

The Joy of Forgetting: Faster Anytime Search via Restarting

Silvia Richter

IIIS, Griffith University, Australia
and NICTA
silvia.richter@nicta.com.au

Jordan T. Thayer and Wheeler Ruml

Department of Computer Science
University of New Hampshire
{jtd7, ruml} at cs.unh.edu

Abstract

Anytime search algorithms solve optimisation problems by quickly finding a (usually suboptimal) first solution and then finding improved solutions when given additional time. To deliver an initial solution quickly, they are typically greedy with respect to the heuristic cost-to-go estimate h . In this paper, we show that this low- h bias can cause poor performance if the greedy search makes early mistakes. Building on this observation, we present a new anytime approach that restarts the search from the initial state every time a new solution is found. We demonstrate the utility of our method via experiments in PDDL planning as well as other domains, and show that it is particularly useful for problems where the heuristic has systematic errors.

Introduction

Heuristic search is a widely used framework for solving shortest-path problems. In particular, given a consistent heuristic and sufficient resources, the A* algorithm (Hart *et al.* 1968) can be used to find an optimal solution with maximum efficiency (Dechter and Pearl 1988). However, in large problems we may not be able to afford the resources necessary for calculating an optimal solution. A popular approach in this circumstance is complete anytime search, in which a (typically suboptimal) solution is found quickly, followed over time by a sequence of progressively better solutions, until the search is terminated or solution optimality has been proven.

A* can be modified to trade solution quality for speed by weighting the heuristic by a factor $w > 1$ (Pohl 1970). The resulting algorithm, called *weighted A** or WA* for short, searches more greedily the larger w is. Assuming an admissible heuristic, the cost of the solution it finds differs from the optimal by no more than the factor w . This ability to balance speed against quality while providing bounds on the suboptimality of solutions makes WA* an attractive basis for anytime algorithms (Hansen *et al.* 1997; Likhachev *et al.* 2004; Hansen and Zhou 2007; Likhachev *et al.* 2008). For example, WA* can be run first with a very high weight, resulting in a greedy best-first search that tries to find a solution as quickly as possible. Then the search can be continued past the first solution to find better ones. By reducing the weight over time, the search can be made progressively less greedy and more focused on quality.

Complete anytime algorithms that are not based on WA* have also been proposed (Zhang 1998; Zhou and Hansen 2005; Aine *et al.* 2007). In these cases, an underlying global search algorithm is usually made more greedy and local by initially restricting the set of states that can be expanded, thus giving a depth-first emphasis to the search which can lead to finding a first solution quickly. Whenever a solution is found, the restrictions are relaxed, allowing for higher-quality solutions to be found.

We start by noting that the greediness employed by anytime algorithms to quickly find a first solution can also create problems for them. If the heuristic goal distance estimates are inaccurate, for example, the search can be led into an area of the search space that contains no goal or only poor goals. Anytime algorithms typically *continue* their search after finding the first solution, rather than starting a new search. This seems reasonable as it avoids duplicate effort. However, a key observation that we will discuss in detail below is that continued search may perform badly in problems where the initial solution is far from optimal and finding a significantly better solution requires searching a different part of the search space. In such cases *restarting* the search can be beneficial, as it allows changing bad decisions that were made near the start state more quickly. Reconsidering these early decisions means the quality of the best-known solution may improve faster with restarts than if we continued to explore the area of the search space around the previous solution.

Building on this observation that forgetfulness can be bliss, we propose an anytime algorithm that incorporates helpful restarts while retaining most of the knowledge discovered in previous search phases. Our approach differs substantially from existing anytime approaches, as previous work on anytime algorithms has concentrated explicitly on *avoiding* duplicate work (Likhachev *et al.* 2004; 2008). We show that the opposite can be useful, namely discarding information that may be biased due to greediness. We show that our method outperforms existing anytime algorithms in PDDL planning and that it is competitive in other benchmark domains. Using an artificial search space, we elucidate conditions under which our approach performs well.

Previous Approaches

We first present some existing approaches and discuss the circumstances under which their effectiveness may suffer.

A* and WA* Several anytime algorithms are based on Weighted A* (WA*), which is in turn based on A*. The A* algorithm uses two data structures: Open and Closed. In the beginning, Open contains only the start state and Closed is empty. Iteratively one of the states from Open is selected, moved from Open to Closed, and expanded, and its successor states are inserted into Open. The state expanded at each step is one that minimises the function $f(s) = g(s) + h(s)$, where $g(s)$ is the cost of the best path currently known from the start state to state s , and $h(s)$ is the heuristic value of s , i. e., an estimate of the cost of the best path from s to a goal.

WA* is a variant of A* using the selection function $f'(s) = g(s) + w \cdot h(s)$, with $w > 1$. Even if the heuristic function h is *admissible*, i. e., never overestimates the true goal distance, the weighting by w introduces inadmissibility. However, the solution is guaranteed to be *w-admissible*, i. e. the ratio between its cost and the cost of an optimal solution is bounded by w .

Anytime Algorithms Based on WA* Anytime algorithms based on WA* run WA* but do not stop upon finding a solution. Rather, they report the solution and continue the search, possibly adjusting their data structures and parameters. A straightforward adjustment, for example, is to use the cost of the best known solution for pruning: after we have found a solution of cost c , we do not need to consider any states that are guaranteed to lead to solutions of cost c or more. This is the approach taken by Hansen and Zhou in their Anytime WA* algorithm (2007). Updating the cost bound is the only adjustment in Anytime WA*; the weight w is never decreased. (While the authors discuss this option, they did not find it to improve results on their benchmarks.)

The Anytime Repairing A* (ARA*) algorithm by Likhachev et al. (2004; 2008) assumes an admissible heuristic and reduces the weight w each time it proves w -admissibility of the incumbent solution. This implicitly prunes the search space as no state is ever expanded whose f' -value is larger than that of the incumbent solution. When reducing w , ARA* updates the f' -values of all states in Open according to the new weight. Furthermore, ARA* avoids re-expanding states within search phases (where we use “phase” to denote the part of search between two weight changes). Whenever a shorter path to a state is discovered, and that state has already been expanded in the current search phase, the state is not re-expanded immediately. Instead it is suspended in a separate list, which is inserted into Open only at the beginning of the next phase. The rationale behind this is that even without re-expanding states the next solution found is guaranteed to be within the current sub-optimality bound (Likhachev et al. 2004; 2008). Of course, it may be that this solution is worse than it would have been had we re-expanded.

Other Anytime Algorithms Recent anytime algorithms that are not based on WA* include the A*-based Anytime Window A* algorithm (Aine et al. 2007) and beam-stack

search (Zhou and Hansen 2005) which is based on breadth-first search. In both cases, the underlying global search algorithm is made more greedy by initially restricting the set of states that can be expanded. In Window A*, this is done by using a “window” that slides downwards in a depth-first fashion: whenever a state with a larger g -value (level) than previously expanded is expanded, the window slides down to that level of the search space. From then on, only states in that level and the k levels above can be expanded, where k is the height of the window. Initially k is zero and it is increased by 1 every time a new solution is found. Anytime Window A* can suffer from its strong depth-first focus if the heuristic estimates are inaccurate.

In beam-stack search, only the b most promising nodes in each level of the search space are expanded, where the beam width b is a user-supplied parameter. The algorithm remembers which nodes have not yet been expanded and returns to them later. However, beam-stack search explores the entire search space below the chosen states before it backtracks on its decision. This may be inefficient if the heuristic is (locally) quite inaccurate and a wrong set of successor states is chosen, e. g. states from which no near-optimal solution can be reached.

The Effect of Low- h Bias

Anytime WA* and ARA* keep the Open list between search phases (possibly re-ordering it). After finding a goal state s_g , Open will usually contain many states that are close to s_g in the search space because the ancestors of s_g have been expanded; furthermore, those states are likely to have *low heuristic values* because of their proximity to s_g . Hence, even if we update Open with new weights, we are likely to expand most of the states around s_g before considering states that are close to the start state. This *low- h bias* can be a critical mistake if early decisions strongly influence final solution quality.

Consider the small gridworld example in Fig. 1. The task is to reach a goal state ($G1$ or $G2$) from the start state S , where the agent can move with cost 1 to each of the 8 neighbours of a cell if they are not blocked. The heuristic has a systematic bias, underestimating goal distances in the left part of the grid. We start with weight 2 in Fig. 1a. Because the heuristic values to the left of S are lower than to the right of S , the search expands the states to the left and finds goal $G1$ with cost 6. The grey cells are generated, but not expanded in this search phase, i. e., they are in Open. After finding a solution with a cost of 6, in Fig. 1b the search continues with a reduced weight of 1.5. A solution with cost 5 consists in turning right from S and going to $G2$. However, the search will first expand all states in Open that have an f' -value smaller than 7, including many nodes near the original path. After expanding a substantial number of states, the second solution it finds is a suboptimal path that shares several states with the first solution, again with cost 6. If we instead *restart* with an empty Open list after finding the first solution (Fig. 1c), far fewer states are expanded. The critical state to the right of S is expanded quickly and the optimal path is found. We will see this phenomenon again below in the experimental section of this paper.

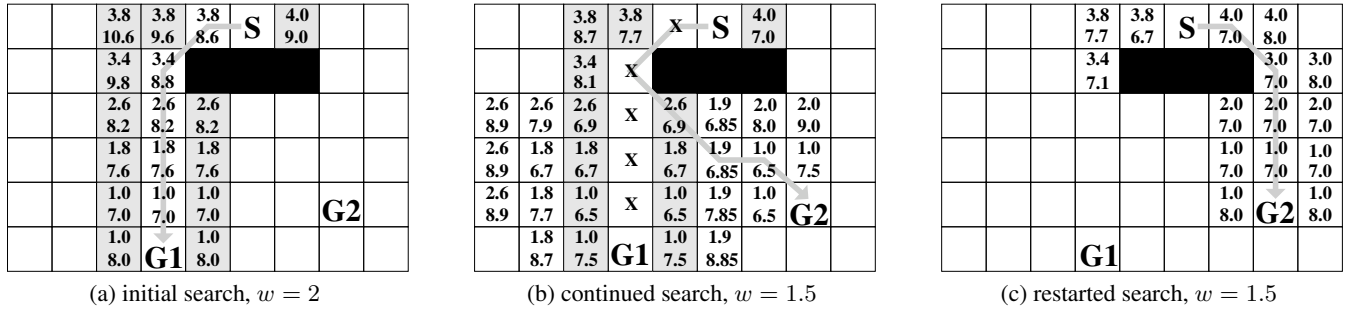


Figure 1: The effect of low- h bias. For all grid states generated by the search, h -values are shown above f' -values. (a) Initial WA* search finds a solution of cost 6. (b) Continued search expands many states around the previous Open list (grey cells), finding another sub-optimal solution of cost 6. (c) Restarted search quickly finds the optimal solution of cost 5.

Restarting WA* (RWA*)

To overcome low- h bias, we propose Restarting WA*, or RWA* for short. It runs iterated WA* with decreasing weight, always re-expanding states when it encounters a cheaper path. RWA* differs from ARA* and Anytime WA* by not keeping the Open list between phases. Whenever a better solution is found, the search empties Open and restarts from the initial state. It does, however, re-use search effort in the following way: besides Open and Closed, we keep a third data structure “Seen”. When the new search phase begins, the states from the old Closed list are moved to Seen. When RWA* generates a state in the new search, there are three possibilities: Case 1: The state has never been generated before (it is neither in Open nor Closed nor Seen). In this case, RWA* behaves like WA*, calculates the heuristic value of the state and inserts it into Open. Case 2: The state has been encountered in previous search phases but not in the current phase (it is in Seen). In this case, RWA* can look up the heuristic value of the state rather than having to calculate it. In addition, RWA* checks whether it has found a cheaper path to the state or whether the previously found path is better, and keeps the better one. The state is removed from Seen and put into Open. Case 3: The state has already been encountered in this search phase (it is in Open or Closed). In this case, RWA* again behaves like WA*, re-inserting the state into Open only if it has found a shorter path to the state. Complete pseudo-code is shown in Fig. 2. This strategy can be implemented in an efficient way by maintaining a boolean value “seen” in each state of the Open/Closed list, rather than keeping the seen states in a separate list.

In short, previous search effort is re-used by not calculating the heuristic value of a state more than once and by making use of the best path to a state found so far. Compared to ARA* and Anytime A*, our method may have to re-expand many states that were already expanded in previous phases. However, in cases where the calculation of the heuristic accounts for the largest part of computation time (e. g., in our planning experiments it is 80%), the duplicated expansions do not have much effect on runtime. Thus, RWA* re-uses most of the previous search effort, but its restarts allow for

more flexibility in discovering different solutions.

Restarts in Other Search Paradigms

Restarts are a well-known and successful technique in combinatorial search and optimisation, e. g. in the areas of propositional satisfiability and constraint-satisfaction problems. Together with randomisation, they are used in systematic search (Gomes *et al.* 1998) as well as local search (Selman *et al.* 1992) to escape from barren areas of the search space. A restart is typically executed if no solution has been found after a certain number of steps. In most of these approaches restarting would be useless without randomisation, as the algorithm would behave exactly the same each time. By contrast, our algorithm is deterministic, and its restarts serve the purpose of revisiting nodes near the start of the search tree to promote exploration when the node evaluation function has changed.

This is perhaps most closely related to the motivation behind limited-discrepancy search (LDS) (Harvey and Ginsberg 1995; Furcy and Koenig 2005). LDS attempts to overcome the tendency of depth-first search to visit solutions that differ only in the last few steps. It backtracks whenever the current solution path exceeds a given limit on the number of discrepancies from the greedy path, restarting from the root each time this limit is increased. Like random restarts and the backtracking in limited-discrepancy search, our restarts mitigate the effect of bad heuristic recommendations early in the search. Our use of restarts differs from limited-discrepancy backtracking in that we restart upon finding a solution, rather than after expanding all nodes with a given number of discrepancies from the greedy path. As opposed to random restarts, we restart upon success (when finding a new solution) rather than upon failure (when no solution could be found). A depth-first approach like LDS is however not competitive for shortest-path problems like planning.

Lastly, incomplete methods can be made complete by restarting with increasingly relaxed heuristic restrictions (Zhang 1998; Aine *et al.* 2007). By contrast, our restarts occur within an algorithm that is already complete.

To our knowledge, restarts have not been used in a complete best-first search before. Our use of restarts in an A*-

```

RWA*( $w_0, \phi$ )
1   $bound \leftarrow \infty, w \leftarrow w_0, Seen \leftarrow \emptyset$ 
2  while not interrupted and not failed
3  do  $Closed \leftarrow \emptyset, Open \leftarrow \{startstate\}$ 
4  while not interrupted and  $Open$  not empty
5  do remove  $s$  with minimum  $f'(s)$  from  $Open$ 
6  for  $s' \in SUCCESSORS(s)$ 
7  do if heuristic admissible and  $f(s') \geq bound$ 
8  or  $g(s') \geq bound$ 
9  then continue
10  $cur\_g \leftarrow g(s) + TRANSITION\_COST(s, s')$ 
11 switch
12 case  $s' \notin (Open \cup Seen \cup Closed)$  :
13    $g(s') \leftarrow cur\_g, pred(s') \leftarrow s$ 
14   calculate  $h(s')$  and  $f'(s')$ 
15   insert  $s'$  into  $Open$ 
16 case  $s' \in Seen$  :
17   if  $cur\_g < g(s')$ 
18   then  $g(s') \leftarrow cur\_g, pred(s') \leftarrow s$ 
19   move  $s'$  from  $Seen$  to  $Open$ 
20 case  $cur\_g < g(s')$  :
21    $g(s') \leftarrow cur\_g, pred(s') \leftarrow s$ 
22   if  $s' \in Open$ 
23   then update  $s'$  in  $Open$ 
24   else move  $s'$  from  $Closed$  to  $Open$ 
25 if  $IS\_GOAL(s')$ 
26 then break
27 if new solution  $s^*$  was found (line 25)
28 then report  $s^*$ 
29    $bound = g(s^*)$ 
30    $w \leftarrow \max(1, w \times \phi)$ 
31    $Seen \leftarrow Seen \cup Open \cup Closed$ 
32 else return failure

```

Figure 2: Pseudo-code for the RWA* algorithm

type algorithm is remarkable because at first glance, it would seem unnecessary: A* keeps a queue of all generated but yet unexpanded nodes, and is thus not thought to suffer from premature commitment. Our insight is that weighted A* effectively makes such commitments due to its low- h bias.

Empirical Evaluation

We compare the performance of RWA* with several existing complete anytime approaches: Anytime A* (Hansen and Zhou 2007), ARA* (Likhachev *et al.* 2004), Anytime Window A* (Aine *et al.* 2007), and beam-stack search (Zhou and Hansen 2005). For beam-stack search, we implemented the regular version rather than the memory-conserving divide-and-conquer variant (memory-conserving techniques could be applied to all the algorithms here). All parameters of the competitor algorithms were carefully tuned for best performance. The algorithms based on WA* (RWA*, Anytime A* and ARA*) all share the same code base, as they differ only in few details; and they perform best for the same starting weights. For beam-stack search, we experimented with beam width values between 5 and 10,000 and plot the best of those runs. We also conducted experiments with iteratively broadening beam widths, but did not obtain better results than with the best fixed value. Beam-stack search expects an initial upper bound on the solution

cost. After consulting its authors, we chose a safe but large number as bound, as it is not clear how to come up with a reasonable upper bound for our experiments up-front.

In addition we compare against an alternative version of Anytime A*. For this algorithm, which originally does not decrease its weight between search phases, we experienced improved performance if we do decrease the weight. Thus, we also report results for a variant dubbed “Anytime A* WS”, (WS for weight schedule), where the weight is decreased as in RWA* or ARA*.

Classical PDDL Planning

Planning is a notoriously hard problem where optimal solutions can be prohibitively expensive to obtain even under very optimistic assumptions (Helmert and Röger 2008). To efficiently find a (typically suboptimal) solution, inadmissible heuristics are employed. In our experiment, we use the popular FF heuristic and all classical planning tasks from the International Planning Competitions between 1998 and 2006.

Experimental Setup We implemented each of the search algorithms within the Fast Downward planning framework (Helmert 2006) used by the winners of the satisficing track in the International Planning Competitions in 2004 and 2008. This framework originally uses greedy best-first search, which we replaced with the respective search algorithms. All other settings of Fast Downward were held fixed to isolate the effect of changing the search mechanism.

Fast Downward has two options built in that can enhance the search for a solution. Since our goal is to achieve best possible performance for planning, we make use of these search enhancements. The first is *delayed heuristic evaluation*, which means that states are stored in Open with their parent’s heuristic value rather than their own. This technique is helpful if the computation of the heuristic is very expensive compared to the other components of the search (like expansions of states). If many more states are generated than expanded, delayed evaluation leads to a great reduction in the number of times we need to calculate heuristic estimates, if at a loss in accuracy.

The second search enhancement in Fast Downward is the use of *preferred operators* (Helmert 2006). Preferred operators denote which successors of a state are deemed to lie on a promising path to the goal. They can be used as a second source of search information, complementing the heuristic. In particular, when run with the preferred-operator option, Fast Downward uses two Open lists, one with all operators and one with only preferred operators, and selects the next state from the two lists in an alternating fashion. For the A*-based algorithms, the preferred-operator mechanism can be used as defined and it improves results notably; it was thus incorporated. Beam-stack search is the one algorithm that does not use preferred operators as it is not obvious how to incorporate them.

For the WA*-based algorithms, we experimented with several starting weights and corresponding weight sequences, yet found that the relative performance did not




Figure 3: Anytime performance in PDDL planning.

change much. The weight sequence which gave overall best results and is used in the results below is 5, 3, 2, 1.5, 1. ARA* was adapted to use a weight schedule like RWA* and Anytime A* WS, since the original mechanism for decreasing weights in ARA* depends on the heuristic being admissible (ARA* originally reduces its weight by a very small amount whenever it proves that the incumbent solution is w -admissible for the current weight w). For beam-stack search, a beam width of 500 resulted in best anytime performance.

The time and memory limits were 30 minutes and 3 GB respectively for each task, running on a 3 GHz Intel Xeon X5450 CPU. In order to compare on meaningful benchmark tasks only, we select a subset of them as follows: Firstly, we exclude tasks that none of the algorithms could solve. Secondly, since we are interested in showing the improvement of solutions over time, we exclude tasks where all algorithms found only one solution (and hence no improvement occurred). Thirdly, we exclude “trivial” tasks where all algorithms find the same best solution within less than one second. This leaves 1096 of the original 1612 tasks.

Results Fig. 3 summarises the results. We plot average normalised quality scores following the definition used in the most recent planning competition (Helmert *et al.* 2008): Q^*/Q , where Q is the cost of the algorithm’s current solution and Q^* is cost of the overall best solution found by any of the algorithms. Hence, the score ranges between 0 (task not solved yet) and 1 (A has found the best solution quality Q^*). These scores are then averaged over all tasks, and error bars in the plots give 95% confidence intervals on the mean. It is intentional that unsolved tasks negatively influence the score of an algorithm. The alternative would be to compare only on the tasks solved by all algorithms. However, that would give an advantage to methods that are less greedy, i. e., those that solve fewer problems but find solutions of better quality. Hence that would not reward good anytime performance (for example, breadth-first search would appear superior).

Legends in our plots are sorted in decreasing order of final score. As the left and centre panels of Fig. 3 show, RWA* performs best, outperforming the other algorithms by a substantial margin. In the left panel, we compare against the other WA*-based algorithms. The weight-decreasing variant of Anytime A* (Anytime A* WS) is slightly bet-

ter than ARA*, whereas the original Anytime A* algorithm performs worse. It is notable that while all four WA*-based algorithms are very closely related, differing only in a few lines of code, the simple addition of the restarts leads to a significant improvement in performance. “Base WA*” depicts a WA* search that stops after finding the first solution and is provided for reference.

Although the confidence intervals for all four anytime algorithms overlap to some extent in the left-most panel of figure 3 a paired student’s t-test reveals that the difference between each pair is statistically significant, having p-values ranging between 10^{-5} in the most discernible case and 10^{-3} in the least discernible case.

The centre panel shows the results for Window A* and beam-stack search. Beam-stack search and Window A* perform significantly worse than the WA*-based algorithms. This is mainly due to solving fewer problems, as a significant determiner in gross performance is the number of problems solved by a certain time. By the time cut-off, all WA*-based algorithms solve (essentially the same) 84% of the tasks, while beam-stack search and Window A* solve only 78% and 63%, respectively. 3% (3%) of the tasks are solved by beam-stack search (Window A*) but not by the WA* algorithms. The opposite is true for 9% (24%) of the tasks.

The right panel of Fig. 3 shows for the best version of each algorithm how often it achieved (at least) a certain solution quality. As can be seen, RWA* achieves high-quality solutions significantly more often than the other algorithms. The best quality is achieved by RWA* in 65% of the tasks, respectively, while all others achieve it in $\leq 51\%$ of the tasks. The average number of solutions found was 3 for the weight-decreasing WA* methods and beam-stack search, 5 for non-decreasing Anytime A* and 1 for Window A*.

When examining the distribution of RWA*’s performance over the different domains, we found that RWA* outperforms the other WA* algorithms in 40% of the domains, while being on par in the remaining 60%. In no domain did RWA* perform notably worse than any of the other WA* methods. Beam-stack search and Window A* show different strengths and weaknesses compared to the WA*-based algorithms. In the Optical Telegraph domain, beam-stack search performs well whereas all other algorithms perform badly. Beam-stack search and Window A* perform worst

	1st solution				2nd solution				3rd (final) solution			
	<i>len</i>	<i>c.i.</i>	<i>exp</i>	<i>t</i>	<i>len</i>	<i>c.i.</i>	<i>exp</i>	<i>t</i>	<i>len</i>	<i>c.i.</i>	<i>exp</i>	<i>t</i>
RWA*	165	-	1142	0.07	163	153	2370	0.08	125	1	16919	0.84
ARA*	165	-	1142	0.06	163	153	1220	0.07	161	145	5743	0.28
Anytime A* WS	165	-	1142	0.07	163	153	1231	0.08	161	145	5629	13.25
Anytime A*	165	-	1142	0.07	163	153	1255	0.07	161	145	206480	0.28

Table 1: Solution sequences for the largest task in the gripper domain. The plan length is denoted by *len*. The first step in which a solution deviates from the previous is denoted by *c.i.* (change index); *exp* is the number of states expanded until the solution is found, and *t* the runtime in seconds.

on large problems, e. g. in the Schedule and Logistics 1998 domains.

To summarise, we find that the WA*-based methods perform significantly better than Window A* and beam-stack search. We believe that this is due to the *global* WA*-based algorithms being better able to deal with inadmissible and locally highly varying heuristic values, whereas Window A* and beam-stack search commit early to certain areas of the search space. Exhausting a state-space area is particularly difficult here as inadmissible heuristic values cannot be used for pruning. Furthermore, beam-stack search cannot make use of the planning-specific preferred operators enhancement.

We also performed experiments without preferred operators and delayed evaluation for all algorithms. There, the performance of all algorithms becomes drastically worse (e. g., the WA*-based methods leave twice as many problems unsolved and in addition show worse anytime performance on the solved problems). RWA* still performs better than all other algorithms, though by a much smaller margin than with the search enhancements. With fewer solved tasks and fewer solutions per task, RWA* cannot improve on the other WA* algorithms as much. Also, while the search enhancements are very helpful in finding a solution, they make the search slightly less informed and more greedy, contributing to low-*h* bias and thus making restarts more effective.

We furthermore conducted experiments with other planning heuristics (not shown), including the recent landmark heuristic (Richter *et al.* 2008) and the context-enhanced additive heuristic (Helmert and Geffner 2008). Compared to the experiments described above, which use the FF heuristic, the performance of all algorithms was worse with these other heuristics. However, the *relative* performance of the algorithms was similar, with RWA* outperforming the other methods in all cases, though by smaller margins.

A Detailed Example We argue that RWA* shows such good results because restarts encourage changes in the beginning of a plan rather than the end. The largest task in the gripper domain provides an illustration. Gripper tasks consist in transporting balls (here 42) from one room to another with the help of a two-armed robot. The initial WA* search phase finds a plan in which all balls are transported separately, whereas in the optimal plan the robot always carries two balls at a time. Window A* does not solve this problem, while beam-stack search finds only one solution (the optimal) after 2.5 seconds. The WA* algorithms find

three improving solutions, with the first solution found after less than one second, see Table 1. All these plans are found very quickly, in less than 15 seconds. However, the last plan found by RWA* is optimal, whereas the other WA* algorithms do not find any improved solutions during the remaining 29 minutes.

The noteworthy aspect in this example is the *indices of change* for the plans, i. e. the steps in which subsequent plans first differ from their predecessors. For ARA* and the two Anytime A* variants these change indices are fairly high (≥ 145), denoting the fact that their subsequent plans only differ in the last 20 actions from the first plan found. For RWA*, the third solution has a change index of 1, i. e. it differs in the first action from the previous plans. This suggests that the big jump in solution quality for RWA* indeed results from further exploration near the start state, whereas the other algorithms unsuccessfully spend their effort in deeper areas of the search space.

Other Benchmarks

We also ran experiments in two other benchmark domains: the robotic arm domain and the sliding-tile puzzle, which we chose because of their previous use as anytime benchmarks (robotic arm), or their standing as traditional benchmarks in the search literature (sliding-tile puzzle).

Robot Motion Planning This domain concerns motion planning for a simulated 20-degree-of-freedom robotic arm (Likhachev *et al.* 2004). The base of the arm is fixed, and the task is to move its end-effector to the goal while navigating around obstacles. An action is defined as a change of the global angle at a particular joint. The workspace is discretised by overlaying it with a grid of 50×50 cells, and the heuristic of a state is the distance from the current location of the robot’s end-effector to the goal, taking obstacles into account. The size of the state space is over 10^{26} , and in most instances it is infeasible to find an optimal solution.

As before, we eliminate trivial instances, leaving 17 of the 22 tasks kindly made available to us by Maxim Likhachev. For ARA*, we use the settings suggested by Likhachev *et al.*, starting with a weight of 3 and decreasing it by 0.02 each time; the other WA*-based algorithms also start with weight 3. Note that because the heuristic is admissible, ARA* can prove suboptimality bounds and make informed decisions on when to reduce its weight (namely, whenever a new bound is proven). For RWA*, on the other hand, a decrease in weight incurs a substantial overhead each time due

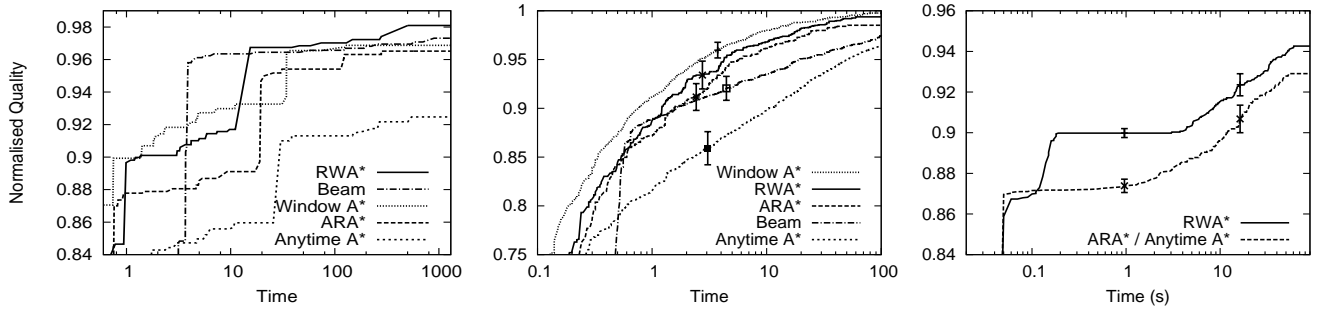


Figure 4: Anytime performance for robot motion planning (left), the 15-puzzle (middle), and an artificial search space (right).

to the restart. This is why larger and less frequent weight decreases work better for it. We use a factor of 0.84 to decay the weight for RWA*, resulting in a similar weight schedule as in the planning experiment, reducing the weight whenever an improved solution is found.

The results are shown in the left panel of Fig. 4. RWA* outperforms ARA* and Anytime A* for all timeouts above 1 second and shows overall very good anytime performance, including reaching the best final quality. Window A* does well in this domain and shows best results in the early stages of search. Anytime A* performs notably worse than ARA* here, though we found that its weight-decreasing variant (not shown) achieves similar performance as ARA*. Beam-stack search, here with a beam width of 100, takes longer than most of the other methods to achieve high performance, but then comes up steeply and achieves almost the same final quality as RWA*. (Error bars are not shown here due to the high variance across these hand-designed instances.)

Sliding-Tile Puzzle In the sliding-tile puzzle domain, we tested on the 100 15-puzzle instances from Korf (1985) using the Manhattan distance heuristic. The starting weight for the WA* algorithms was 3, the weight sequence for RWA* being 3, 2, 1.5, 1.25, 1 while ARA* uses small decrements of 0.2. The beam width in beam-stack search was 500.

The results are shown in the middle panel of Fig. 4. There are comparatively few different solutions, leading to similar performance of all WA*-based algorithms that decrease weight. Window A* performs best, with RWA* second and ARA* third. With smaller beam widths, the score of beam-stack search rises faster in the beginning but it takes longer to achieve perfect quality; with larger beam widths the reverse holds.

Summary While in the robot motion domain and the sliding-tile puzzle RWA* does not dominate its competitors as notably as in the PDDL planning experiment, we note that its performance is continually very good. Compared to its most similar competitors, the other WA* algorithms, RWA* is always better or on par. This is also the case in a third domain (gridworld pathplanning), as we demonstrated previously (Richter *et al.* 2009). This may suggest that even in cases where the restarts are not helpful, they do little harm in

terms of computation time. The methods that are not based on WA* (Window A* and beam-stack search) perform well in some domains but badly in others. By contrast, RWA* shows robustly good performance over all domains.

An Artificial Search Space

To get a clearer picture of when restarts are helpful, we conducted experiments on a manually-designed search tree. We fixed a branching factor of 25 and a typical solution depth of 500. Nodes are characterised by their *approximate goal distance* (*agd*). The root has an *agd*-value of 500 and nodes with a value of 0 are labelled as goals. Edge costs are chosen randomly (with uniform probabilities) between 0 and 10, and for a given edge of cost c the *agd*-value of the reached child varies randomly by up to c from the parent’s value. This is analogous to the costs of walking in a gridworld, where a move of c steps may take an agent c steps closer to the goal, c steps further away or anything in between, depending on the direction the agent walks.

The heuristic values were chosen to enforce systematic biases. The heuristic underestimates the *agd*-values of nodes by up to a certain percentage. The h -value of the root is chosen randomly with this constraint, and the resulting error factor (the ratio of h -value to *agd*-value) is recorded. For all other nodes, their h -values are correlated with their parent’s h -values such that a child’s h -value differs by no more than 1 from the value that would result when using the same error factor as the parent. This is achieved by first computing the “parent-induced h -value” that would result if a node had the same error factor as its parent. Then we determine a random value for the child within the constraint of heuristic accuracy. The actual h -value for the child is obtained by moving the parent-induced value by 1 into the direction of the random value.

Averaged results for 100 runs with different random seeds for a heuristic that underestimates by up to 20% are shown in the right panel of Fig. 4, where RWA* and ARA* were run with a weight schedule of 2, 1.5, 1.25, 1.125, 1. The relative results are similar for heuristics of different accuracy. Note that the weight-decreasing variant of Anytime A* is equivalent to ARA* here, as there are no cycles in the search space. RWA* has a clear advantage over the other algorithms, and we found the same effect as in planning, that the change in-

dices of solutions for RWA* are substantially lower than for the other algorithms.

When weakening the correlation between the h -values of parents and children, the advantage of RWA* diminishes. For example, if all h -values are chosen completely at random, RWA* has no advantage over ARA*/Anytime A* unless a very accurate heuristic is used. This is due to the fact that without systematic errors in the heuristic, all algorithms explore more nodes in the top of the search tree, rather than committing quickly to a bad path. This is in line with observations by Zahavi et al. (2007), who note that inconsistent heuristics can be beneficial because they make it less likely that the search gets stuck in a region of bad heuristic estimates. In that case, the greedy search makes fewer early mistakes and restarts consequently provide less benefit.

Additionally, we found that RWA* gains more advantage if the heuristic underestimates nodes near the root by a stronger factor than nodes that are lower in the tree. In none of the scenarios we looked at did RWA* perform notably worse than the other algorithms.

Discussion

As we saw in the example tasks from gridworld and the gripper domain, restarting anytime search overcomes the low- h bias of greedy anytime search that tends to expand nodes near a known goal. The desire to revisit nodes near the start state stems from the fact that a heuristic evaluation function is often less informed near the root of the search tree than near a goal and that early mistakes can be important to correct (Harvey and Ginsberg 1995; Furcy and Koenig 2005). For example, in problems with multiple goal states the heuristic may misjudge the relative distance of the goals and lead a greedy search into a sub-optimal part of the search space. Similarly, domains with inaccurate heuristics that have systematic errors can lead to early mistakes and thus benefit from restarts.

Conclusion

We have demonstrated an important dysfunction in conventional approaches to anytime search: by trying to re-use previous search effort from a greedy search, traditional anytime methods suffer from low- h bias and tend to improve the incumbent solution starting from the end of the solution path. As we showed in this paper, this can be a poor strategy when the heuristic makes early mistakes. The counter-intuitive approach of clearing the Open list, as in the Restarting Weighted A* algorithm, can lead to much better performance in such domains while doing little harm in others.

Acknowledgements

We thank Malte Helmert, Charles Gretton, and Patrik Haslum for helpful input and Maxim Likhachev for making his code available. NICTA is funded by the Australian Government, as represented by the Department of Broadband, Communications and the Digital Economy, and the Australian Research Council, through the ICT Centre of Excellence program. We also gratefully acknowledge support from NSF grant IIS-0812141 and the DARPA CSSG program.

References

- Sandip Aine, P. P. Chakrabarti, and Rajeev Kumar. AWA* – A window constrained anytime heuristic search algorithm. In *Proc. IJCAI 2007*, pages 2250–2255, 2007.
- Rina Dechter and Judea Pearl. The optimality of A*. In Laveen Kanal and Vipin Kumar, editors, *Search in Artificial Intelligence*, pages 166–199. Springer-Verlag, 1988.
- David Furcy and Sven Koenig. Limited discrepancy beam search. In *Proc. IJCAI 2005*, pages 125–131, 2005.
- Carla P. Gomes, Bart Selman, and Henry A. Kautz. Boosting combinatorial search through randomization. In *Proc. AAAI 1998*, pages 431–437, 1998.
- Eric A. Hansen and Rong Zhou. Anytime heuristic search. *JAIR*, 28:267–297, 2007.
- Eric A. Hansen, Shlomo Zilberstein, and Victor A. Danilchenko. Anytime heuristic search: First results. Technical Report CMP-SCI 97-50, University of Massachusetts, Amherst, September 1997.
- Peter E. Hart, Nils J. Nilsson, and Bertram Raphael. A formal basis for the heuristic determination of minimum cost paths. *IEEE Transactions on Systems Science and Cybernetics*, 4(2):100–107, 1968.
- William D. Harvey and Matthew L. Ginsberg. Limited discrepancy search. In *Proc. IJCAI 1995*, pages 607–615, 1995.
- Malte Helmert and Héctor Geffner. Unifying the causal graph and additive heuristics. In *Proc. ICAPS 2008*, pages 140–147, 2008.
- Malte Helmert and Gabriele Röger. How good is almost perfect? In *Proc. AAAI 2008*, pages 944–949, 2008.
- Malte Helmert, Minh Do, and Ioannis Refanidis. IPC 2008, deterministic part. Web site, <http://ipc.informatik.uni-freiburg.de>, 2008.
- Malte Helmert. The Fast Downward planning system. *JAIR*, 26:191–246, 2006.
- Richard E. Korf. Depth-first iterative-deepening: An optimal admissible tree search. *AIJ*, 27(1):97–109, 1985.
- Maxim Likhachev, Geoffrey J. Gordon, and Sebastian Thrun. ARA*: Anytime A* with provable bounds on sub-optimality. In *Proc. NIPS-03*, 2004.
- Maxim Likhachev, Dave Ferguson, Geoffrey J. Gordon, Anthony Stentz, and Sebastian Thrun. Anytime search in dynamic graphs. *AIJ*, 172(14):1613–1643, 2008.
- Ira Pohl. Heuristic search viewed as path finding in a graph. *AIJ*, 1:193–204, 1970.
- Silvia Richter, Malte Helmert, and Matthias Westphal. Landmarks revisited. In *Proc. AAAI 2008*, pages 975–982, 2008.
- Silvia Richter, Jordan T. Thayer, and Wheeler Ruml. The joy of forgetting: Faster anytime search via restarting. In *International Symposium on Combinatorial Search (SoCS 2009)*, 2009.
- Bart Selman, Hector J. Levesque, and David G. Mitchell. A new method for solving hard satisfiability problems. In *Proc. AAAI-92*, pages 440–446, 1992.
- Uzi Zahavi, Ariel Felner, Jonathan Schaeffer, and Nathan R. Sturtevant. Inconsistent heuristics. In *Proc. AAAI 2007*, pages 1211–1216, 2007.
- Weixiong Zhang. Complete anytime beam search. In *Proc. AAAI 1998*, pages 425–430, 1998.
- Rong Zhou and Eric A. Hansen. Beam-stack search: Integrating backtracking with beam search. In *Proc. ICAPS 2005*, pages 90–98, 2005.