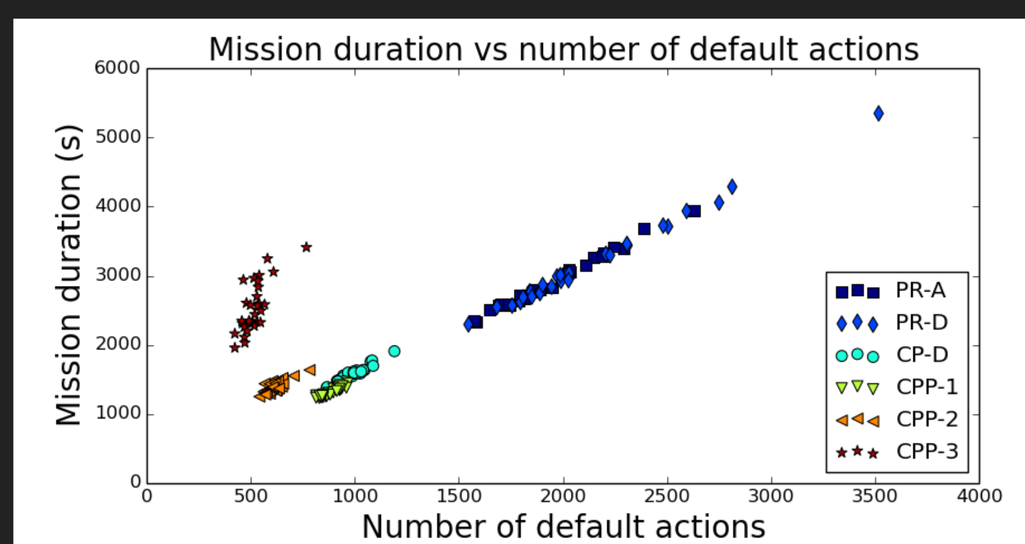	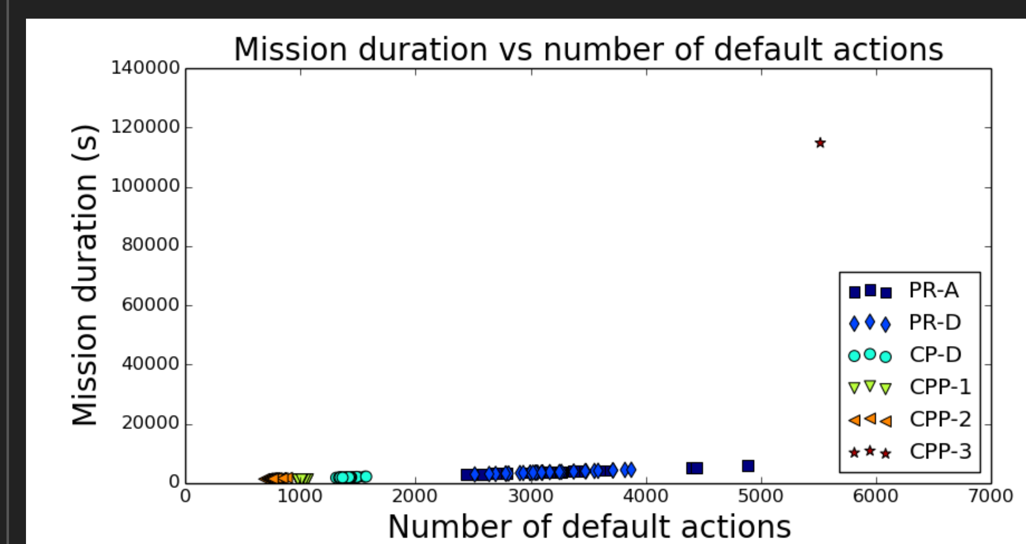

Reminding the Concepts Used for the Tests

Reminding the Concepts Used for the Tests

Proactive planning

9 additional hypotheses for generating multiple solution-plans

Reminding the Concepts Used for the Tests

Proactive planning

- 9 additional hypotheses for generating multiple solution-plans

Hypotheses

- Global Hypothesis (+ obstacle prediction)
- Sub-path region hypotheses: selecting sub-goals (+ obstacle prediction)

Reminding the Concepts Used for the Tests

Proactive planning

9 additional hypotheses for generating multiple solution-plans

Hypotheses

Global Hypothesis (+ obstacle prediction)
Sub-path region hypotheses: selecting sub-goals (+ obstacle prediction)

Using multiple planners

The most appropriate:
D* Lite for adapting the previous plan (the goal remains the same)
A* for the sub-path region hypotheses (the goal changes)

Reminding the Concepts Used for the Tests

Proactive planning

9 additional hypotheses for generating multiple solution-plans

Hypotheses

Global Hypothesis (+ obstacle prediction)

Sub-path region hypotheses: selecting sub-goals (+ obstacle prediction)

Using multiple planners

The most appropriate:

D* Lite for adapting the previous plan (the goal remains the same)

A* for the sub-path region hypotheses (the goal changes)

Tackling uncertainties with multiple hypotheses

Planning for some possible futures separately

Conclusion from the Tests

Conclusion from the Tests

Proactive planning with multiple hypotheses may be able to improve the performances compared to commonly used strategies.

(CPP-1 & CPP-2 for the tests)

Conclusion from the Tests

Proactive planning with multiple hypotheses may be able to improve the performances compared to commonly used strategies.

(CPP-1 & CPP-2 for the tests)

But the hypotheses to plan for and the actions to execute must be selected carefully, otherwise unexpected behavior may appear.

(CPP-3 for the tests)

Critical Aspects

**For the multiple hypotheses paradigm,
some aspects are critical:**

Critical Aspects

**For the multiple hypotheses paradigm,
some aspects are critical:**

The choice and construction of the hypotheses to plan for.

The construction of interesting hypotheses is very difficult without enough semantic information.

Critical Aspects

For the multiple hypotheses paradigm,
some aspects are critical:

The choice and construction of the hypotheses to plan for.

The construction of interesting hypotheses is very difficult without enough semantic information.

The planner selected to solve for a hypothesis must be efficient enough to return a solution before the hypothesis becomes obsolete.

Critical Aspects

For the multiple hypotheses paradigm,
some aspects are critical:

The choice and construction of the hypotheses to plan for.

The construction of interesting hypotheses is very difficult without enough semantic information.

The planner selected to solve for a hypothesis must be efficient enough to return a solution before the hypothesis becomes obsolete.

Selecting the appropriate action to execute must be done carefully.

Otherwise unwanted behavior may appear (getting stuck in a local optimum, doubling back to often, ...).

Critical Aspects

**For the multiple hypotheses paradigm,
some aspects are critical:**

The choice and construction of the hypotheses to plan for.

The construction of interesting hypotheses is very difficult without enough semantic information.

The planner selected to solve for a hypothesis must be efficient enough to return a solution before the hypothesis becomes obsolete.

Selecting the appropriate action to execute must be done carefully.

Otherwise unwanted behavior may appear (getting stuck in a local optimum, doubling back to often, ...).

- **The multiple hypotheses paradigm needs enough semantic information to be efficient.**

Semantic Information

Semantic Information

With enough semantic information available, a robot could largely enhance its capabilities by inferring automatically:

Semantic Information

With enough semantic information available, a robot could largely enhance its capabilities by inferring automatically:

more interesting hypotheses;

(e.g. taking into account that a person can let the robot pass, a chair can be moved, some places would better be avoided, ...)

Semantic Information

With enough semantic information available, a robot could largely enhance its capabilities by inferring automatically:

- **more interesting hypotheses;**
  (e.g. taking into account that a person can let the robot pass, a chair can be moved, some places would better be avoided, ...)
- **which hypothesis to instantiate in which context;**

Semantic Information

With enough semantic information available, a robot could largely enhance its capabilities by inferring automatically:

- **more interesting hypotheses;**
  (e.g. taking into account that a person can let the robot pass, a chair can be moved, some places would better be avoided, ...)
- **which hypothesis to instantiate in which context;**
- **which planner to use for solving a particular hypothesis;**

Semantic Information

With enough semantic information available, a robot could largely enhance its capabilities by inferring automatically:

- **more interesting hypotheses;**
  (e.g. taking into account that a person can let the robot pass, a chair can be moved, some places would better be avoided, ...)
- **which hypothesis to instantiate in which context;**
- **which planner to use for solving a particular hypothesis;**
- **how to create the right input for a planner and a hypothesis;**

Semantic Information

With enough semantic information available, a robot could largely enhance its capabilities by inferring automatically:

- **more interesting hypotheses;**
  (e.g. taking into account that a person can let the robot pass, a chair can be moved, some places would better be avoided, ...)
- **which hypothesis to instantiate in which context;**
- **which planner to use for solving a particular hypothesis;**
- **how to create the right input for a planner and a hypothesis;**
- **how to schedule the planning episodes;**
  (e.g. for hypotheses with different time horizons, priorities, ...)

Semantic Information

With enough semantic information available, a robot could largely enhance its capabilities by inferring automatically:

- **more interesting hypotheses;**
  (e.g. taking into account that a person can let the robot pass, a chair can be moved, some places would better be avoided, ...)
- **which hypothesis to instantiate in which context;**
- **which planner to use for solving a particular hypothesis;**
- **how to create the right input for a planner and a hypothesis;**
- **how to schedule the planning episodes;**
  (e.g. for hypotheses with different time horizons, priorities, ...)
- **and how to select the most appropriate action.**

Semantic Information

With enough semantic information available, a robot could largely enhance its capabilities by inferring automatically:

more interesting hypotheses;

(e.g. taking into account that a person can let the robot pass, a chair can be moved, some places would better be avoided, ...)

which hypothesis to instantiate in which context;

which planner to use for solving a particular hypothesis;

how to create the right input for a planner and a hypothesis;

how to schedule the planning episodes;

(e.g. for hypotheses with different time horizons, priorities, ...)

and how to select the most appropriate action.

- Therefore, we are investigating tools for acquiring, storing and managing semantic information, such as *RoboSherlock*[1], *KnowRob*[2] and *CRAM*[3].

1. Michael Beetz et al. “RoboSherlock : Unstructured Information Processing for Robot Perception”, in ICRA 2015
2. Moritz Tenorth et Michael Beetz. “Representations for robot knowledge in the KnowRob framework”, in Artificial Intelligence 2015
3. Michael Beetz, Lorenz Mösenlechner et Moritz Tenorth. “CRAM—A Cognitive Robot Abstract Machine for everyday manipulation in human environments”, in IROS 2010