

Faculty of Mathematics and Physics
Charles University
November 12th 2024



Artificial Intelligence for Computer Games

More search algorithms

Bidirectional search

Partially unknown environments

Rapid random trees


Goal Bounding

Pruning the space




Idea roughly


When searching towards the goal, do not open nodes that leads outside your goal



A*



JPS+



JPS+ Goal Bounding

Image(s) from the presentation:

http://www.gameapro.com/Rabin_AISummitGDC2015_JPSPlusGoalBounding.zip


Goal Bounding

Pruning the space




Idea roughly


When searching towards the goal, do not open nodes that leads outside your goal



A*



JPS+



JPS+ Goal Bounding

Image(s) from the presentation:

http://www.gameapro.com/Rabin_AISummitGDC2015_JPSPlusGoalBounding.zip

Goal Bounding

Pruning the space



- A method to prune the search space
 - Requires preprocessing the search space offline, so map must be static
 - Assumes 2D maps, but extensible to arbitrary dimensions
- Usable for regular grids as well as navmeshes
- Two sources:
 - Rabin, S. 2015. JPS+ now with Goal Bounding: Over 1000 × Faster than A*, GDC 2015. [[PPTX](#)]
 - Rabin, S., Sturtevant, N.R. 2017. Faster A* with Goal Bounding, Game AI Pro 3. [[PDF](#)]

Goal Bounding


Pruning the space



Idea

For each oriented edge, store bounding box of the area that contains all nodes that are part of all optimal paths leading through that edge. Use it to prune edges during expansion.

Nodes on optimal paths reachable via going left from the point



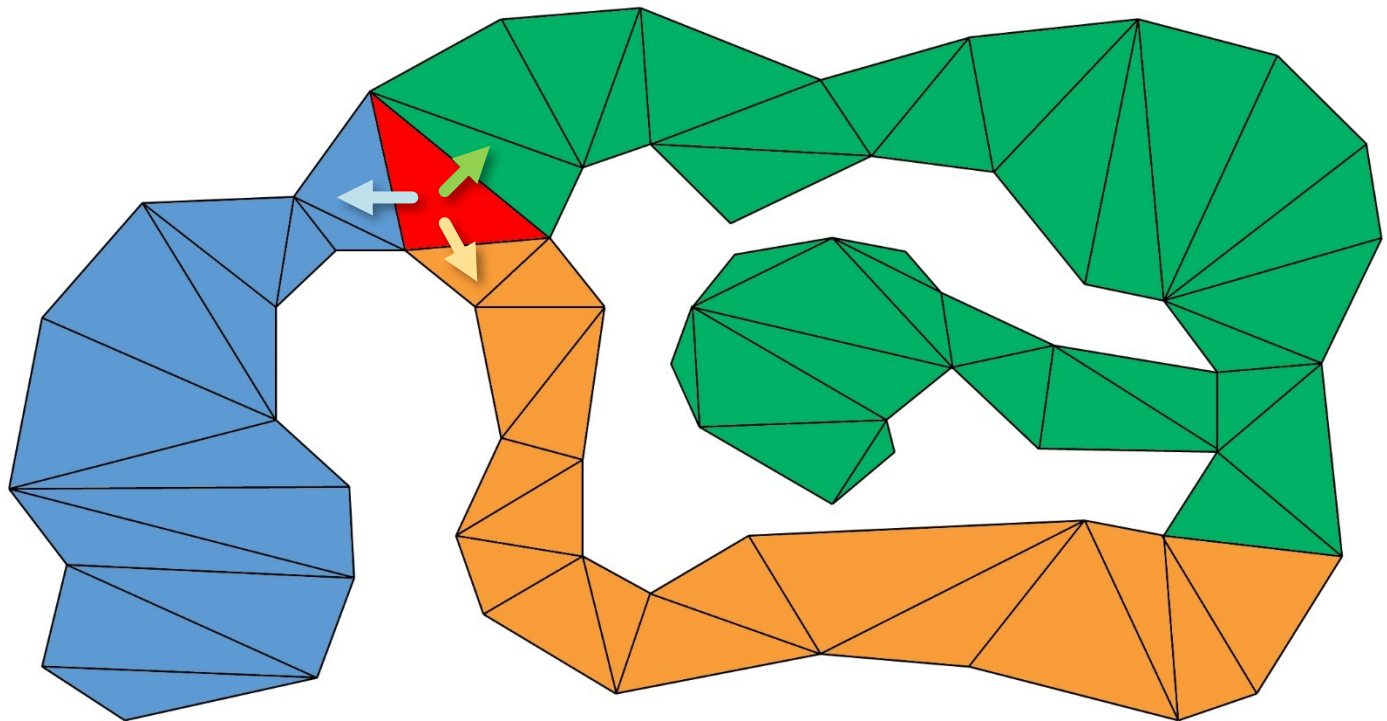
Bounding box of those nodes.

Goal Bounding

Pruning the space



Nodes
optimally
reachable
through
edges of red
triangle.




Goal Bounding

Pruning the space



Nodes
optimally
reachable
through
edges of red
triangle.

Respective
bounding
boxes
enclosing
optimally
reachable
areas.



Goal Bounding

Graph search algorithm integration



Goal bounding can be **integrated** into general graph-search algorithm template! Goal bounding box check is fast $\sim O(1)$.

Algorithm template

1. **make** open-list
2. **push** start into open-list
3. **while** open-list **not empty**
4. **extract** node from open-list according to "strategy"
5. **if** node is target
6. **return** path to node
7. **else**
8. **expand** node by checking its direct neighbors,
 ignoring neighbors whose goal bounding box
 does NOT contain the target,
 possibly adding **those who do** into open-list
9. **move** expanded node to closed-list


Goal Bounding

Precomputation phase



- Precomputation must be done for each graph node
- Can be easily run in parallel for each node

Nodes on optimal paths reachable via going left from the point



Bounding box of those nodes.


Goal Bounding

Precomputation phase



Precomputation idea - step 1

For the given node, run Dijkstra's algorithm in flood fill mode (no target) marking each node reached with the first edge of the path towards that node.




Goal Bounding

Precomputation phase



Precomputation idea - step 2

For each edge, compute the bounding box of the nodes marked in previous step, store it.



Pathfinding




Bidirectional search

Bidirectional search


Primer



- Idea: search from both ends (start \rightarrow target; target \rightarrow start) until the searches meet



Dijkstra's



Bidirectional Dijkstra's

Bidirectional search

Shortest paths



- Given a directed graph G with n vertices, m edges
- Every edge $v \rightarrow w$ has a length $l(v, w)$
- Let $\text{dist}(v, w)$ be the shortest-path distance from v to w
- Goal: find path from start vertex s to goal vertex t with distance $\text{dist}(s, t)$

Bidirectional search

Review: Dijkstra's algorithm



- for every vertex v , remembers
 - $d(v)$: shortest known distance from s to v
 - $p(v)$: parent
 - $S(v)$: status = unreached, frontier, expanded
- Initially $d(s) = 0$, $p(s) = \text{nil}$, $S(s) = \text{frontier}$
 - For all other vertices v : $d(v) = \infty$, $p(v) = \text{nil}$, $S(v) = \text{unreached}$
- In each iteration:
 - choose frontier vertex v with smallest $d(v)$
 - for each edge (v, w) in graph:
 - if $d(w) > d(v) + l(v, w)$:
 - set $d(w) = d(v) + l(v, w)$
 - set $p(w) = v$
 - set $S(w) = \text{frontier}$
 - set $S(v) = \text{expanded}$

Bidirectional search

Dijkstra's algorithm: properties



- Every vertex is expanded only once
- When a vertex v is expanded, $d(v)$ is the shortest distance from the start s to v
- Vertices are expanded in non-decreasing order of distance from s

Bidirectional search

Dijkstra's algorithm




- A **forward search** from s to t or a **reverse search** from t to s will produce the same result
- We can run both at once!
 - Each has its own priority queue
 - Each stores independent values for $d(v)$, $p(v)$, $S(v)$
 - We write $d_f(v)$, $d_r(v)$ for the forward/reverse distances
- We alternate steps of both searches
- Stop when the searches meet
 - What does this mean, exactly?

Bidirectional search

Dijkstra's algorithm



An example



Bidirectional search

Dijkstra's algorithm



- We remember the shortest path seen so far, and its length μ (initially ∞)
- When we discover an edge (v, w) where v and w have already been expanded:
 - If $d_f(v) + l(v, w) + d_r(w) < \mu$ then update μ and the path
- Let top_f and top_r be the smallest values in the forward and reverse priority queues
- We can stop when $top_f + top_r \geq \mu$. Then we have already found the shortest path.
 - If we ever expand any vertex in both directions, the stopping condition will always be true

Bidirectional search

Dijkstra's algorithm



- Why can we stop when $top_f + top_r \geq \mu$?
 - Suppose there is a path P with length less than μ
 - So for every vertex x on P , we must have
 - $dist(s, x) < top_f$ or $dist(x, t) < top_r$
 - So P must contain an edge (v, w) such that
 - $dist(s, v) < top_f$ and $dist(w, t) < top_r$
 - So we have already expanded v and w
 - When we expanded the second of these, we would have already discovered this path and set μ to its length!

Bidirectional search

Review: A*



- A* uses a heuristic function $h(v)$
- h is **admissible** (optimistic) if for every vertex v ,
 $h(v) \leq \text{dist}(v, t)$
- h is **consistent** if for every edge (v, w) , $h(v) \leq l(v, w) + h(w)$
- Every consistent heuristic is admissible

Bidirectional search

Review: A*



- for every vertex v , A* remembers
 - $g(v)$: shortest known distance from s to v
 - $p(v)$: parent
 - $S(v)$: status = unreached, frontier, expanded
- Initially $g(s) = 0$, $p(s) = \text{nil}$, $S(s) = \text{frontier}$
 - For all other vertices v : $g(v) = \infty$, $p(v) = \text{nil}$, $S(v) = \text{unreached}$
- In each iteration:
 - choose frontier vertex v with smallest $f(v) = g(v) + h(v)$
 - for each edge (v, w) in graph:
 - if $g(w) > g(v) + l(v, w)$:
 - set $g(w) = g(v) + l(v, w)$
 - set $p(w) = v$
 - set $S(w) = \text{frontier}$
 - set $S(v) = \text{expanded}$

Bidirectional search

A*: graph transformation



- If h is consistent, then $h(v) \leq l(v, w) + h(w)$
 - So $l(v, w) - h(v) + h(w) \geq 0$
- Define $l_h(v, w) = l(v, w) - h(v) + h(w)$
- Consider a graph G' that's like G , but uses length function l_h
- Let $\text{dist}_h(x, y)$ be the shortest distance from x to y in G'
- For all x and y , $\text{dist}_h(x, y) = \text{dist}(x, y) - h(x) + h(y)$
- A path from x to y in G is a shortest path iff it is a shortest path from x to y in G'

Bidirectional search

A*: graph transformation



- We know that
 - for all x and y , $\text{dist}_h(x, y) = \text{dist}(x, y) - h(x) + h(y)$
- A* on graph G is the same as Dijkstra's on G' !
 - A* on G picks vertex with smallest
 - $f(v) = g(v) + h(v) = \text{dist}(s, v) + h(v)$
 - Dijkstra's on G' picks vertex with smallest
 - $d_h(v) = \text{dist}_h(s, v) = \text{dist}(s, v) - h(s) + h(v)$
 - $h(s)$ is constant, so these are the same

Bidirectional search

Bidirectional A*



- We need two heuristic functions
 - $h_f(v)$: estimate of $\text{dist}(v, t)$
 - $h_r(v)$: estimate of $\text{dist}(s, v)$
- We want these functions to produce the same transformed graph
 - so we can use the stopping criterion from bidirectional Dijkstra's
- Forward: $l_f(v, w) = l(v, w) - h_f(v) + h_f(w)$
- Reverse: $l_r(w, v) = l(v, w) - h_r(w) + h_r(v)$
- For all edges (v, w) , we need $l_f(v, w) = l_r(w, v)$
 - so $h_f(v) + h_r(v) = h_f(w) + h_r(w)$
 - The function $(h_f + h_r)$ must be constant!
 - Most heuristics (e.g. Euclidean distance) are not like that

Bidirectional search

Bidirectional A*



- Given heuristic functions h_f, h_r
- Define
 - $p_f(v) = (h_f(v) - h_r(v)) / 2$
 - $p_r(v) = (h_r(v) - h_f(v)) / 2$
 - Then $p_f(v) + p_r(v) = 0$
 - We can show that p_f and p_r are consistent / admissible
- Refinement
 - $p_f(v) = (h_f(v) - h_r(v) + h_r(t)) / 2$
 - $p_r(v) = (h_r(v) - h_f(v) + h_f(s)) / 2$
 - Now $p_f(t) = p_r(s) = 0$
 - $p_f + p_r$ is still a constant function
 - p_f and p_r are still consistent / admissible

Bidirectional search

Bidirectional A*



- Using heuristic functions p_f and p_r , A* is equivalent to Dijkstra's on the graph G'
- Bidirectional Dijkstra's stops when $\text{top}_f' + \text{top}_r' \geq \mu'$
- Let v_f be the top element in the forward heap
 - $\text{top}_f = \text{dist}(s, v_f) + p_f(v_f)$
 - $\text{top}_f' = \text{dist}_{p_f}(s, v_f) = \text{dist}(s, v_f) - p_f(s) + p_f(v_f)$
 - So $\text{top}_f' = \text{top}_f - p_f(s)$
 - Similarly, $\text{top}_r' = \text{top}_r - p_r(t)$
- $\mu' = \text{dist}_{p_f}(s, t) = \mu - p_f(s) + p_f(t)$
- So we can stop when
 - $[\text{top}_f - p_f(s)] + [\text{top}_r - p_r(t)] \geq \mu - p_f(s) + p_f(t)$
- Simplifying and using $p_f(t) = 0$, we have
 - $\text{top}_f + \text{top}_r \geq \mu + p_r(t)$

Pathfinding

The Algorithms



How to find a path in an unknown environment?

Dynamic searches
(peeking into the field of robotics)

Dynamic pathfinding

Problem statement



The Problem

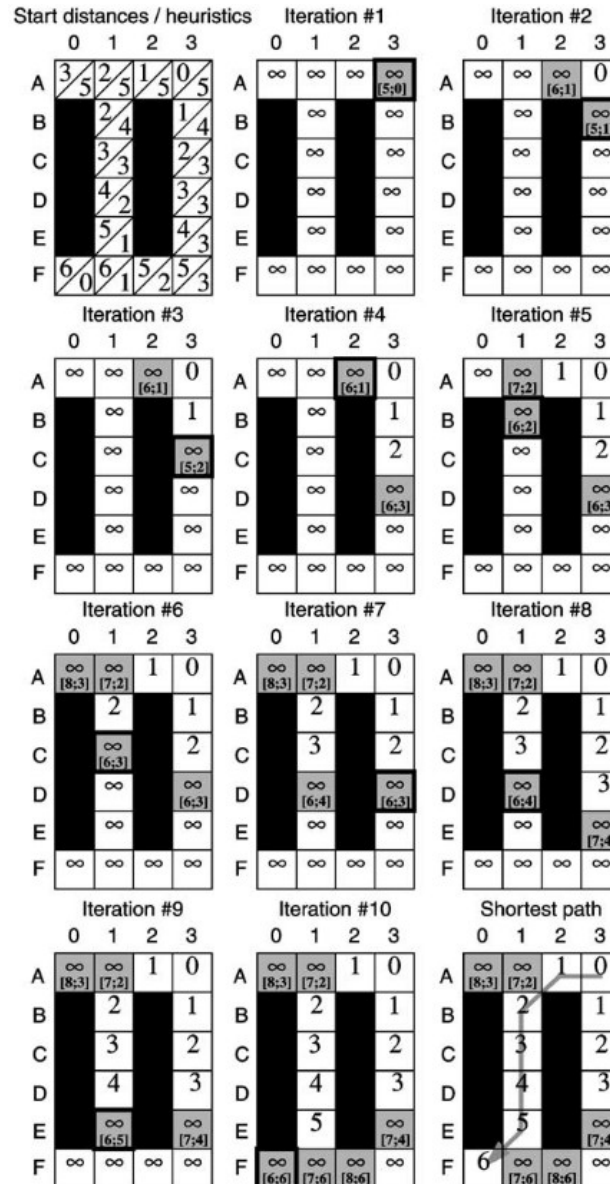
The environment is not known in advance and is being updated online based on sensor readings. How to find a path to the target?

D* Lite; based on Lifelong Planning A* (LPA*)

Koenig, S., & Likhachev, M. (2002). [D* Lite](#). *Aaai/iaai*, 15.


Lifelong Planning A*

Example: first search



Lifelong Planning A*

Example: second search



Lifelong Planning A* (LPA*)



- For each vertex s , LPA* maintains two estimates of the shortest distance to it:
 - $g(s)$: like in A*
 - $rhs(s)$: a one-step lookahead
 - $rhs(s_{start}) = 0$
 - for other states s ,
 - $rhs(s) = \min (g(s') + c(s', s))$ over all neighbors s'
- s is called **locally consistent** if $g(s) = rhs(s)$
 - If $g(s) < rhs(s)$, it is **locally underconsistent**. If $g(s) > rhs(s)$, it is **locally overconsistent**.
- If all vertices are locally consistent, then all values $g(s)$ are shortest-path distances
- The open list holds all locally inconsistent vertices
- It is a priority queue ordered lexicographically by $k(s) = [k_1(s); k_2(s)]$, where
 - $k_1(s) = \min(g(s), rhs(s)) + h(s)$
 - $k_2(s) = \min(g(s), rhs(s))$

Lifelong Planning A*

procedure CalculateKey(s)

{01} return [$\min(g(s), rhs(s)) + h(s)$; $\min(g(s), rhs(s))$];

procedure Initialize()

{02} $U = \emptyset$;

{03} for all $s \in S$ $rhs(s) = g(s) = \infty$;

{04} $rhs(s_{start}) = 0$;

{05} $U.Insert(s_{start}, [h(s_{start}); 0])$;

procedure UpdateVertex(u)

{06} if ($u \neq s_{start}$) $rhs(u) = \min_{s' \in pred(u)} (g(s') + c(s', u))$;

{07} if ($u \in U$) $U.Remove(u)$;

{08} if ($g(u) \neq rhs(u)$) $U.Insert(u, CalculateKey(u))$;

procedure ComputeShortestPath()

{09} while ($U.TopKey() < CalculateKey(s_{goal})$ OR $rhs(s_{goal}) \neq g(s_{goal})$)

{10} $u = U.Pop()$;

{11} if ($g(u) > rhs(u)$)

{12} $g(u) = rhs(u)$;

{13} for all $s \in succ(u)$ $UpdateVertex(s)$;

{14} else

{15} $g(u) = \infty$;

{16} for all $s \in succ(u) \cup \{u\}$ $UpdateVertex(s)$;

procedure Main()

{17} Initialize();

{18} forever

{19} ComputeShortestPath();

{20} Wait for changes in edge costs;

{21} for all directed edges (u, v) with changed edge costs

{22} Update the edge cost $c(u, v)$;

{23} UpdateVertex(v);

D* Lite



D* Lite (Koenig, 2002)

- An adaptation of LPA* for dynamic planning
- Now the start vertex may change each time we replan!



D* Lite

- First modification of LPA*: we must search **backwards** from the goal to the start
 - because we always want shortest paths to the fixed goal
 - Now the heuristic function estimates the distance from any position to the robot
- Second modification: every time the robot moves, the heuristic function changes, so we must recalculate all priorities

D* Lite:
Basic Version

```
procedure CalculateKey(s)
{01'} return [ $\min(g(s), rhs(s)) + h(s_{start}, s)$ ;  $\min(g(s), rhs(s))$ ];

procedure Initialize()
{02'}  $U = \emptyset$ ;
{03'} for all  $s \in S$   $rhs(s) = g(s) = \infty$ ;
{04'}  $rhs(s_{goal}) = 0$ ;
{05'}  $U.Insert(s_{goal}, CalculateKey(s_{goal}))$ ;

procedure UpdateVertex(u)
{06'} if ( $u \neq s_{goal}$ )  $rhs(u) = \min_{s' \in Succ(u)} (c(u, s') + g(s'))$ ;
{07'} if ( $u \in U$ )  $U.Remove(u)$ ;
{08'} if ( $g(u) \neq rhs(u)$ )  $U.Insert(u, CalculateKey(u))$ ;

procedure ComputeShortestPath()
{09'} while ( $U.TopKey() < CalculateKey(s_{start})$  OR  $rhs(s_{start}) \neq g(s_{start})$ )
{10'}    $u = U.Pop()$ ;
{11'}   if ( $g(u) > rhs(u)$ )
{12'}      $g(u) = rhs(u)$ ;
{13'}     for all  $s \in Pred(u)$  UpdateVertex(s);
{14'}   else
{15'}      $g(u) = \infty$ ;
{16'}     for all  $s \in Pred(u) \cup \{u\}$  UpdateVertex(s);

procedure Main()
{17'} Initialize();
{18'} ComputeShortestPath();
{19'} while ( $s_{start} \neq s_{goal}$ )
{20'}   /* if ( $g(s_{start}) = \infty$ ) then there is no known path */
{21'}    $s_{start} = \arg \min_{s' \in Succ(s_{start})} (c(s_{start}, s') + g(s'))$ ;
{22'}   Move to  $s_{start}$ ;
{23'}   Scan graph for changed edge costs;
{24'}   if any edge costs changed
{25'}     for all directed edges (u, v) with changed edge costs
{26'}       Update the edge cost  $c(u, v)$ ;
{27'}       UpdateVertex(u);
{28'}     for all  $s \in U$ 
{29'}        $U.Update(s, CalculateKey(s))$ ;
{30'}     ComputeShortestPath();
```



D* Lite

- Updating all queue priorities is expensive
- We can modify the algorithm so each queue priority is only a **lower bound** on the actual priority
- When the robot moves from s to s' , actual priorities can decrease by at most $h(s, s')$
- Instead of subtracting $h(s, s')$ from all priorities, we can accumulate the decrease into a variable k_m and then apply it to any vertex when it comes out of the queue
- For details and pseudocode, see the D* Lite paper

Pathfinding

The Algorithms



How to find the path in an unknown environment faster?

Beating D* Lite (almost every time)
(peeking into the field of robotics)

Multipath Adaptive A* (MPAA*)

Speeding up A* to outperform D* Lite



Multipath Adaptive A*

Hernández, C., Baier, J. A., & Asín, R. (2014, May). [Making A* run faster than D*-Lite for path-planning in partially known terrain.](#)

- Proposes Multipath Adaptive A* as a simpler approach, usually faster than D* Lite

Adaptive A*



- Let $gd[s]$ be the minimal cost from state s to the goal
- Let $f^* = gd[s_{start}]$ be the minimal cost found by an A* search
- For any state s that was expanded,
 - $g[s]$ is the minimal cost from the start to s
 - $f[s] = g[s] + h[s]$
- Now
 - $f^* \leq g[s] + gd[s]$
 - $f^* - g[s] \leq gd[s]$
 - So $f^* - g[s]$ is an admissible estimate of $gd[s]$
 - We can use this as a new heuristic value for s
- Also, since s was expanded, we have
 - $f[s] \leq f^*$
 - $g[s] + h[s] \leq f^*$
 - $h[s] \leq f^* - g[s]$
 - So the new heuristic value $f^* - g[s]$ dominates the old

Adaptive A*: pseudocode (part 1)

```
1 procedure InitializeState(s)
2   if search(s) ≠ counter then
3     |  $g(s) \leftarrow \infty$ 
4   |  $search(s) \leftarrow counter$ 
5 procedure A* (sinit)
6   InitializeState(sinit)
7   parent(sinit) ← null
8    $g(s_{init}) \leftarrow 0$ 
9   Open ← ∅
10  insert sinit into Open with f-value  $g(s_{init}) + h(s_{init})$ 
11  Closed ← ∅
12  while Open ≠ ∅ do
13    remove a state s from Open with the smallest f-value  $g(s) + h(s)$ 
14    if GoalCondition(s) then
15      | return s
16    insert s into Closed
17    for each  $s' \in succ(s)$  do
18      InitializeState(s')
19      if  $g(s') > g(s) + c(s, s')$  then
20        |  $g(s') \leftarrow g(s) + c(s, s')$ 
21        | parent(s') ← s
22        | if s' is in Open then
23          | | set priority of s' in Open to  $g(s') + h(s')$ 
24          | else
25          | | insert s' into Open with priority  $g(s') + h(s')$ 
26  | return null
```

Adaptive A*: pseudocode (part 2)

```
27 procedure BuildPath( $s$ )
28   while  $s \neq s_{start}$  do
29      $next(parent(s)) \leftarrow s$ 
30      $s \leftarrow parent(s)$ 

31 procedure Observe( $s$ )
32   for each  $arc(t, t')$  in the range of visibility from  $s$  do
33     if cost of  $(t, t')$  has increased then
34       update  $c(t, t')$ 
35        $next(t) \leftarrow null$ 

36 procedure main()
37    $counter \leftarrow 0$ 
38   Observe( $s_{start}$ )
39   for each state  $s \in S$  do
40      $search(s) \leftarrow 0$ 
41      $h(s) \leftarrow H(s, s_{goal})$ 
42      $next(s) \leftarrow null$ 

43   while  $s_{start} \neq s_{goal}$  do
44      $counter \leftarrow counter + 1$ 
45      $s \leftarrow A^*(s_{start})$ 
46     if  $s = null$  then
47       return "goal is not reachable"

48     for each  $s' \in Closed$  do
49        $h(s') \leftarrow g(s) + h(s) - g(s') / * heuristic$ 
50        $BuildPath(s)$ 
51     while no action cost has just increased in  $path[s_{start}]$  do
52        $t \leftarrow s_{start}$ 
53        $s_{start} \leftarrow next(s_{start})$ 
54        $next(t) \leftarrow null$ 
55       Move agent to  $s_{start}$ 
56       Observe( $s_{start}$ )
```

Multipath Adaptive A*



- Suppose we select a state s such that
 - s belongs to a previously found path σ
 - the suffix of σ starting in s is a provably optimal path from s to the goal
- Then we can stop the search immediately
- We can check these conditions easily

```
1 function GoalCondition( $s$ )
2   while  $next(s) \neq null$  and  $h(s) = h(next(s)) + c(s, next(s))$  do
3      $s \leftarrow next(s)$ 
4   return  $s_{goal} = s$ 
```

Pathfinding

The Algorithms



How to find a path in a continuous space?
Rapidly exploring random trees
(RRT, RRT*)
(peeking into the field of robotics)

Rapidly-exploring Random Trees

RRT and RRT*



RRT

LaValle, S. M., & Kuffner Jr, J. J. (2001). [Randomized kinodynamic planning](#). *The international journal of robotics research*, 20(5), 378-400.

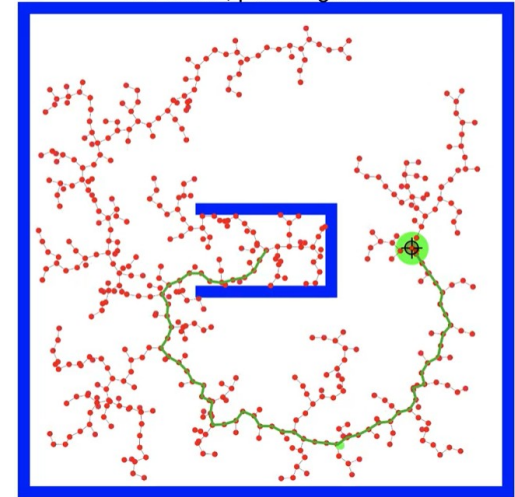
IDEA: Randomly throw a point, find nearest existing point of RRT, make a step from that point towards the random point. Repeat.

RRT*


Karaman, S., & Frazzoli, E. (2011). [Sampling-based algorithms for optimal motion planning](#). *The international journal of robotics research*, 30(7), 846-894.

IDEA: Do RRT but then try to rewiring the tree around new point to be more optimal.

706 nodes, path length 59.92



289 nodes, path length 40.37



Rapidly-exploring Random Trees

RRT and RRT*

Beautiful explanation available on [YouTube](#)

by Aaron Becker.

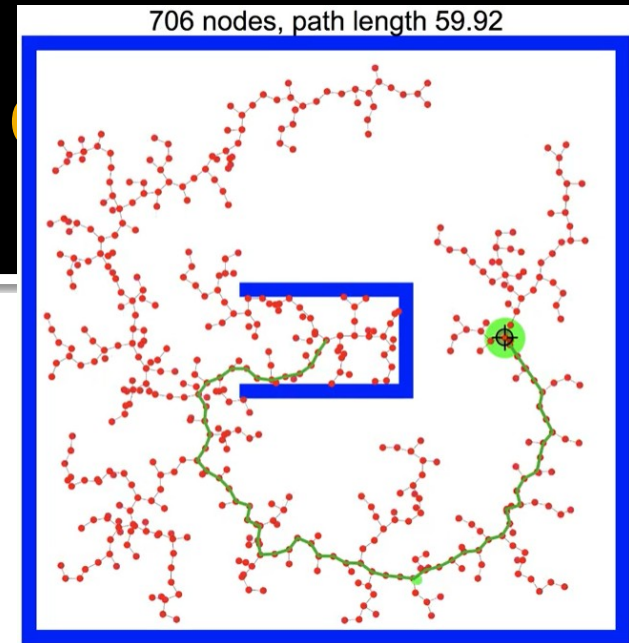


- YouTube video: "RRT, RRT* & Random Trees" (Aaron Becker)

Rapidly-exploring Random Algorithms

■ RT (trying to find path to *target* with certain *epsilon*)

1. `G.init(root: Point)`
2. **while** path to epsilon-area around target not found
3. `point := random`
4. `nearest := G.nearest_vertex(point)`
5. `new_vertex = nearest + (point - nearest).normalized * step`
6. **if** no obstacle nearest->new_vertex
`G.add_edge(nearest, new_vertex)`




Rapidly-exploring Random Algorithms

■ **RRT*** (trying to find path to *target* with certain *epsilon*)

```
1. G.init(root: Point)
2. while path to epsilon-area
   around target not found
3.   point := random
4.   nearest := G.nearest_vertex(point)
5.   new_vertex =
       nearest + (point - nearest).normalized
               * step
6.   if no obstacle nearest->new_vertex
7.     min_cost_vertex =
         find vertex
         | from G.vertices_around(new_vertex)
         | with min path cost from root
         | and no obstacle on vertex->new_vertex
8.   G.add_edge(min_cost_vertex, new_vertex)
9.   for vertex in G.vertices_around(new_vertex)
10.    if obstacle on vertex->new_vertex
11.      continue
12.    if path_cost(root, vertex) >
        path_cost(root, new_vertex) + |vertex,new_vertex|
13.      G.remove_edge(parent(vertex), vertex)
14.      G.add_edge(new_vertex, vertex)
```

289 nodes, path length 40.37





EVROPSKÁ UNIE
Evropské strukturální a investiční fondy
Operační program Výzkum, vývoj a vzdělávání



Material has been produced within and supported by the project
„Zvýšení kvality vzdělávání na UK a jeho relevance pro potřeby trhu práce“
kept under number CZ.02.2.69/0.0/0.0/16_015/0002362.