

To Max or not to Max: Online Learning for Speeding Up Optimal Planning

C. Domshlak E. Karpas S. Markovitch

Faculty of Industrial Engineering and Management

Faculty of Computer Science
Technion

May 27, 2010



Outline

- 1 Motivation
- 2 Theoretical Model
- 3 From Model to Practice
 - Dealing with Model Assumptions
 - Learning
 - Using the Classifier
- 4 Experimental Evaluation



Motivation

- Domain independent optimal planning
 - A^* + admissible heuristic (almost always)
 - Which heuristic to use?
- Sample results:

Domain	h_{LA}	h_{LM-CUT}
airport	25	38
freecell	28	15



Combining Admissible Heuristics

- Why use only one heuristic?
- Simplest combination method: \max_h
- Sample results:

Domain	h_{LA}	h_{LM-CUT}	\max_h
airport	25	38	36
freecell	28	15	22

- Other combination methods exist (additive heuristics, additive/disjunctive, ...)



Combining Admissible Heuristics (2)

- The problem with \max_h
 - We need to compute many heuristic functions
 - The heuristic value is the result of only one computation
 - Some computation is wasted
- Possible solution: learn a classifier which predicts which heuristic will be the “winner”



Informative vs. Fast Heuristics

- Sometimes spending a lot of time to compute the most informed heuristic is not the best thing to do



Outline

- 1 Motivation
- 2 Theoretical Model
- 3 From Model to Practice
 - Dealing with Model Assumptions
 - Learning
 - Using the Classifier
- 4 Experimental Evaluation



Theoretical Model - Which Heuristic to Compute When?

Assumptions


- State space is a tree
- Single goal state
- Uniform cost actions
- Constant branching factor b
- Perfect knowledge

Two heuristics: h_1 and h_2


- Consistent
- Evaluating h_i takes time t_i




Theoretical Model - Which Heuristic to Compute When?




Theoretical Model - Which Heuristic to Compute When?




Theoretical Model - Which Heuristic to Compute When?




Theoretical Model - Which Heuristic to Compute When?




Theoretical Model - Which Heuristic to Compute When?




Theoretical Model - Which Heuristic to Compute When?



Theoretical Model - Which Heuristic to Compute When?




Theoretical Model - Which Heuristic to Compute When?



Theoretical Model - Which Heuristic to Compute When?

Using h_2 - 1 evaluation, t_2 time


Using h_1



Theoretical Model - Which Heuristic to Compute When?

Using h_2 - 1 evaluation, t_2 time

Using h_1




We need to expand this region



Theoretical Model - Which Heuristic to Compute When?

Using h_2 - 1 evaluation, t_2 time


Using h_1
Expand region



Theoretical Model - Which Heuristic to Compute When?

Using h_2 - 1 evaluation, t_2 time


Using h_1
Expand region
Evaluate leaves



Theoretical Model - Which Heuristic to Compute When?

Using h_2 - 1 evaluation, t_2 time

Using h_1
Expand region
Evaluate leaves
 $b^l t_1$ time




l - number of levels to expand



Theoretical Model - Which Heuristic to Compute When?

Using h_2 - 1 evaluation, t_2 time


Using h_1 - $b^l t_1$ time



Theoretical Model - Which Heuristic to Compute When?

Using h_2 - 1 evaluation, t_2 time
Using h_1 - $b^l t_1$ time

Best decision - use h_2 iff $t_2 < b^l t_1$




Theoretical Model - Which Heuristic to Compute When?


Using h_2 - 1 evaluation, t_2 time
Using h_1 - $b^l t_1$ time

Best decision - use h_2 iff $t_2 < b^l t_1$


i.e. $l > \log_b(t_2/t_1)$




Theoretical Model - Which Heuristic to Compute When?



Theoretical Model - Which Heuristic to Compute When?

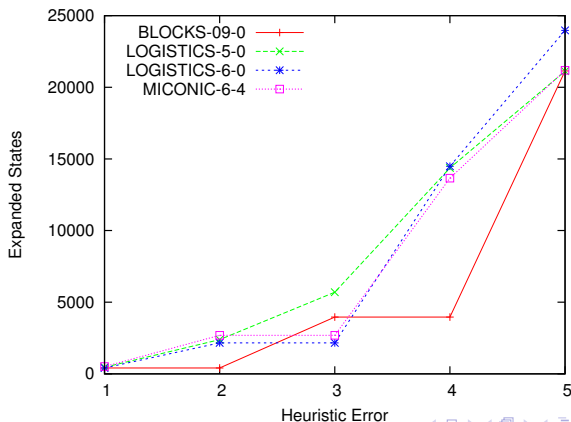


Theoretical Model - Which Heuristic to Compute When?



Justifying the Rule

- The decision rule derived from the model can be justified by some empirical results (Helmert and Röger, 2008)



Outline

- 1 Motivation
- 2 Theoretical Model
- 3 From Model to Practice**
 - Dealing with Model Assumptions
 - Learning
 - Using the Classifier
- 4 Experimental Evaluation



Dealing with Model Assumptions

Assumptions

- State space is a tree
- Single goal state
- Uniform cost actions
- Constant branching factor b
- Perfect knowledge

Two heuristics: h_1 and h_2

- Consistent
- Evaluating h_i takes time t_i



Dealing with Model Assumptions

Assumptions

- State space is a tree - doesn't change the rule
- Single goal state - doesn't change the rule
- Uniform cost actions - doesn't change the rule
- Constant branching factor b
- Perfect knowledge

Two heuristics: h_1 and h_2

- Consistent - doesn't change the rule
- Evaluating h_i takes time t_i



Dealing with Model Assumptions

Assumptions

- State space is a tree - doesn't change the rule
- Single goal state - doesn't change the rule
- Uniform cost actions - doesn't change the rule
- Constant branching factor b - estimate
- Perfect knowledge

Two heuristics: h_1 and h_2

- Consistent - doesn't change the rule
- Evaluating h_i takes time t_i - estimate



Dealing with Model Assumptions

Assumptions

- State space is a tree - doesn't change the rule
- Single goal state - doesn't change the rule
- Uniform cost actions - doesn't change the rule
- Constant branching factor b - estimate
- Perfect knowledge - use decision rule at every state

Two heuristics: h_1 and h_2

- Consistent - doesn't change the rule
- Evaluating h_i takes time t_i - estimate




Learning

- Collecting training examples
- Labeling training examples
- Generating features
- Building a classifier



Collecting Training Examples

- State space is sampled using stochastic hill climbing “probes”
 - Depth limit = $2 * h(s_0)$
 - Probability of expanding successor $s = 1/h(s)$
- All *generated* states are added to the training set
- Probing stops when enough training examples are collected



Labeling Training Examples

- b, t_1, t_2 are estimated from the collected examples
- $h_2 - h_1$ is calculated for each state
- Each example is labeled by h_2 iff $h_2 - h_1 > \alpha \log_b(t_2/t_1)$
- WLOG $t_2 > t_1$ - the choice is always whether to evaluate the more expensive heuristic



Generating Features

- We use the simplest features - values of state variables
- Better features will probably lead to better results




Building a Classifier

- We use the Naive Bayes classifier
 - Very fast
 - Incremental
 - Provides probability distribution for classification




Using the classifier

State Evaluation



Using the classifier

State Evaluation



Final Remarks

- This is an active online learning scheme
- Using SAS^+ helps, because it reduces dependence between state variables
- This approach can be easily extended to multiple heuristics
 - Learn a classifier for each pair
 - Decide which heuristic to use by voting



Outline

- 1 Motivation
- 2 Theoretical Model
- 3 From Model to Practice
 - Dealing with Model Assumptions
 - Learning
 - Using the Classifier
- 4 Experimental Evaluation



Evaluation

- We used two state of the art heuristics
 - $h_{\text{LM-CUT}}$ - Helmert and Domshlak 2009
 - h_{LA} - Karpas and Domshlak 2009
- Parameters
 - $\alpha = 1$ - decision rule bias
 - $\rho = 0.6$ - confidence threshold
 - Training set size = 100



Results - Solved Problems

Domain	h_{LA}	h_{LM-CUT}	max_h	rnd_h	sel_h
airport	25	38	36	29	36
blocks	20	28	28	28	28
depots	7	7	7	7	7
driverlog	14	14	14	14	14
freecell	28	15	22	15	28
grid	2	2	2	2	2
gripper	6	6	6	6	6
logistics-2000	19	20	20	20	20
logistics-98	5	6	6	5	6
miconic	140	140	140	140	140
mprime	21	25	25	19	25
mystery	13	17	17	14	17
openstacks	7	7	7	7	7
pathways	4	5	5	4	5
psr-small	48	49	48	48	48
pw-notankage	16	17	17	17	17
pw-tankange	9	11	11	10	11
rovers	6	7	7	6	7
satellite	7	8	9	7	9
tpp	6	6	6	6	6
trucks	7	10	9	7	9
zenotravel	9	12	12	10	12
Total	419	450	454	421	460




Results - Time

Domain	h_{LA}	h_{LM-CUT}	max_h	rnd_h	sel_h
airport (25)	125.96	35.36	73.80	54.78	68.44
blocks (20)	66.01	3.71	6.39	6.44	5.59
depots (7)	196.91	65.99	103.26	155.14	94.36
driverlog (14)	66.67	110.87	86.04	120.84	81.31
freecell (15)	6.04	249.28	23.93	44.22	9.25
grid (2)	12.05	33.78	44.27	38.3	40.26
gripper (6)	71.6	106.48	264.79	161.98	77.07
logistics-2000 (19)	73.32	152.27	255.36	153.89	79.17
logistics-98 (5)	18.84	24.11	29.55	28.69	24.43
miconic (140)	2.03	8.04	10.08	5.67	7.65
mprime (19)	17.52	17.9	15.68	111.48	8
mystery (12)	7.55	1.61	2.03	57.93	2.49
openstacks (7)	15.93	72.3	75.83	52.69	17.11
pathways (4)	5.38	0.08	0.14	1.15	0.18
psr-small (48)	3.55	4.05	7.92	5.73	4.87
pw-notankage (16)	48.8	71.34	71.49	73.92	59
pw-tankage (9)	211.43	173.61	189.89	172.99	130.98
rovers (6)	122.7	5.23	8.79	45.72	7.97
satellite (7)	46.22	3.47	4.51	21.95	3.58
tpp (6)	108.54	14.36	5.9	56.32	5.69
trucks (7)	238.85	11.69	16.48	39.64	15.56
zenotravel (9)	9.84	0.91	1.33	8.27	1.28
Average (Problem)	39.65	38.59	41.39	42.6	24.53
Average (Domain)	67.08	53.02	58.97	64.44	33.83



Anytime Behavior



Thank You

- Thank You

