Creating Planning Portfolios with Predictive Models

Defense

March 23, 2017

Isabel Cenamor

icenamor@inf.uc3m.es

Advisors: Tomás de la Rosa and Fernando Fernández

Departamento de Informática

uc3m | Universidad Carlos III de Madrid

Outline

Automated Planning

Given a planning task:

- A description of the **initial state**
- A description of the **goals**
- A description of a set of **actions**

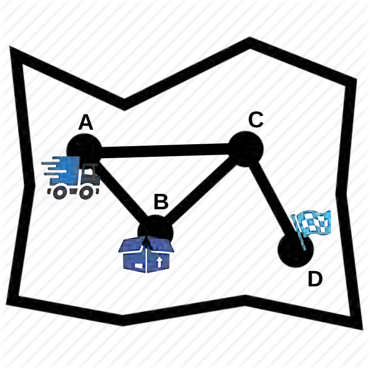

Find a sequence of actions (a **plan**) from the initial state to a final state in which the goal conditions fulfill

Planning Task → **Planner** → **Plan**

Introduction

- Planning Community organizes the International Planning Competition (IPC)
- Each IPC presents different **tracks**: optimal, temporal, satisficing...

Introduction

- Planning Community organizes the International Planning Competition (IPC)
- Each IPC presents different **tracks**: optimal, temporal, satisficing...
- IPC creates a perfect **framework** to fix the standard

Introduction

- Planning Community organizes the International Planning Competition (IPC)
- Each IPC presents different **tracks**: optimal, temporal, satisficing...
- IPC creates a perfect **framework** to fix the standard
- There is **no** single planner which is **always the best** planner for all planning tasks!

- Planning Community organizes the International Planning Competition (IPC)
- Each IPC presents different **tracks**: optimal, temporal, satisficing...
- IPC creates a perfect **framework** to fix the standard
- There is **no** single planner which is **always the best** planner for all planning tasks!
- A set of planners could be aggregated to create a **portfolio**

Portfolio Definition

Planning Portfolio

Given a set of base planners, $\{pl_1, \ldots, pl_n\}$, and a maximum execution time, T, a planning portfolio can be considered as a sequence of m pairs $< pl_1, t_1 >, \ldots, < pl_m, t_m >$, where $pl_i \in \{pl_1, \ldots, pl_n\}$ and $\sum_{j=1}^{m} t_j \leq T$.

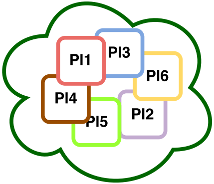

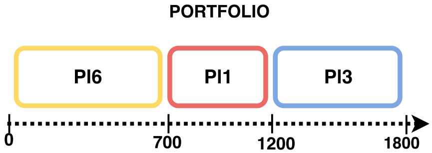

Portfolio
Challenges

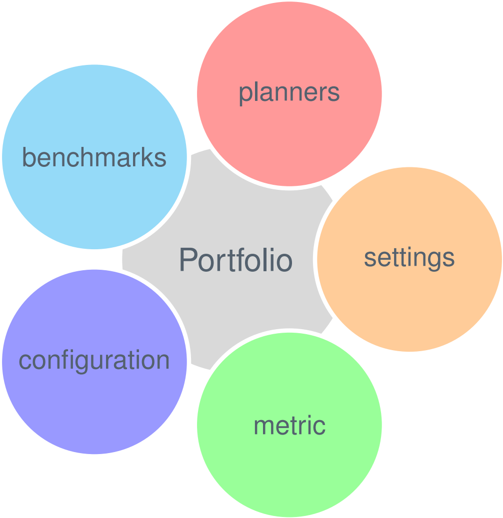

Choose the planning algorithms to consider for the portfolio

- Select and combine heuristics and search algorithms: FDSS [HRS+11], Cedalion [SSHH15], Uniform [SBGH12], ...
- Domain-optimized portfolio planners: PbP [GSV14], AGAP [VCK14]
- A group of independent planners: BUS [HDH+99], MIPlan [NBL15], ArvandHerd [VNM+14], ...

Configuration target: domain-independent (static), domain-specific, instance-specific

- Domain independent configuration (static): FDSS, MIPlan, Cedalion, Uniform, ArvandHerd, . . .
- Domain-specific configuration: PbP, AGAP
- Instance-specific configuration: BUS, AllPACA [MWK14]

Criteria of planner selection and execution order

- Maximizes the coverage: FDSS, Cedalion
- Knowledge with round-robin: PbP
- Predictive models: BUS, AllPACA
- Sorted planners in function of their contribution: MIPlan

Discussion

- Static Portfolio configurations are **suboptimal**
- Instance-specific configurations require an oracle
  - Given a problem $\rightarrow$ which is the **best** planner and how much **time** does it need
- Selected planners
  - **Many**
  - Low diversity
- Oracle
  - Predictive Models are **not perfect**
  - **Uncorrelated** shallow features
  - BUS portfolio

Objectives

1. Renew the idea of **dynamic portfolios** per instance
2. Find a diverse **subset of planners** with a multi-criteria approach
3. **Characterize** the planning task as a function of easily computable features
4. Model the **planner performance** with machine learning
5. Exploit the **predictive models** in a portfolio configuration
6. **Analyze** the features in **homogeneous problems** test sets
7. Extrapolate the general approach to **temporal planning**

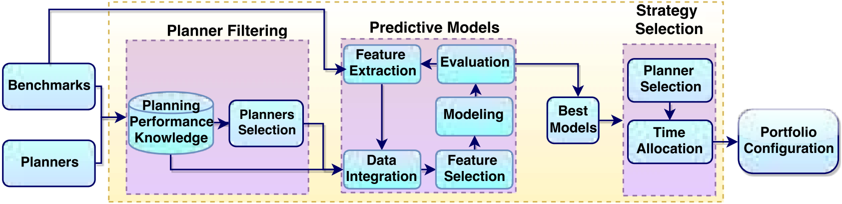

Benchmarks
Planners
Planner Filtering
Planning Performance Knowledge
Planners Selection
Predictive Models
Feature Extraction
Evaluation
Modeling
Data Integration
Feature Selection
Best Models
Strategy Selection
Planner Selection
Time Allocation
Portfolio Configuration

Planner Filtering

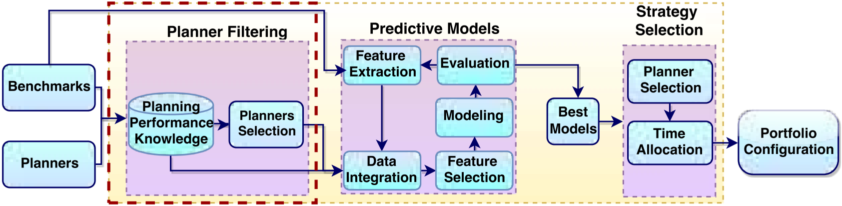

Initial Idea: follow IPC criteria

- Coverage
- Time
- Quality

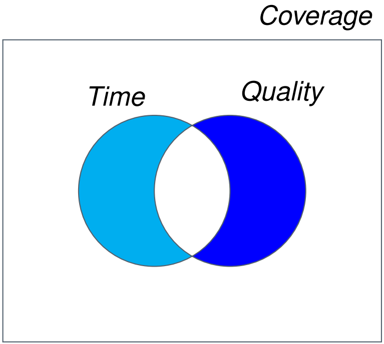

Time vs. Quality
Metric

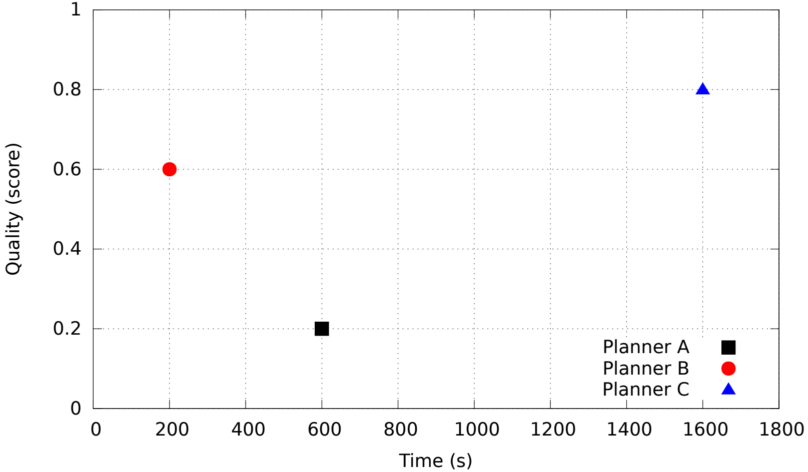

QT-Pareto Score Filtering

Our proposal

QT-Pareto dominance

A planner p_1 gets a tuple $\langle Q, T \rangle$ in a problem π , and a planner p_2, in the same problem, gets $\langle Q', T' \rangle$. The planner p_1 dominate p_2 if and only if $Q \geq Q'$ and $T < T'$.

QT-Pareto Score

Planner p gets $\frac{N}{N^*}$ points, where N is the number of tuples where p Pareto-dominates another planner, and N^* is the number of different tuples in which planner p appears.

QT-Pareto dominance

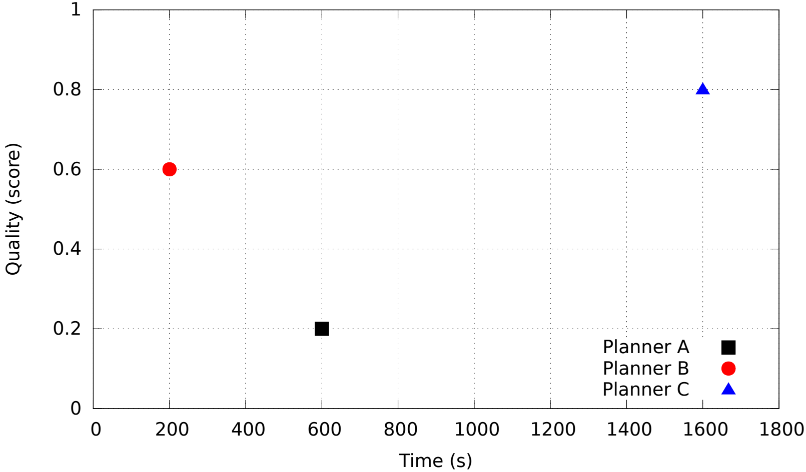

Metric Scope
Filtering Method

- Problem
- Domain
- IPC Ranking

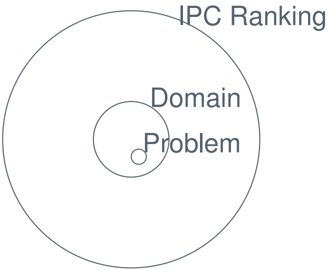

Planner Selection in Parcprinter domain

Best planners per problem in terms of **quality** score

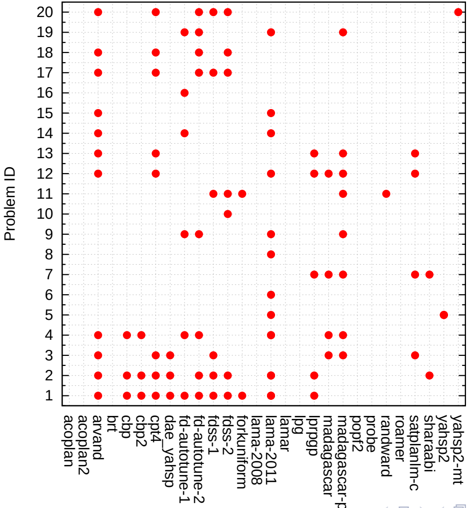

Planner Selection

Select a planner p as candidate when it gets the highest Score Filtering in a domain.

Experimental Setting
Evaluate planner selection

Training phase:

- Base planners: IPC-2011 and LPG-TD
- Benchmark domains: IPC-2011

Test phase:

- Time limit: 1800 seconds
- Memory limit: 4 GB RAM
- Benchmark domains: IPC-2014

Configurations:

- Portfolios: uniform time with arbitrary order
  1. ***QT***: portfolio using QT-Pareto
  2. ***Q***: portfolio using Quality
  3. ***T***: portfolio using Time
  4. ***C***: portfolio using number of solved problems (coverage)
  5. ***OET***: portfolio including 28 planners

Results of the Planner Filtering

Quality - Static Portfolio Configurations

Domains	QT	Q	T	C	OET
Hiking	19.14	**19.38**	18.56	19.12	18.17
Barman	**19.64**	17.65	19.14	19.38	16.74
Thoughtful	**19.54**	18.79	18.53	18.61	14.51
GED	19.17	18.52	**19.29**	19.08	18.28
Openstacks	19.66	**19.99**	19.50	14.88	15.44
Parking	18.99	**19.00**	16.99	9.72	17.64
Maintenance	15.53	**16.84**	13.89	16.46	15.00
Tetris	15.22	**15.89**	7.38	12.51	4.99
CityCar	**13.50**	12.69	7.82	8.68	5.99
Visitall	**16.90**	9.02	9.12	3.94	13.25
Childsnack	**18.73**	5.37	8.24	7.53	11.95
Transport	**19.95**	5.98	5.40	5.69	8.92
Floortile	**17.00**	3.43	1.88	3.43	4.81
CaveDiving	6.39	0.00	**7.00**	**7.00**	0.00
Total	**239.35**	182.56	172.73	166.03	165.68

Analysis of the Filtering Results
Ranking, planners selection and diversity

Ranking	Planner	QT	Q	T	C	FD
1	LAMA-2011	√	√	√	√	√
2	FDSS-1		√		√	√
3	FDSS-2	√		√	√	√
4	FD-AUTOTUNE-1	√	√		√	√
5	ROAMER		√		√	√
6	FORKUNIFORM			√	√	√
7	FD-AUTOTUNE-2	√	√		√	√
8	PROBE	√		√	√	
9	ARVAND	√	√		√	√
10	LAMA-2008	√			√	√
11	LAMAR	√	√	√	√	√
16	YAHSP2			√	√	
17	YAHSP2-MT	√		√	√	
20	MADAGASCAR-P			√	√	
22	MADAGASCAR	√				
24	LPG-TD	√	√	√	√	
Total	28	11	9	10	22	12

Predictive Models

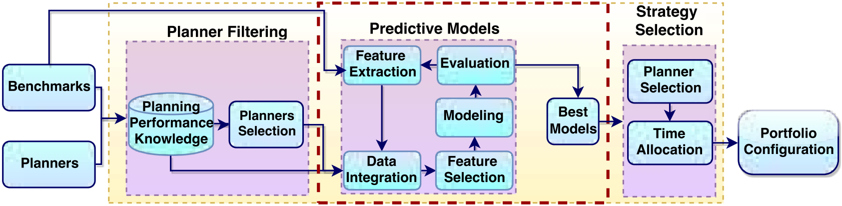

Feature Extraction I

Type, number, example

There are 114 features:

- **PDDL** (8): number of objects in the problem, ...
- **FD Instantiation** (16): number of generated rules in the translation process to SAS^+ task, ...
- **Heuristics** (16): FF heuristic in the initial state, ...
- **Landmark** (14): number of landmarks included in the merged landmark graph, ...
- **SAS^+** (50): number of variables of the CG, ...
- **Fact Balance** (10): number of times that a fact in the initial state is deleted in the computation of the relaxed plan, ...

Feature Extraction II

Summary of the extracted features

Process	Success	Average (s.)	Median (s.)
Tranlate (PDDL)	97%	5.98	0.36
Preprocess (FD & SAS$^+$)	97%	1.10	0.06
Fact Balance	93%	0.73	0.03
Heuristics	87.54%	13.15	0.68
Landmarks	87.54%	1.72	0.24
Mercury	97%	0.01	0.00
Extra time		0.44	0.22
Total		**23.11**	1.60

Total Instances

- **45** different **domain** descriptions: IPC-2006:2011 & Learning IPC-2008:2011
- Input: Features (problems and domains) + performance data (planner, solved, time)

Classification Task

- **Input**: Features + Planner
- **Output**: Solved / Unsolved task

Regression Task

- **Input**: Features + Planner
- **Output**: Time best solution

Results Modeling

Each training algorithm using 10-fold cross-validation technique

Classification

- **Accuracy**
- Standard Deviation

Regression

- **Relative Absolute Error**
- Standard Deviation

Rotation Forest
Accuracy = 90.50 %

Decision Table
Relative Absolute Error = 64.13%

Configuration Strategies

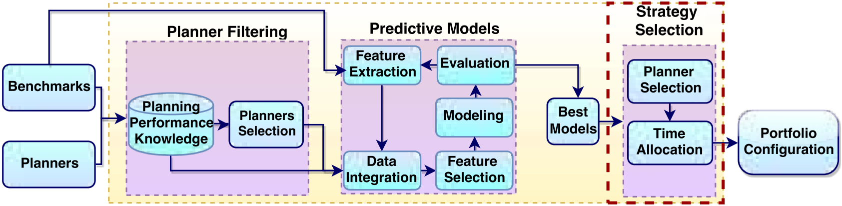

Strategy Selection

Problems with predictive models

How to transform the predictions of the best models into an actual portfolio configuration.

Include the previous knowledge in different strategies:

- Not using any predictive model
- Using **classification** model
- Using **classification** and **regression** models

But, there are two problems...

- ✗ If all planners get a **positive** prediction
- ✗ If all planners get a **negative** prediction

Solution: to use the **confidence** to predict the positive class

Estimated Number of Planners

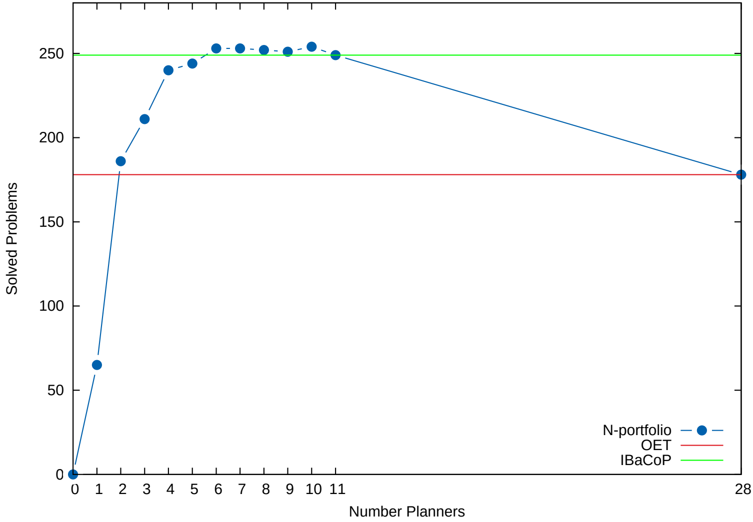

Strategy Selection

Our approximation for the IPC

- IBaCoP: QT-Pareto Score Filtering with uniform time
- IBaCoP2: Best N confidence strategy where N=5
- IBaCoP2-B5E: Estimated time to the previous selected planners

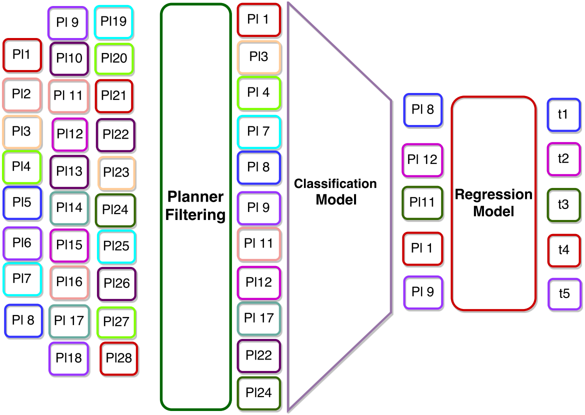

Experimental Setting

- Time limit: 1800 seconds
- Time limit for **feature extraction**: 300 seconds
- Memory limit: 4GB RAM
- Test benchmark domains: IPC-2014

Two baseline portfolios:

- Random 5 Planners (Rand): Run for 5 times from IBaCoP
- Best 5 Planners (Def): LAMA-2011, PROBE, FD-AUTOTUNE-1, LAMA-2008 and FD-AUTOTUNE-2

Two planners:

- Mercury: Second planner in terms of quality
- Jasper: Second planner in terms of coverage

Domains	Mercury	Jasper	Def	Rand	IBaCoP	IBaCoP2	B5E
Hiking	18.96	18.17	18.78	18.07	19.25	**19.63**	19.63
Openstacks	**19.64**	18.76	19.25	17.23	17.35	17.38	17.37
Thoughtful	12.73	16.37	19.15	17.60	**19.17**	18.15	18.23
GED	**19.46**	17.95	16.40	14.22	17.31	17.70	17.70
Parking	18.14	17.22	**18.18**	12.47	17.89	18.16	18.17
Barman	14.61	**18.97**	17.17	14.10	16.79	16.85	16.87
Maintenance	5.72	10.79	12.52	15.27	**16.45**	16.21	16.25
Tetris	**16.37**	16.14	9.37	11.49	13.60	15.69	13.55
Childsnack	0.00	0.00	2.67	10.16	**19.50**	19.23	19.36
CityCar	4.10	11.03	4.96	9.77	11.43	**14.36**	12.57
Visitall	**20.00**	15.36	13.68	12.72	15.24	9.94	8.01
Transport	**19.87**	12.02	6.90	8.51	10.25	11.53	11.13
CaveDiving	7.00	**8.00**	7.00	7.00	6.30	7.00	7.00
Floortile	2.00	2.00	4.14	9.39	16.22	15.28	**17.46**
total	178.59	182.78	170.16	177.99	216.75	**217.11**	213.31

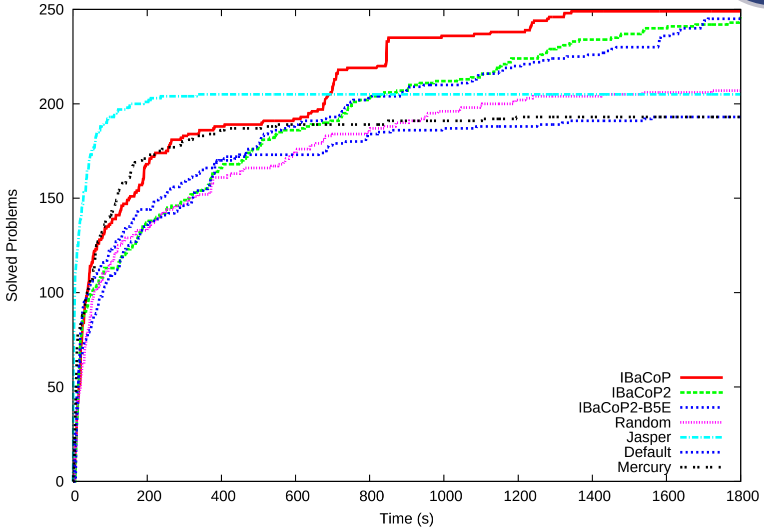
250
200
150
100
50
0
Solved Problems
0
200
400
600
800
1000
1200
1400
1600
1800
Time (s)
IBaCoP
IBaCoP2
IBaCoP2-B5E
Random
Jasper
Default
Mercury

Number of times each planner has been selected in a domain

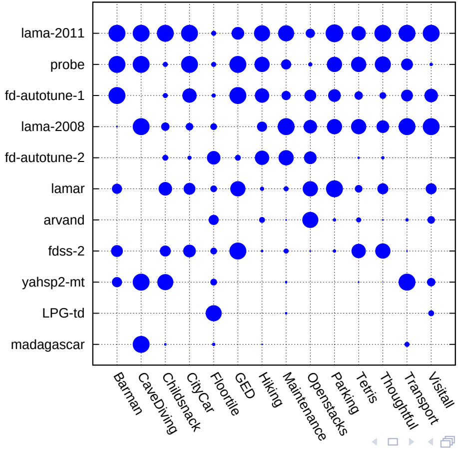

Empirical Performance Modeling

Empirical Performance Modeling may **encode knowledge** as a combination of the following capabilities:

- **Domain** discrimination
- **Size** discrimination
- **Search space** discrimination

Search Space Discrimination

- Planning EPMs have been usually trained using a set of **available benchmarks**
- Under these circumstances is **very hard** to isolate the effect of different **discrimination types**

1. **Generate** 200 problems (D) with the same size P_p
2. **Run** the problems with each planner
3. **Label** the data with different cut-off (c)
4. Apply feature **filtering criteria** with c = 66%

Execution time for the 200 problems
Barman domain with MERCURY planner

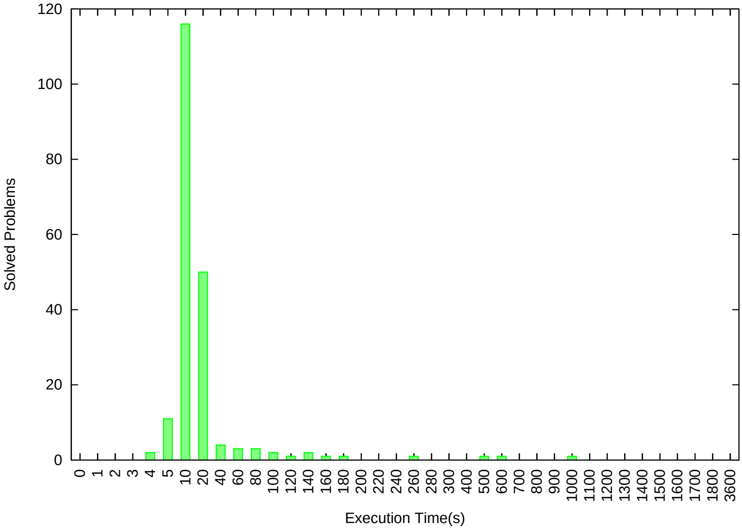

Execution time for the 200 problems

Barman domain with MERCURY planner

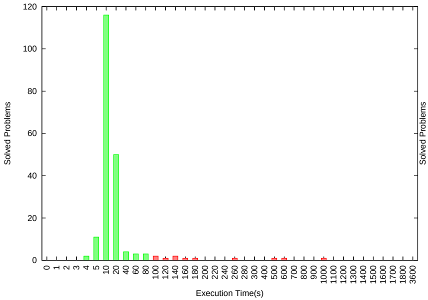

95%

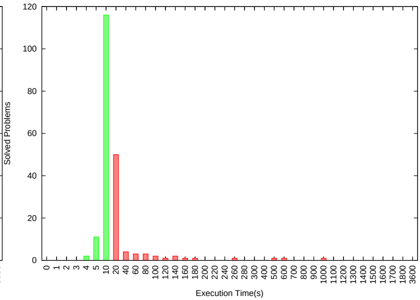

66 %

Algorithm	95% Acc	95% AUROC	66% Acc	66% AUROC
ZeroR	95.0	0.50	66.0	0.50
J48	94.5	0.50	68.0	0.62
NaiveBayes	77.0	**0.68**	67.0	**0.71**
RandomForest	94.0	0.67	66.5	0.65
RotationForest	95.0	0.51	70.0	0.64

The area under the curve (AUROC) is equal to the probability that a classifier will rank a randomly chosen positive instance higher than a randomly chosen negative one

General Feature Analysis

- **Landmark**: number of edges
- **Heuristic**: Causal Graph, FF, Landmark-cut
- **Fact Balance**: Balance distortion, Balance Ratio

Temporal Approximation

Handicaps:

- There are no features to temporal problems in the current state of the art
- State-of-the-art planning EPMs mainly focus on classical planning

Proposal:

- A **new** set of features which are specific to temporal problems
- Predict the performance of **temporal planners**

Proposal

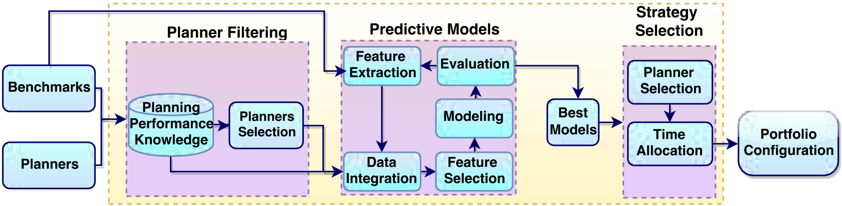
Planner Filtering
Predictive Models
Strategy Selection
Benchmarks
Planners
Planning Performance Knowledge
Planners Selection
Feature Extraction
Evaluation
Modeling
Data Integration
Feature Selection
Best Models
Planner Selection
Time Allocation
Portfolio Configuration

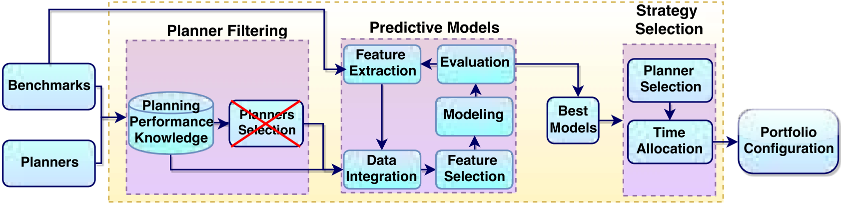

- **Planners**: 8 planners LPG-TD, POPF2, YAHSP2, YAHSP2-MT, TEMPORAL FAST DOWNWARD, ITSAT, YAHSP3 and YAHSP3-MT
- **Benchmarks**: temporal problems from IPC 2002, 2004, 2006, 2008, 2011 and 2014

Planning Task Characterization

There are **68 features** from the general procedure

Common

- PDDL
- SAS$^+$

There are **71 new** ones that are specific to temporal planning problems

New

- Temporal SAS$^+$
- Temporal PDDL
- Temporal Fast Downward

Configuration Strategies

Classification Portfolio: select **the planner** with the best confidence

Regression Portfolio: select the **faster** planner

Experimental Setting

Benchmarks:

- **Training**: IPC 2006-2011
- **Test**: IPC 2014

Additional Comparatives:

- **B4P**: is a portfolio with always best planners
- **LPG-td**: is the best planner in terms of coverage
- **Yahsp2**: is the best planner in terms of quality
- **VBS**: is the virtual best solver

Coverage and Time Score Results

	Classification	Regression	LPG-td	Yahsp2	B4P	VBS
TMS	18	18	0	0	0	18
Turn&Open	12	17	0	0	15	17
Storage	17	17	17	9	17	17
Driverlog	7	13	13	9	12	13
Floortile	20	20	20	8	20	20
MatchCellar	19	20	0	0	20	20
MapAnalyser	10	7	7	20	20	20
RTAM	0	20	20	20	20	20
Satellite	12	20	20	20	20	20
Parking	14	20	20	20	20	20
Coverage	129	**172**	117	106	164	185
IPC-Score	91.8	**129.3**	62.1	86.2	72.5	185

Conclusions

- The **multi-criteria** planner filtering method achieves a good selection without reducing **diversity**

Conclusions

- The **multi-criteria** planner filtering method achieves a good selection without reducing **diversity**
- The created **features** properly characterize the planning tasks

Conclusions

- The **multi-criteria** planner filtering method achieves a good selection without reducing **diversity**
- The created **features** properly characterize the planning tasks
- The **predictive models** based on these features have good results

Conclusions

- The **multi-criteria** planner filtering method achieves a good selection without reducing **diversity**
- The created **features** properly characterize the planning tasks
- The **predictive models** based on these features have good results
- The **configuration strategies** take advantage from the predictive models

Conclusions

- The **multi-criteria** planner filtering method achieves a good selection without reducing **diversity**
- The created **features** properly characterize the planning tasks
- The **predictive models** based on these features have good results
- The **configuration strategies** take advantage from the predictive models
- IBaCoP2 shows benefits over IBaCoP

Conclusions

- The **multi-criteria** planner filtering method achieves a good selection without reducing **diversity**
- The created **features** properly characterize the planning tasks
- The **predictive models** based on these features have good results
- The **configuration strategies** take advantage from the predictive models
- IBaCoP2 shows benefits over IBaCoP
- The portfolios achieve **remarkable** results

Conclusions

- The **multi-criteria** planner filtering method achieves a good selection without reducing **diversity**
- The created **features** properly characterize the planning tasks
- The **predictive models** based on these features have good results
- The **configuration strategies** take advantage from the predictive models
- IBaCoP2 shows benefits over IBaCoP
- The portfolios achieve **remarkable** results
- First **Temporal Approximation**

Conclusions

- The **multi-criteria** planner filtering method achieves a good selection without reducing **diversity**
- The created **features** properly characterize the planning tasks
- The **predictive models** based on these features have good results
- The **configuration strategies** take advantage from the predictive models
- IBaCoP2 shows benefits over IBaCoP
- The portfolios achieve **remarkable** results
- First **Temporal Approximation**
- The **relevance** of each **feature** is not dominant across different domains and planners

Future Work

- The **automated selection** of the number of planners per planning task
- Incorporate the **synergy** between different automated planners for the portfolio configuration
- Incorporate **new features** to regression tasks
- Evaluate a **portfolio** in homogeneous problems sets

Publications

- *Tomás de la Rosa, Isabel Cenamor and Fernando Fernández*, **'Performance Modelling of Planners from Homogeneous Problem Sets'**. In the 27th International Conference on Automate Planning and Scheduling 2017.
- *Isabel Cenamor, Tomás de la Rosa, and Fernando Fernández*, **'The IBaCoP planning system: Instance-based configured portfolios'**, Journal of Artificial Intelligence Research (JAIR) N 56.
- *Isabel Cenamor, Tomás de la Rosa, and Fernando Fernández*, **'Learning Predictive Models to Configure Planning Portfolios'**, Workshop Planning and Learning ICAPS-2013
- *Isabel Cenamor, Tomás de la Rosa, and Fernando Fernández*, **'Mining IPC-2011 Results'**, Workshop on International Planning Competition ICAPS-2012

Awards
in the International Planning Competition

- ★ Winner at Sequential Satisficing track
- ★ Runner up at Sequential Satisficing Multi-core track

Thank you for your attention!

Creating Planning Portfolios with Predictive Models
Isabel Cenamor
Advisors: Tomás de la Rosa and Fernando Fernández

uc3m | Universidad Carlos III de Madrid

[GSV14] Alfonso Gerevini, Alessandro Saetti, and Mauro Vallati.
Planning through automatic portfolio configuration: The PbP approach.
Journal of Artificial Intelligence Research, 50:639–696, 2014.

[HDH+99] Adele E. Howe, Eric Dahlman, Christoper Hansen, Michael Scheetz, and Anneliese von Mayrhauser.
Exploiting competitive planner performance.
In Susanne Biundo and Maria Fox, editors, *Recent Advances in AI Planning, 5th European Conference on Planning, ECP'99, Durham, UK, September 8-10, 1999, Proceedings*, volume 1809 of *Lecture Notes in Computer Science*, pages 62–72. Springer, 1999.

[HRS+11] Malte Helmert, Gabriele Röger, Jendrik Seipp, Erez Karpas, Jörg Hoffmann, Emil Keyder, Raz Nissim, Silvia Richter, and Matthias Westphal.
Fast downward stone soup.
The Seventh International Planning Competition, IPC-7 planner abstracts:38, 2011.

[MWK14] Yuri Malitsky, David Wang, and Erez Karpas.
The AllPACA planner: All planners automatic choice algorithm.
IPC 2014 planner abstracts, pages 71–73, 2014.

[NBL15] Sergio Núñez, Daniel Borrajo, and Carlos Linares López.
Automatic construction of optimal static sequential portfolios for AI planning and beyond.
Artificial Intelligence, 226:75–101, 2015.

[SBGH12] Jendrik Seipp, Manuel Braun, Johannes Garimort, and Malte Helmert.
Learning portfolios of automatically tuned planners.
In Lee McCluskey, Brian Williams, José Reinaldo Silva, and Blai Bonet, editors, *Proceedings of the Twenty-Second International Conference on Automated Planning and Scheduling, ICAPS 2012, Atibaia, São Paulo, Brazil, June 25-19, 2012*. AAAI, 2012.

[SSHH15] Jendrik Seipp, Silvan Sievers, Malte Helmert, and Frank Hutter.
Automatic configuration of sequential planning portfolios.
In Blai Bonet and Sven Koenig, editors, *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, January 25-30, 2015, Austin, Texas, USA.*, pages 3364–3370. AAAI Press, 2015.

[VCK14] Mauro Vallati, Lukáš Chrpa, and Diane Kitchin.
ASAP: an automatic algorithm selection approach for planning.
International Journal on Artificial Intelligence Tools, 23(06):1460032, 2014.

[VNM+14] Richard Valenzano, Hootan Nakhost, Martin Müller, Jonathan Schaeffer, and N Sturtevant.
Arvandherd 2014.
IPC 2014 planner abstracts, pages 11 – 14, 2014.

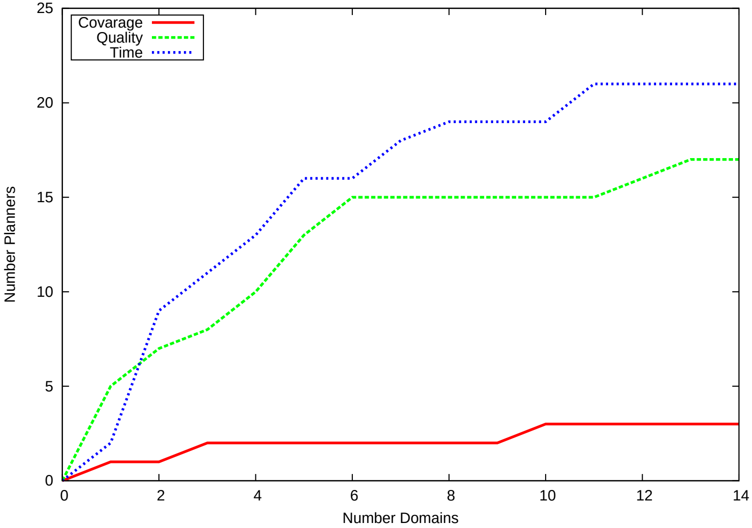

Covarage
Quality
Time
Number Planners
Number Domains
0
5
10
15
20
25
0
2
4
6
8
10
12
14

Algorithm for computing the positive and negative balance footprints for a layer of the RPG.

- *RP_init* Minimum, average and variance of the number of times that a fact in the initial state is deleted in the computation of the relaxed plan. $(\mathcal{B}(p, \pi^{\pm}_{s_0}), \forall p \in \mathcal{S})$. (3)
- *RP_goal*Minimum, average and variance of the number of times that a goal is deleted in the computation of the relaxed plan. $(\mathcal{B}(g, \pi^{\pm}_{s_0}), \forall g \in s_\star)$(3)
- *Ratio_ff* Ratio between the value of the max and FF heuristic. This proportion shows the idea of parallelization of the relaxed plan.
- *RP Balance Ratio* Aggregate the value of each layer multiplying it by a weight that represents the proportion of actions that appear in each particular layer of the occurrences in which a fact has a positive balance.
  $\sum_{i=1}^{layers(RPG)} \frac{|A_{i-1}|}{|\mathcal{A}|} \times fp_i^+$
- *RP Unbalance Ratio* Aggregate the value of each layer multiplying it by a weight that represents the proportion of actions that appear in each particular layer of the occurrences in which a fact has a negative balance. $\sum_{i=1}^{layers(RPG)} \frac{|A_{i-1}|}{|\mathcal{A}|} \times fp_i^-$
- *Balance Distorsion* Aggregate the value of each layer for the distorsion of unbalanced facts.
  $\sum_{i=1}^{layers(RPG)} dist\_fp_i$

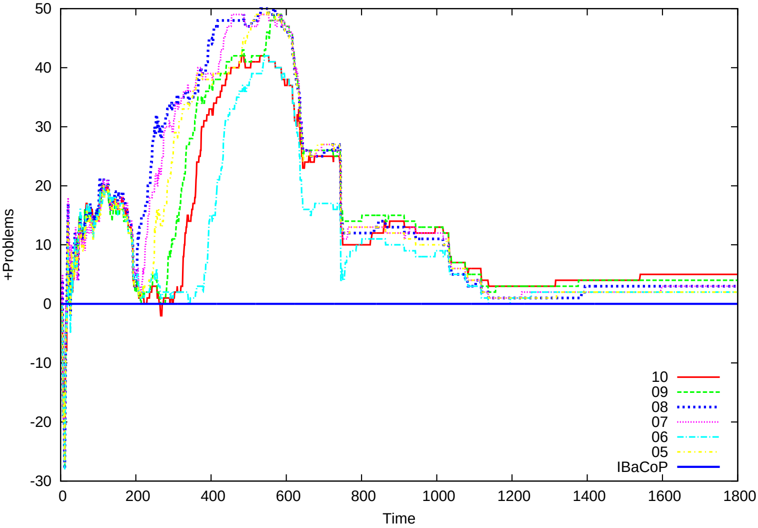

50
40
30
20
10
0
-10
-20
-30
0
200
400
600
800
1000
1200
1400
1600
1800
Time
+Problems
10
09
08
07
06
05
IBaCoP