

FF: The Fast-Forward Planning System







Introduction

- The winner of the most successful automatic planner at the Fifth International Conference on AI Planning and Scheduling (AIPS'00)
- Builds upon Heuristic Search Planner (HSP)
 - Winner of AIPS'98
- Key differences between FF and HSP:
 - More sophisticated **heuristic** evaluation
 - weight value computation instead of relaxed plan extraction
 - Uses systematic **search** for escaping local minima
 - \circ $\,$ enforced hill climbing vs hill climbing
 - Use of "helpful action pruning"

Heuristic HSP

For every fact f If $f \in$ search state S : Set weight(f) = 0 Else Set weight(f) = ∞ For all actions with precondition pre(o) that adds a fact f Update weight(f) = min(weight(f), weight(pre(o))+1)

S represents the search state and G represents the goal state

 $\text{heuristic}_{HSP}(S) = \sum_{g \in G} \text{weight}(g)$





Heuristic FF

- <u>Construct planning graph</u> that it alternates between fact and action layers, begins with the initial state as the first fact layer and the final layer represents all goal facts
- Extract a relaxed plan s.t. starting at the goal fact layer, at each layer i
 - If a goal is NOT present in previous layer i-1, select an action from the previous layer that adds the goal
 - Else, simply insert it into the goals to be achieved at i-1
- Given relaxed plan (O₀, O₁, ..., O_{m-1}) such that O_i is set of actions selected at time i

 $\text{heuristic}_{FF}(S) = \sum_{i=0,...,m-1} |O_i|$





Example







Search / Hill Climbing

- FF uses enforced hill climbing:
 - 1. Evaluate all direct successors of current state
 - 2. If a better state (lower heuristic) is found, move to it
 - 3. If no better state is found, expand next level in a BFS fashion
 - 4. Continue until goal is reached





Helpful Actions

- Relaxed plan computed during heuristic calculation
- The first set of actions in this relaxed plan are identified as <u>helpful actions</u> because they contribute directly to the next goal(s) in the relaxed plan.
- The search is restricted to only consider successors generated by these actions instead of expanding all applicable actions.

 $H(S) = \{o \mid pre(o) \subseteq S, add(o) \cap G_1 \neq \emptyset\}$

 o is an action, pre(o) action's preconditions, add(o) are action's add effects, and G₁ is set of subgoals identified at first level of relaxed plan





- Test performance of FF vs HSP 1.0 focusing on 3 major differences:
 - 1. <u>Heuristics</u>: Relaxed plan extraction vs weight value computation
 - 2. Search: Enforced hill climbing vs hill climbing
 - 3. Helpful action: Helpful action pruning vs no such pruning
- Each difference is attached to a switch that can be turned on or off
- Each planner ran on large set of benchmarks across 20 different domains



Results

Distance		Hill Clim	bing		Enforced Hill Climbing				
Estimate	All Actio	ons	Helpful	Helpful Actions		All Actions		Helpful Actions	
	Time	Length	Time	Length	Time	Length	Time	Length	
HSP distance	2	2	1	2	2	0	1	0	
FF distance	12	2	12	5	11	9	9	11	
Search Strategy	All Actions				Helpful Actions				
0,	HSP Distance		FF Distance		HSP Distance		FF Distance		
	Time	Length	Time	Length	Time	Length	Time	Length	
Hill Climbing	5	1	9	1	3	2	1	2	
Enforced Hill Climbing	9	8	8	10	16	6	16	9	
Pruning Technique		Hill Clim	bing		Enforced Hill Climbing				
	HSP Distance		FF Distance		HSP Distance		FF Distance		
	Time	Length	Time	Length	Time	Length	Time	Length	
HSP Distance	2	0	3	0	2	1	2	0	
EF Distance	13	7	14	8	15	5	15	3	

 Table 1. Comparison of Related Planners When Varying on Goal Distance Estimates, Search Strategies, or Pruning Technique, from Top to Bottom.



Distance Estimates

Distance Estimate		Hill Clim	bing		Enforced Hill Climbing				
	All Actions		Helpful Actions		All Actions		Helpful Actions		
	Time	Length	Time	Length	Time	Length	Time	Length	
HSP distance	2	2	1	2	2	0	1	0	
FF distance	12	2	12	5	11	9	9	11	

- FF's estimate improve run-time for about half of the domains across all alignments
- With enforced hill climbing, FF's advantage has a much higher advantage



Enforced hill climbing vs hill climbing

Search Strategy	All Actions				Helpful Actions				
	HSP Distance		FF Distance		HSP Distance		FF Distance		
	Time	Length	Time	Length	Time	Length	Time	Length	
Hill Climbing	5	1	9	1	3	2	1	2	
Enforced Hill Climbing	9	8	8	10	16	6	16	9	

- Without pruning helpful actions, enforced hill climbing degrades as many times as it improves it
- When helpful actions pruned, enforced hill climbing is faster in 80% of the 20 domains



Helpful actions vs All actions

Pruning Technique	Hill Climbing				Enforced Hill Climbing			
	HSP Distance		FF Distance		HSP Distance		FF Distance	
	Time	Length	Time	Length	Time	Length	Time	Length
HSP Distance	2	0	3	0	2	1	2	0
FF Distance	13	7	14	8	15	5	15	3

- Helpful action pruning is faster in about 75% of domains across all alignments
- Only one domain showed increase in solution length when using helpful actions





Conclusion

- FF simple but effective improvement on HSP
- Key differences: Heuristic quality, enforced hill climbing, helpful actions
- Future improvements needed for more complex domains

