

Rectangle Search: An Anytime Beam Search

Sofia Lemons^{1,2}, Wheeler Ruml¹, Rob Holte³, Carlos Linares López⁴

¹ University of New Hampshire

² Earlham College

³ University of Alberta, Alberta Machine Intelligence Institute (Amii)

⁴ Computer Science and Engineering Department, Universidad Carlos III de Madrid
sofia@snlemons.com, ruml@cs.unh.edu, rholte@ualberta.ca, carlos.linares@uc3m.es

Abstract

Anytime heuristic search algorithms try to find a (potentially suboptimal) solution as quickly as possible and then work to find better and better solutions until an optimal solution is obtained or time is exhausted. The most widely-known anytime search algorithms are based on best-first search. In this paper, we propose a new algorithm, rectangle search, that is instead based on beam search, a variant of breadth-first search. It repeatedly explores alternatives at all depth levels and is thus best-suited to problems featuring deep local minima. Experiments using a variety of popular search benchmarks suggest that rectangle search is competitive with fixed-width beam search and often performs better than the previous best anytime search algorithms.

Introduction

It is often convenient to have a heuristic search algorithm that can flexibly make use of however much time is available. The search can be terminated whenever desired and returns the best solution found so far. Dean and Boddy (1988) termed these *anytime* algorithms. Russell and Zilberstein (1991) further differentiated between *interruptible* algorithms, which quickly find a solution and then find better solutions as time passes, eventually finding an optimal plan if given sufficient time, and *contract* algorithms, which are informed of the termination time in advance and thus need only find a single solution before that time. Anytime algorithms have been proposed as a useful tool for building intelligent systems (Zilberstein 1996; Zilberstein and Russell 1996). While only a few contract search algorithms have been proposed (Dionne, Thayer, and Ruml 2011), interruptible algorithms have been widely investigated and applied (Likhachev and Ferguson 2009).

As we review below, the most well-known interruptible anytime heuristic search algorithms are based on best-first search. Best-first search is attractive as it is the basis for the optimally-efficient optimal search algorithm A* (Hart, Nilsson, and Raphael 1968) and it is relatively well understood. However, because anytime algorithms are intended for use cases in which the solutions found do not need to be proven optimal, and are not even expected to be optimal, it is not

obvious that best-first search is the most appropriate choice of algorithmic architecture.

In this paper, we propose an interruptible algorithm called rectangle search that is inspired by beam search, which is based on breadth-first search rather than best-first search. Rectangle search is straightforward and simple to implement: the beam incrementally widens and deepens. We study rectangle search’s performance experimentally on seven popular heuristic search benchmarks. We find that, overall, rectangle search outperforms previously-proposed anytime search algorithms. Furthermore, it tends to find solutions of comparable cost at similar times when compared to fixed-width beam search, with the added benefit of reducing solution cost over time. We investigate when it performs well or poorly, finding that it performs well in problems with large local minima. Overall, rectangle search appears to mark a new state of the art for anytime search, while also serving as a replacement for fixed-width beam search.

Background

Before presenting rectangle search, we first review relevant prior work in anytime search and beam search.

Anytime Search

Most previous anytime algorithms are based on weighted A* (Pohl 1973), which is a best-first search using $f'(n) = g(n) + w \times h(n)$, where $g(n)$ is the cost to reach n from the start, $h(n)$ is the estimated cost-to-go to the goal from n , and $w \geq 1$. For example, anytime weighted A* (AWA*) (Hansen and Zhou 2007) runs weighted A* but retains a current incumbent solution and continues searching for better solutions until there are no open nodes with $f(n) = g(n) + h(n) < g(\text{incumbent})$, thus proving that the incumbent is optimal.

Anytime Repairing A* (ARA*) (Likhachev, Gordon, and Thrun 2004) is similar, but applies a schedule of decreasing weights ending with $w = 1$. When a solution is found, it decreases the weight according to the schedule and reorders the open list. It terminates after finding a solution with $w = 1$ or after exhausting all nodes with $f(n) < g(\text{incumbent})$.

Anytime EES (AEES) (Thayer, Benton, and Helmert 2012) requires no weighting parameter and explicitly works to minimize the time between finding new solutions by using distance-to-go estimates $d(n)$. $d(n)$ is an estimate of the




Figure 1: The exploration of rectangle search.

number of actions along a minimum-cost path from node n to a goal. Often this can be computed while calculating $h(n)$. AEES maintains an open list ordered on an error-adjusted evaluation function $\hat{f}(n)$, a focal list ordered on an error-adjusted distance-to-go measurement $\hat{d}(n)$, and a cleanup list ordered on $f(n)$. It compares the current incumbent solution's cost to the lowest f -value among open nodes to determine a bound w on solution suboptimality. The focal list maintains only nodes n with $\hat{f}(n) < w \cdot \hat{f}(b)$ where b is the lowest \hat{f} -valued node, because these are predicted to lead to a solution that is better than the current incumbent. AEES was demonstrated to outperform other anytime algorithms including ARA*, Anytime Nonparametric A* (ANA*) (Van Den Berg et al. 2011), and Restarting Weighted A* (RWA*) (Richter, Thayer, and Ruml 2010).

Beam Search

Beam search (Bisiani 1987) is an incomplete and suboptimal variant of breadth-first search. It expands a fixed maximum number of nodes at each depth level of the search, referred to as the beam width b . All nodes from the beam at the current level are expanded and then the best b among those descendants are selected for the next level's beam. Only nodes at the same depth in the tree are compared. It continues to search until either a solution is found or until no new states are reachable from the current depth. A variant of beam search using $d(n)$ for node selection called bead outperforms beam search using $f(n)$ or $h(n)$ in non-unit cost domains (Lemons et al. 2022).

One issue with beam search is that when the beam width b is increased, sometimes a higher cost solution is returned. The algorithm monobead (Lemons et al. 2022) addresses this issue. It regards the beam as an ordered sequence of numbered *slots*. Given nodes on the beam at current depth d , to fill beam slot i at the next depth level, the node at slot i of depth d 's beam is expanded, its children are added to a priority queue for depth $d + 1$, and the best node from that queue is selected. This iterates for values of i from 1 to b , with the priority queue retaining any children that were not selected to fill previous slots. The node selected for slot i at depth $d + 1$ is thus restricted to be a child of a node in slots 1 through i at depth d . This careful selection order prevents children of nodes at later slots from supplanting children of earlier slots and preserves any solutions that would have been found by searches with a narrower beam width.

Rectangle Search

One way to view rectangle search is as an iteratively widening and deepening monobead search that obviates the need

Algorithm 1: Pseudocode for rectangle search.

```

1 begin
2    $openlists \leftarrow [\emptyset]$ 
3    $closed \leftarrow \emptyset$ 
4    $incumbent \leftarrow \text{node with } g = \infty$ 
5   expand start and add children to  $openlists[0]$ 
6    $depth \leftarrow 1$ 
7   while non-empty lists exist in  $openlists$  do
8     for  $i = 0 \dots \text{length}(openlists) - 2$  do
9       | select & expand from  $openlists[i]$ 
10      extend  $openlists$  with  $\emptyset$  aspect times
11      for  $j = i + 1 \dots \text{length}(openlists) - 2$  do
12        | for  $k = 1 \dots depth$  do
13          | | select & expand from  $openlists[j]$ 
14         $depth \leftarrow depth + aspect$ 
15      trim empty lists from start & end of  $openlists$ 
16   return  $incumbent$ 

```

Algorithm 2: Node selection & expansion.

```

Data: closed, openlists, i, incumbent
17 begin
18   do
19     |  $n \leftarrow \text{remove first node from } openlists[i]$ 
20     while  $f(n) \geq g(incumbent)$  ;
21     add  $n$  to  $closed$ 
22      $children \leftarrow \text{expand}(n)$ 
23     for each child in children do
24       | if  $f(child) < g(incumbent)$  then
25         | if child is a goal then
26           | |  $incumbent \leftarrow child$ 
27           | | report new incumbent
28         | else
29           | |  $dup \leftarrow \text{child's entry in } closed$ 
30           | | if child not in  $closed$  or
31           | |  $g(child) < g(dup)$  then
32           | | | add child to  $openlists[i + 1]$ 

```

to pre-specify b . In each iteration i it is allowed to expand one new node at each previously explored depth level and to expand i nodes at a new depth level. This last step makes the number of expansions at each allowed depth level equal, resulting in a rectangular shape of the state space being explored, as shown in Figure 1. By default, the algorithm explores only one additional depth level at each iteration (resulting in a square), but this number of new levels can be adjusted by setting an *aspect* parameter. We use $\text{rectangle}(x)$ to refer to rectangle search with $aspect = x$.

Figure 1 demonstrates three iterations of $\text{rectangle}(1)$. In the first iteration, node A is selected from the front of $open1$ (the open list for nodes at depth 1) and expanded, with its children being inserted into $open2$. The second iteration begins with node B being selected from $open1$ and expanded and its children also inserted into $open2$. The algorithm is now allowed to expand two nodes at a new depth level, depth 2. At this level, nodes D and E are selected from the front of

open2 and expanded, with their children being inserted into *open3*. The third iteration consisted of node C being selected at depth 1, node F being selected at depth 2, and nodes G, H, and I being selected at depth 3. If no solution has been found yet, the children of the nodes at depth 3 would be inserted into a new open list for depth level 4 (not shown).

Algorithms 1 and 2 show the design of rectangle search. It creates an ordered collection of priority queues, called *openlists*, which initially contains only a single entry for storing nodes at depth level 1 (line 2). The main loop at line 7 continues so long as there are open nodes at any depth level. For all depth levels except the deepest, a single node is selected and expanded, with its children being inserted into the next depth level's open list (line 31). Because of this structure, there will not have been any nodes previously expanded from the deepest depth level represented in *openlists*. New open lists are added to represent deeper levels (line 10), and then up to *depth* many nodes will be selected from all but the deepest depth level now represented in *openlists*. For *rectangle(1)*, this will result in *depth* many nodes being expanded at a single new depth level at each iteration of the main loop, and a single new priority queue being added to *openlists* to hold the nodes at the new depth level to be explored in the next iteration. While rectangle search could order its priority queues on any criterion, we use $d(n)$ to encourage finding a first solution quickly.

Properties of Rectangle Search

Theorem 1. *With an admissible heuristic, rectangle search is complete.*

Proof. Rectangle search begins from the start state and terminates when no states remain at any depth level (line 7). Because it begins with an incumbent cost value of ∞ , it will not prune any non-duplicate nodes until a solution is found (lines 20 and 24). It will search the entire reachable state space if necessary to find a solution. \square

Theorem 2. *With an admissible heuristic, the last solution returned by rectangle search is optimal.*

Proof. Rectangle search prunes only nodes with f -values greater than or equal to the current incumbent (lines 20 and 24), which cannot lead to a better solution. It terminates when no nodes remain at any depth (line 7). Therefore, it expands all nodes with f -value less than or equal to the optimal solution and the last solution it returns will be optimal. \square

Rectangle search is related to monobead search because both algorithms select nodes for expansion at each level in a well-defined order, with respect to beam slots. Rectangle does not hold this strict ordering when expanding nodes at a new depth level for the first time. It also performs duplicate elimination without regard to the slot at which a node was seen. However, once a depth level has had nodes selected and expanded, only a single node will be selected and expanded at a time for that level, preserving a monobead-like order from that point onward. Therefore, the solutions found by rectangle may not follow the monotonic ordering ensured by monobeam. To highlight the similarity, we define

a variant of rectangle search called strict rectangle that expands nodes at the deepest current depth level one at a time, based on their slot index in the beam, and selects nodes to fill those slots at the next depth level immediately after the current node in that slot is expanded. Strict rectangle also only eliminates duplicate nodes if the previously seen version was expanded from the same slot as the duplicate or earlier. With these restrictions, the node selection for a specific slot and depth are exactly the same for monobead and strict rectangle search.

To estimate the overhead of rectangle search in relation to a fixed-width beam search, we analyze the number of nodes expanded by strict rectangle search relative to the number of nodes expanded by monobead search. For simplicity, we assume strict rectangle with *aspect* = 1. Let the first solution found by monobead be generated when it expands a node x in slot s at depth d . To provide an upper bound on the extra work done by strict rectangle to find this solution, we assume the beam width for monobead is equal to s , no other solutions exist, and strict rectangle never runs out of nodes to fill its beam. If $s = d$, strict rectangle will search exactly the same region of the search space as monobead and the number of nodes expanded by each will be equal. If $s > d$, then strict rectangle will expand x during iteration s and will have searched $s - d - 1$ superfluous depth levels (x will be expanded before any node at depth s is expanded), each with $s - 1$ nodes expanded. And if $s < d$, strict rectangle will expand x during iteration d and will have expanded $d - s$ superfluous nodes at depth levels 1 to $d - 1$. This yields:

$$\text{overhead} \leq \begin{cases} 0 & s = d \\ (s - d - 1)(s - 1) & s > d \\ (d - 1)(d - s) & s < d \end{cases}$$

In summary, the overhead of strict rectangle over monobead is at most quadratic in whichever of d or s is greater.

As we will see in our experimental evaluation, *rectangle(1)* is not well-suited for domains in which the heuristic is very accurate and solutions are very deep (e.g., $d \gg s$). Rectangle search must do $O(d^2)$ work to reach depth d , whereas a best-first search might only need to do $O(d)$ if the heuristic is accurate. In this sense, *rectangle(1)* errs too much on the side of exploration in such domains. Using a larger *aspect* may help to reach deeper regions of the search space earlier, but will still require some extraneous exploration at higher levels.

On the other hand, rectangle search is not obliged to expand open nodes in order of their heuristic evaluation value (be it f , h , or d). A *local minimum* (or *crater*) is a strongly connected set of states with deceptive heuristic values that are lower than the highest heuristic value necessary to encounter along a path from them to a goal (Wilt and Ruml 2014; Heusner 2019). Intuitively, they have low heuristic values but are not actually close to a goal and expanding these nodes does not necessarily represent progress. Unlike many best-first searches, rectangle search does not need to expand all nodes in a local minimum before expanding a node with a higher heuristic value. Furthermore, we conjecture that rectangle's similarity to monotonic beam search

may aid in retaining diverse nodes during the search, preventing a large local minimum from displacing all other nodes in the beam and dominating the search.

Although it is based on breadth-first search, rectangle search’s behavior has similarities to depth-first approaches. When run with a large *aspect* value (e.g., 500 is evaluated below), rectangle search’s behavior is similar to limited-discrepancy search (Korf 1996). First, it repeatedly expands the best-looking child of the previously expanded node, forming a deep but thin probe into the state space. This is similar to hill-climbing (Hoffmann and Nebel 2001) and depth-first search. If the heuristic is accurate or goals are plentiful, this will quickly find a goal. When a goal is found, its cost establishes an f pruning bound that limits the depth of subsequent search and forces the algorithm to widen more rapidly. Rectangle search widens by expanding nodes at all previous depth levels, allowing it, for example, to escape a deep local minimum in d without expanding all nodes in the minimum. This is similar in spirit to limited-discrepancy search, a depth-first-search-based method in which alternatives to the child preferred by a node ordering heuristic are explored at every level of a tree, except that, by virtue of its *openlists*, rectangle search has the flexibility to explore alternatives anywhere in a depth layer.

Experimental Evaluation

To better understand rectangle search’s behavior, we implemented it and other algorithms in C++¹ and compared their anytime behavior on seven classic search benchmarks (and for most, multiple variations in cost or action model). All algorithms were given a 7.5GB memory limit and a 5 minute time limit on a 2.6 GHz Intel Xeon E5-2630v3.

Relevant Comparators

In addition to best-first algorithms like ARA* and AEES, we examined depth-first algorithms that can give anytime results. DFS* (Vempaty, Kumar, and Korf 1991) is an anytime algorithm based on iterative deepening and depth-first branch-and-bound. It performs a depth-first search with an increasing cost bound (beginning with $h(start)$ and doubling at each iteration) until it finds an initial solution. It then performs a branch-and-bound search with the incumbent solution’s cost as the new bound. We implemented DFS* both with and without child ordering.

We also implemented a similar variation of Improved Limited Discrepancy Search (ILDS) (Korf 1996) which we call ILDS*. Because ILDS requires a depth bound, ILDS* begins with a depth bound equal to $d(start)$ and performs ILDS with an increasing number of discrepancies until all nodes within the depth bound have been explored. When a goal is found, it is retained as an incumbent. The search continues until no node with $f(n) \leq g(incumbent)$ is pruned based on the depth bound or discrepancy bound.

Another anytime algorithm with similarities to rectangle search is Complete Anytime Beam Search (CABS) (Zhang 1998). As adapted by Libralesso et al. (2020), CABS performs a sequence of increasing-width beam searches, dou-

bling the beam width at each iteration, and retaining the best solution so far as an incumbent. It terminates when no node excluded from the beam has $f(n) \leq g(incumbent)$. While CABS has been primarily used in depth-bounded problems, many of our domains do not have fixed depths and therefore our implementation of CABS uses a closed list.

To keep our plots clear, we only include the competitive algorithms: ARA*, AEES, CABS, and rectangle search. Plots for all algorithms in all domains are given in Lemons et al. (2023). We ran ARA* in several previously-proposed configurations: initial weights of 10 and 2.5 with a decrement of 0.02 (Likhachev, Gordon, and Thrun 2004) and a weight schedule of 5, 3, 2, 1.5, 1 (Thayer, Benton, and Helmert 2012). Rectangle search was tested with *aspect* values of 1 and 500 (selected through preliminary tests).

The Sliding Tile Puzzle

Six different cost models of sliding tiles were used: unit cost; heavy cost, where moving tile number t costs t ; sqrt cost, moving t costs \sqrt{t} ; inverse cost, $1/t$; reverse cost, $\#spaces - t$, and reverse inverse cost, $1/(\#spaces - t)$. We tested both the 15-puzzle (4x4) and the 24-puzzle (5x5), using a cost-weighted Manhattan distance heuristic.

Selected results are shown in Figures 2a and 2b (the legend is in 2e). We plot anytime performance over time, including before all algorithms have solved all instances. The primary metric we use is average *quality*: the cost of the best known solution (optimal if known, otherwise the best solution cost given by any of the algorithms at any time), divided by the cost of the incumbent solution (or ∞ if no solution has been found yet.) Quality can range from 0 for unsolved instances to a maximum of 1. A dot is included to show when an algorithm has solved all instances. Before this dot, an increase in average quality could be due to either a new instance being solved or a better solution being found for an instance that was already solved.

In both the 15- and 24-puzzle, and across all cost models, rectangle(1) achieves full coverage before other algorithms and provides better solution cost from that point onward. More specifically, both configurations of rectangle reach full coverage first, with rectangle(1) having lower solution costs at that point and rectangle(500) quickly catching up. In particular, in the inverse cost model rectangle(1), rectangle(500), and AEES are the only algorithms to solve all instances within our time bound.

Blocks World

We tested on 100 random blocks world instances with 20 blocks. We included two different action models: ‘blocks world’, where blocks are directly moved to a stack as an action, and ‘deep blocks world’ (Lelis, Zilles, and Holte 2013), where picking up and putting down blocks each use an action, leading to longer solutions. The heuristic used was the number of blocks out of place (any block which is not in a sequence of blocks from the table up which matches the goal state). This heuristic value is doubled for deep blocks. For both variants, rectangle search reaches full coverage earliest and provides the best solution costs after.

¹Code available at <https://github.com/snlemons/search>.

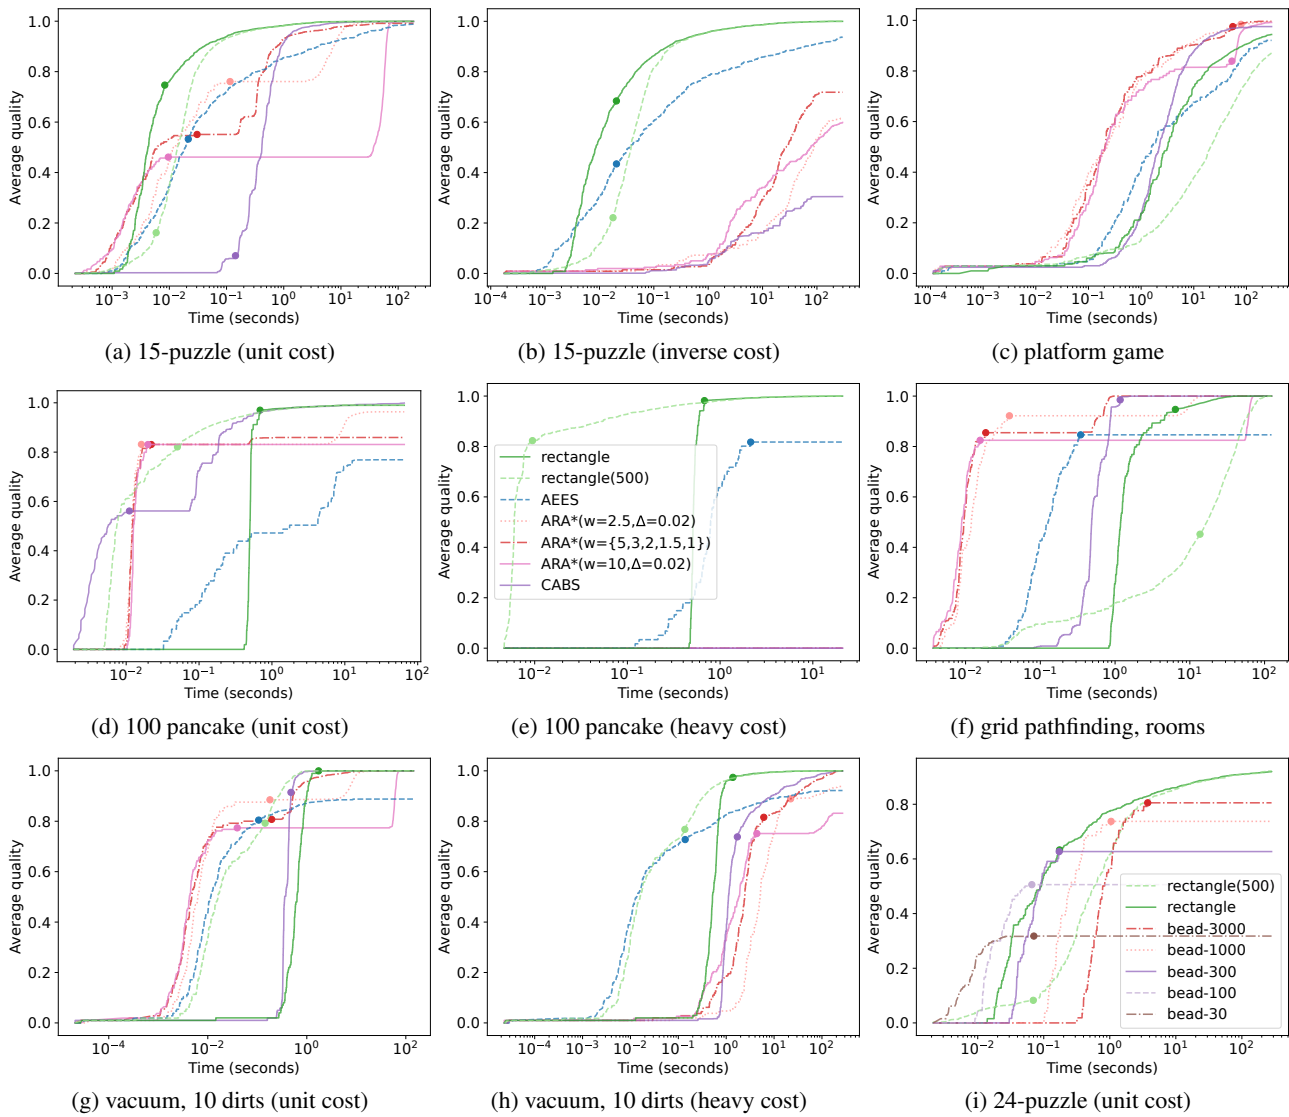


Figure 2: Anytime algorithms (2a-2f, legend in 2e). Rectangle (anytime) versus fixed-width bead (one shot, 2i).

The Pancake Problem

Two cost models were used: unit cost, and heavy cost (Hatem and Ruml 2014), in which each pancake is given an ID number from 1 through N (the number of pancakes), and the cost of a flip is the ID of the pancake above the spatula. 50, 70, and 100 pancakes were used, with the gap heuristic (Helmert 2010), with modifications to include cost per pancake in the heavy cost model.

Selected results are shown in Figures 2d and 2e. In all the sizes of unit pancake, rectangle(1) lags behind the other algorithms in terms of solving. Once it has solved the problems, it provides better solution cost than the other algorithms, but this is not desirable behavior for an anytime algorithm. However, increasing *aspect* improves the performance: rectangle(500) solves unit instances at around the same time as the ARA* variants, and provides better solution cost than all algorithms shortly after. CABS reaches full

coverage slightly earlier than other algorithms, but lacks in terms of quality for much of the time. In heavy cost pancakes, only rectangle(500), rectangle(1) and AEES solve problems reliably, with rectangle(500) being a clear winner.

Vacuum World

In vacuum world (Russell and Norvig 2010) a robot must vacuum up dirt in a grid world. We tested both unit costs and heavy costs, where the cost of movement is equal to the number of dirt that the robot has already vacuumed (Thayer and Ruml 2011). The heuristic sums the number of remaining dirts, the edges of the minimum spanning tree (MST) of the Manhattan distances among the dirts, and the minimum Manhattan distance from the agent to one of the dirts. For the heavy cost model, the distance components were multiplied by the number of dirts cleaned so far. We tested three problem sizes: 200×200 grid with 10 dirts, 500×500 grid

with 20 dirt, and 500×500 grid with 60 dirt.

Selected results are shown in Figures 2g and 2h. ARA* leads in performance in smaller unit cost problems, though rectangle is competitive, reaching full coverage only slightly after ARA*. Rectangle also improves quality beyond ARA* shortly after reaching full coverage. As problem size increases, ARA* loses its lead. In the 500x500 problems with 60 dirt rectangle(1) lags, but rectangle(500) dominates in terms of coverage. CABS provides better quality for some of the time in these largest vacuum problems. In heavy cost, only rectangle and AEES are able to solve reliably in all problem sizes tested, with rectangle(500) clearly superior.

Grid Pathfinding

For 2D grid pathfinding problems we used 5000x5000 maps with random obstacles (grid-random) with 4-way movement (unit cost and life cost (Ruml and Do 2007)), and structured maps of room64 and orz100d (Sturtevant 2012) with 8-way movement. For 4-way movement, we use cost-adjusted Manhattan distance; for 8-way movement, octile distance.

Results in the rooms map are shown in Figure 2f. Overall, ARA* gives best results in all of the maps used. Rectangle provides comparable solution costs for the instances which it solves, but tends to lag significantly on coverage, even more than AEES and CABS. The performance of rectangle(1) is relatively consistent across the maps tested. However, the performance of rectangle(500) is worse in orz100d than in random grids, and worse in 64 room maps than in orz100d. We will return to this phenomenon below.

A Platform Game

The platform domain (Burns, Ruml, and Do 2013) is based on 2D video games in which a character jumps between platforms to reach a goal location. The heuristic uses visibility graphs and actions have unit cost. Not all actions in this domain are reversible, as gravity has an effect on available movements while jumping. We used 100 problem instances defined as 50x50 grids. Results are shown in Figure 2c. Most algorithms give comparable performance on this domain, though ARA* provides slightly better quality and is first to reach full coverage. The other algorithms approach full coverage but do not quite reach it in the time given. It is noteworthy that in this domain, rectangle(500) performs worse than rectangle(1), indicating that reconsideration of earlier steps may matter more than quick, deep exploration.

Dynamic Traffic

The traffic domain (Kiesel, Burns, and Ruml 2015) has both fixed and moving obstacles that the agent must navigate around to reach a goal. The agent knows the future positions of all obstacles. When the agent collides with an obstacle, it is returned to the start location. This search space is directed, because the choices available to the agent change as obstacles move and the same configuration of obstacles may never occur again. We used 100 grid maps of size 100x100 with 5000 total obstacles.

In this domain, all algorithms give comparable performance, except rectangle(1), which solves much later than all

| | worse | better |
|-----------|------------------------------|-----------------|
| rect(500) | non-random grids | pan, hvy vac |
| AEES | 24, pan, bw | |
| rect(1) | pan, vac, traffic, grid | tiles |
| CABS | tiles, hvy pan, bw, grid | |
| ARA* 5 | 15 inv, 24, pan, bw, hvy vac | plat, room, orz |
| ARA* 2.5 | 15 inv, 24, pan, bw, hvy vac | plat, room, orz |
| ARA* 10 | tiles, 24, pan, bw, hvy vac | plat, room, orz |

Table 1: Where each method is worse or better than others.

of the other algorithms. ARA* reaches full coverage slightly before the other algorithms, but rectangle(500) and CABS improve on quality more quickly than the other algorithms after reaching full coverage.

Strengths and Weaknesses of Rectangle Search

Table 1 summarizes the domains in which each algorithm performs exceptionally well or poorly relative to the others. When a general domain like ‘tiles’ is listed instead of a specific variant like ‘hvy vac’, this indicates that most variants gave similar relative performance. Rectangle search provides competitive results across a wide variety of domains. While it is not the best algorithm in all domains, rectangle(1) gives superior performance in all sliding tiles configurations, and rectangle(500) in pancakes and the heavy vacuum problems. Perhaps more importantly, rectangle(500) is not entirely inferior to other algorithms in any domain. Its only clear failures are on the non-random map variants of grid pathfinding (orz100d and 64room) — we discuss this below. On the other hand, all other algorithms tested are markedly inferior (either failing to solve or providing poor quality) in at least two of the domains tested. For these reasons, rectangle(500) appears promising as a method of first resort on new problem domains.

In cases where the heuristic gives extremely accurate guidance, the additional exploration of rectangle(1) at higher depth levels will be unnecessary and increasing *aspect* can lead to rectangle finding solutions faster. However, not all the settings in which rectangle search struggles to compete are ones in which the heuristic is extremely accurate, and adjusting *aspect* does not adequately solve performance problems in all domains.

To understand how rectangle search can perform better than best-first searches and when it can be expected to perform worse, we focus on heuristic error—in particular, local minima. Consider the grid pathfinding problem shown in Figure 3, with a large room containing a slalom of cup-like obstacles and small exits from the room near the start. States are colored green along the solution path and a gradient from yellow to red shows expansion order, with yellow denoting earliest and dark red latest. If the goal location is outside the room (left panels), most of the room forms one large local minimum. The best path is to move through one of the nearby exits but the heuristic tempts the search inward. A best-first search (GBFS in the figure, as it is the limiting case of ARA* and AEES) will explore the entire local minimum until it has exhausted all states that appear

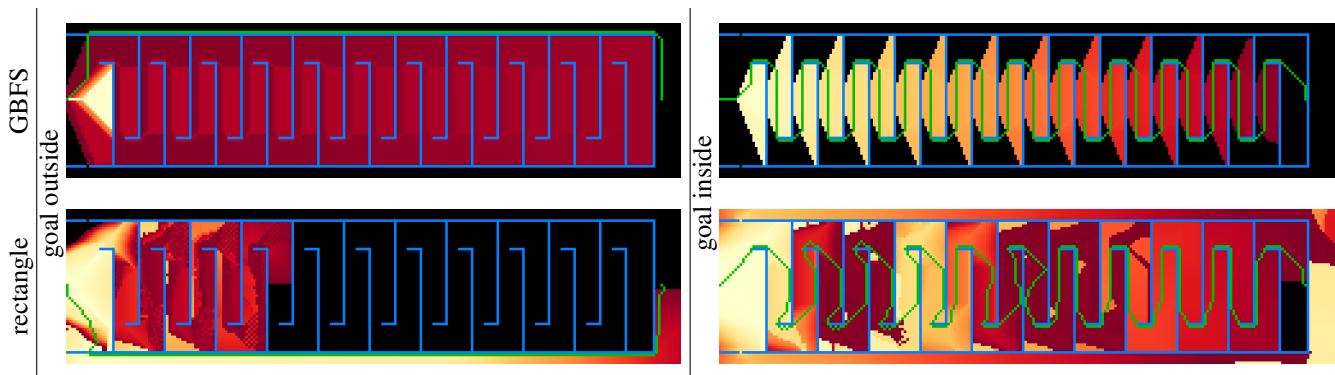


Figure 3: Good and bad scenarios for rectangle search vs best-first search.

better than the exits. Rectangle search (shown here with *aspect=1*), however, widens its exploration of the early regions of the search space and finds a path through an exit and to the goal quickly without exploring most of the local minimum. On the other hand, when the goal is inside the obstacle (right panels), the cups form many small local minima and it is best for a search to focus its efforts inside the room. GBFS quickly fills each minimum and reaches the goal, while rectangle’s skeptical widening causes it to search both inside and outside the room, even re-expanding useless states when it encounters better paths to them.

Rectangle search assumes that it is helpful to re-consider early decisions. This allows it to escape problematic regions of the search space such as a large minimum. But in problems where progress is periodically made and it is better to commit to paths than to reconsider early choices, rectangle search will waste effort. Therefore, rectangle search is strongest in problems where there are deep local minima that it can avoid searching, but is not the best choice when local minima are small and the search can profit from committing to nodes with low heuristic values. For example, in the 64room map of grid pathfinding (Figure 2f), the rooms create a sequence of local minima, causing rectangle search to waste effort exploring adjacent rooms near the start while ARA* and AEES can forget about a room once they escape it. In contrast, the tiny local minima created by randomly placed single-cell obstacles are easy to escape by generating a few successors (considering alternative fringe nodes earlier in the search is not necessary), so rectangle(500) is no worse than other algorithms in such problems.

Comparison with Bead Search

We also performed a comparison of rectangle search with fixed-width bead search to understand how the additional overhead of rectangle search compares to the results obtainable by a well-selected beam width. In the sliding tile puzzle, rectangle search provides competitive quality to bead at a variety of widths, and is able to continue improving its solution quality where bead cannot. Figure 2i shows results for the unit cost 24-puzzle. While rectangle(1) is clearly superior in this setting, rectangle(500) still gives comparable results to many of the fixed-width beam searches. Similar

results were found for the rest of the sliding tiles, blocks world, and the platform game. In the pancake problem, the smaller vacuum problems, grid pathfinding, and traffic, it is rectangle(500) which performs approximately as well as fixed width bead searches or better, with rectangle(1) doing as well as many but not all of the fixed-width bead searches. In only the larger vacuum problems with unit cost do we see both configurations of rectangle performing significantly more poorly than bead searches.

Overall, rectangle search is able to provide solutions about as quickly and with comparable quality to bead search with fixed widths across most of the domains tested. This means that even in settings where anytime behavior is not required, rectangle search can serve as the algorithm of choice.

Discussion

BULB (Furcy and Koenig 2005) behaves like regular beam search until a given depth limit is reached, at which point it uses backtracking to continue the search. In contrast, rectangle search represents a new alternative to conventional beam search and does not require a depth limit. For huge problems where memory capacity is an issue, it would be interesting to integrate BULB-like backtracking with rectangle search.

We have investigated rectangle search’s performance with *aspect=1* and *aspect=500*. Additional research will be necessary to understand how to tune this parameter. Nonlinear increases of width and depth are also a possibility. We leave exploration of these variants to future work.

Conclusions

Rectangle search is an effective anytime algorithm with a simple design. Unlike previous anytime algorithms, which are based on best-first search, rectangle search is instead based on breadth-first search. It has similarities to both monotonic beam search and depth-first algorithms like ILDS. It enforces exploration at a variety of depths in the search tree, which allows it to escape large local minima. Rectangle search is often an effective replacement for fixed-width beam searches. Overall, rectangle search’s promising performance suggests that suboptimal non-best-first heuristic search deserves further exploration.

Acknowledgments

We are grateful for support from the NSF-BSF program (via NSF grant 2008594) and Earlham College (via the Lemann Student/Faculty Collaborative Research Fund).

References

- Bisiani, R. 1987. Beam search. In Shapiro, S. C., ed., *Encyclopedia of Artificial Intelligence*, 56–58. John Wiley and Sons.
- Burns, E.; Ruml, W.; and Do, M. B. 2013. Heuristic Search When Time Matters. *Journal of Artificial Intelligence Research*, 47: 697–740.
- Dean, T. L.; and Boddy, M. S. 1988. An Analysis of Time-Dependent Planning. In *Proceedings of AAAI*, 49–54.
- Dionne, A. J.; Thayer, J. T.; and Ruml, W. 2011. Deadline-Aware Search Using On-line Measures of Behavior. In *Proceedings of SoCS*.
- Furcy, D.; and Koenig, S. 2005. Limited Discrepancy Beam Search. In *Proceedings of IJCAI*.
- Hansen, E. A.; and Zhou, R. 2007. Anytime Heuristic Search. *Journal of Artificial Intelligence Research*, 28: 267–297.
- Hart, P. E.; Nilsson, N. J.; and Raphael, B. 1968. A Formal Basis for the Heuristic Determination of Minimum Cost Paths. *IEEE Transactions of Systems Science and Cybernetics*, SSC-4(2): 100–107.
- Hatem, M.; and Ruml, W. 2014. Bounded suboptimal search in linear space: New results. In *Proceedings of SoCS-14*.
- Helmert, M. 2010. Landmark heuristics for the pancake problem. In *Proceedings of SoCS-10*.
- Heusner, M. 2019. *Search behavior of greedy best-first search*. Ph.D. thesis, University of Basel.
- Hoffmann, J.; and Nebel, B. 2001. The FF Planning System: Fast Plan Generation Through Heuristic Search. *Journal of Artificial Intelligence Research*, 14: 253–302.
- Kiesel, S.; Burns, E.; and Ruml, W. 2015. Achieving goals quickly using real-time search: experimental results in video games. *Journal of Artificial Intelligence Research*, 54: 123–158.
- Korf, R. E. 1996. Improved Limited Discrepancy Search. In *Proceedings of AAAI-96*, 286–291.
- Lelis, L. H.; Zilles, S.; and Holte, R. C. 2013. Stratified Tree Search: A Novel Suboptimal Heuristic Search Algorithm. In *Proceedings of AAMAS-13*, 555–562.
- Lemons, S.; Linares López, C.; Holte, R. C.; and Ruml, W. 2022. Beam Search: Faster and Monotonic. In *Proceedings of ICAPS*.
- Lemons, S.; Ruml, W.; Holte, R. C.; and Linares López, C. 2023. Rectangle Search: An Anytime Beam Search (Extended Version). arXiv:2312.12554.
- Libralesso, L.; Bouhassoun, A.-M.; Cambazard, H.; and Jost, V. 2020. Tree search for the sequential ordering problem. In *Proceedings of ECAI-20*.
- Likhachev, M.; Gordon, G.; and Thrun, S. 2004. ARA*: Anytime A* with Provable Bounds on Sub-Optimality. In *Proceedings of NIPS 16*.
- Likhachev, M.; and Ferguson, D. 2009. Planning long dynamically feasible maneuvers for autonomous vehicles. *The International Journal of Robotics Research*, 28(8): 933–945.
- Pohl, I. 1973. The avoidance of (relative) catastrophe, heuristic competence, genuine dynamic weighting and computational issues in heuristic problem solving. In *Proceedings IJCAI*, 20–23.
- Richter, S.; Thayer, J.; and Ruml, W. 2010. The Joy of Forgetting: Faster Anytime Search via Restarting. In *Proceedings of ICAPS*, 137–144.
- Ruml, W.; and Do, M. B. 2007. Best-first Utility-guided Search. In *Proceedings of IJCAI-07*, 2378–2384.
- Russell, S. J.; and Norvig, P. 2010. *Artificial Intelligence: A Modern Approach*. Pearson.
- Russell, S. J.; and Zilberstein, S. 1991. Composing Real-Time Systems. In *Proceedings of IJCAI-91*, 212–217.
- Sturtevant, N. 2012. Benchmarks for Grid-Based Pathfinding. *Transactions on Computational Intelligence and AI in Games*, 4(2): 144 – 148.
- Thayer, J.; Benton, J.; and Helmert, M. 2012. Better parameter-free anytime search by minimizing time between solutions. In *Proceedings of SoCS*, 120–128.
- Thayer, J. T.; and Ruml, W. 2011. Bounded Suboptimal Search: A Direct Approach Using Inadmissible Estimates. In *Proceedings of IJCAI*.
- Van Den Berg, J.; Shah, R.; Huang, A.; and Goldberg, K. 2011. Anytime Nonparametric A*. In *Proceedings of AAAI-11*, 105–111.
- Vempaty, N. R.; Kumar, V.; and Korf, R. E. 1991. Depth-first vs best-first search. In *Proceedings of AAAI-91*, 434–440.
- Wilt, C.; and Ruml, W. 2014. Speedy Versus Greedy Search. In *Proceedings of SoCS*.
- Zhang, W. 1998. Complete Anytime Beam Search. In *Proceedings of AAAI-98*, 425–430.
- Zilberstein, S. 1996. Using Anytime Algorithms in Intelligent Systems. *AI Magazine*, 17(3): 73.
- Zilberstein, S.; and Russell, S. 1996. Optimal composition of real-time systems. *Artificial Intelligence*, 82(1): 181–213.