

A* in Dynamic Environments

BY ŁUKASZ STAFINIAK

Instytut Informatyki Uniwersytetu Wrocławskiego

Real-Time Adaptive A^*

A^* maintains (“A New Principle for Incremental Heuristic Search: Theoretical Results” Maxim Likhachev Sven Koenig):

- estimated minimal cost $g(s)$ from current state s_{curr} to every state s (initially $+\infty$)
- a heuristic $h(s)$ estimating the goal distance
- a priority queue (**open list=fringe nodes** \mathcal{F}) weighted by $f(s) = g(s) + h(s)$ initially with s_{curr} only
- pop a state s , update $g(s')$ for each successor of s , push onto the queue each successor for which $g(s')$ dropped, put s into **closed list=internal nodes** \mathcal{I} .

Take $s \in \text{closed}$ (an **expanded** state=internal node).

After a complete A^* search, $g(s)$ is the cost-minimal path from s_{start} to s .

Let $gd(s)$ – the goal distance of s , $f^* = gd(s_{\text{start}})$ – the optimal cost. Then:

$$f^* - g(s) \leq gd(s)$$

so the updated heuristic $f^* - g(s)$ is **admissible**. It is a better heuristic:

$$h(s) \leq f^* - g(s)$$

Adaptive A^* repeatedly finds cost-minimal paths for problems with the **same goal vertices** and **non-decreasing edge costs**.

The updated heuristic is **consistent**:

$$h'(s) \leq h'(\text{succ}(s, a)) + c(s, a) \leq h'(\text{succ}(s, a)) + c'(s, a)$$

[Compare admissibility and increasing costs requirement to exploration.]

(“Real-Time Adaptive A^* ” Sven Koenig Maxim Likhachev)

s_{term} – a state that **bounded-lookahead** A^* is about to expand when it terminates.

procedure `realtime_adaptive_astar()`:

 while ($s_{\text{curr}} \notin \text{GOAL}$) do

 lookahead := any desired integer greater than zero;

```

astar();
if  $s = \text{FAILURE}$  then
    return FAILURE;
for all  $s \in \text{CLOSED}$  do
     $h(s) := g(s_{\text{term}}) + h(s_{\text{term}}) - g(s)$ ;
movements := any desired integer greater than zero;
while ( $s_{\text{curr}} \neq s_{\text{term}} \wedge \text{movements} > 0$ ) do
     $a :=$  the action in  $A(s_{\text{curr}})$  on the cost-minimal trajectory
    from  $s_{\text{curr}}$  to  $s_{\text{term}}$ ;
     $s_{\text{curr}} := \text{SUCC}(s_{\text{curr}}, a)$ ;
    movements := movements - 1;
    for any desired number of times (including zero) do
        increase any desired  $c(s, a)$  where  $s \in S$  and  $a \in A(s)$ ;
    if any increased  $c(s, a)$  is on the cost-minimal trajectory
    from  $s_{\text{curr}}$  to  $s_{\text{term}}$  then
        break;
return SUCCESS;

```

Model learning is easy: start with the most optimistic $c(s, a) = \varepsilon$ (e.g., $\varepsilon = 1$), update costs from experience (e.g. $c(s, a) := \infty$ when $a \notin A(s)$).

Because of simplicity, RTAA* is the best method for heavily time-constrained domains.

Prioritized Learning Real-Time A*

(“Real-Time Heuristic Search with a Priority Queue” D. Chris Rayner, Katherine Davison, Vadim Bulitko, Kenneth Anderson, Jieshan Lu)

```
function PLRTA*(s):
```

```
    while  $s \neq s_{\text{goal}}$  do
```

```
        StateUpdate(s)
```

```
        repeat
```

```
             $p = \text{queue.Pop}()$ 
```

```
            if  $p \neq s_{\text{goal}}$  then StateUpdate(p)
```

```
        until  $N$  states are updated or  $\text{queue} = \emptyset$ 
```

```
         $s \leftarrow$  neighbor  $s'$  with lowest  $f(s, s') = c(s, s') + h(s')$ 
```

```

function StateUpdate( $s$ )
    find neighbor  $s'$  with lowest  $f(s, s') = c(s, s') + h(s')$ 
     $\Delta \leftarrow f(s, s') - h(s)$ 
    if  $\Delta > 0$  then
         $h(s) \leftarrow f(s, s')$ 
        for all neighbors  $n$  of  $s$  do
            AddToQueue( $n, \Delta$ )

function AddToQueue( $s, \Delta_s$ )
    if  $s \notin$  queue then
        if queue.Full() then
             $(r, \Delta_r) \leftarrow$  queue.Pop()
            if  $\Delta_r < \Delta_s$  then queue.Push( $s, \Delta_s$ )
            else queue.Push( $r, \Delta_r$ )
        else queue.Push( $s, \Delta_s$ )

```

Moving Target Adaptive A*

- instead of the “optimal” minimax (escaping opponent), just generalize A*
- based on Adaptive A* (see the previous notes file)
- we want linear space complexity (don't store distances between arbitrary states visited)
- correct the h values when the goal state changes; let H be the original heuristic

$$h(s) := \max(H(s, s'_{\text{target}}), h(s) - h(s'_{\text{target}}))$$

- the lazy version accumulates updates and applies them when a state is visited (remembering that the previous update was n times ago)

D* Lite

(“D* Lite” Sven Koenig, Maxim Likhachev)

D* Lite is based on **Lifelong Planning A*** (LPA*) which performs A* and accommodates cost changes by replanning backwards from the change points only backing-up distance from start for relevant states.

- no assumptions about how the costs change (up or down, close or far from the current/start state)
- the priority queue always contains the **locally inconsistent** states $g(s) \neq \text{rhs}(s)$
- Initialize() should perform lazily not to tabulate all states

D* Lite lets the current state move by planning backwards from the goal state, replacing as g the distance from current state by the more stable distance from goal state. Heuristic h now measures distances from the current state.

- h admissible $h(s, s') \leq c^*(s, s')$ and obeys the triangle inequality
- not to reorder the queue, accumulate changes $h(s_{\text{prev}}, s_{\text{curr}})$

Minimax LPA* algorithm results by replacing $\min_{s' \in \text{Succ}(u)} (c(u, s') + g(s'))$ with $\min_{a \in \mathcal{A}(u)} \max_{s' \in \text{Succ}(u)} (c(u, s') + g(s'))$ (and performing $\arg \min_{a \in \mathcal{A}(u)}$).

Lifelong Planning A* algorithm

procedure CalculateKey(s):

 return $(\min(g(s), \text{rhs}(s)) + h(s, s_{\text{goal}}), \min(g(s), \text{rhs}(s)))$

procedure Initialize():

$U \leftarrow \emptyset$;

 for all $s \in \mathcal{S}$ do $\text{rhs}(s) \leftarrow g(s) \leftarrow \infty$

$\text{rhs}(s_{\text{start}}) \leftarrow 0$

$U.\text{Insert}(s_{\text{start}}, \text{CalculateKey}(s_{\text{start}}))$

procedure UpdateVertex(u):

 if $u \neq s_{\text{start}}$ then $\text{rhs}(u) \leftarrow \min_{s' \in \text{Pred}(u)} (g(s') + c(s', u))$

 if $u \in U$ then $U.\text{Remove}(u)$

 if $g(u) \neq \text{rhs}(u) \{ \wedge \text{NotYet}(u) \}$ then $U.\text{Insert}(u, \text{CalculateKey}(u))$

```

procedure ComputeShortestPath():
  while  $U.\text{TopKey}() < \text{CalculateKey}(s_{\text{goal}}) \vee \text{rhs}(s_{\text{goal}}) \neq g(s_{\text{goal}})$ 
     $u \leftarrow U.\text{Pop}()$ 
    if  $g(u) > \text{rhs}(u)$ 
       $g(u) \leftarrow \text{rhs}(u)$ 
      for all  $s \in \text{Succ}(u)$  do  $\text{UpdateVertex}(s)$ 
    else
       $g(u) \leftarrow \infty$ 
      for all  $s \in \text{Succ}(u) \cup \{u\}$  do  $\text{UpdateVertex}(s)$ 

procedure Main():
  Initialize()
  forever
    ComputeShortestPath(); {insert all inconsist. states to  $U$ }
    Wait for changes in edge costs
    for all directed edges  $(u, v)$  with changed costs
      update  $c(u, v)$ 
       $\text{UpdateVertex}(v)$ 

```

D* Lite algorithm

procedure CalculateKey(s):

 return $(\min(g(s), \text{rhs}(s)) + h(s_{\text{start}}, s) + k_m, \min(g(s), \text{rhs}(s)))$

procedure Initialize():

$U \leftarrow \emptyset$

$k_m \leftarrow 0$

 for all $s \in \mathcal{S}$ do $\text{rhs}(s) \leftarrow g(s) \leftarrow \infty$

$\text{rhs}(s_{\text{goal}}) \leftarrow 0$

$U.\text{Insert}(s_{\text{goal}}, \text{CalculateKey}(s_{\text{goal}}))$

procedure UpdateVertex(u):

 if $u \neq s_{\text{goal}}$ then $\text{rhs}(u) \leftarrow \min_{s' \in \text{Succ}(u)} (c(u, s') + g(s'))$

 if $u \in U$ then $U.\text{Remove}(u)$

 if $g(u) \neq \text{rhs}(u) \{ \wedge \text{NotYet}(u) \}$ then $U.\text{Insert}(u, \text{CalculateKey}(u))$

```

procedure ComputeShortestPath():
  while  $U.\text{TopKey}() \dot{<} \text{CalculateKey}(s_{\text{start}}) \vee \text{rhs}(s_{\text{start}}) \neq g(s_{\text{start}})$ 
     $k_{\text{old}} \leftarrow U.\text{TopKey}()$ 
     $u \leftarrow U.\text{Pop}()$ 
    if  $k_{\text{old}} \dot{<} \text{CalculateKey}(u)$ 
       $U.\text{Insert}(u, \text{CalculateKey}(u))$ 
    else if  $g(u) > \text{rhs}(u)$ 
       $g(u) \leftarrow \text{rhs}(u)$ 
      for all  $s \in \text{Pred}(u)$  do  $\text{UpdateVertex}(s)$ 
    else
       $g(u) \leftarrow \infty$ 
      for all  $s \in \text{Pred}(u) \cup \{u\}$  do  $\text{UpdateVertex}(s)$ 
  {insert remaining locally inconsistent states to  $U$ }

```

procedure Main():

$s_{\text{last}} \leftarrow s_{\text{start}}$

Initialize()

ComputeShortestPath()

while $s_{\text{start}} \neq s_{\text{goal}}$

$s_{\text{start}} \leftarrow \arg \min_{s' \in \text{Succ}(s_{\text{start}})} (c(s_{\text{start}}, s') + g(s'))$

Move to s_{start}

Scan graph for any changed edge costs

if any edge cost changed

$k_m \leftarrow k_m + h(s_{\text{last}}, s_{\text{start}})$

$s_{\text{last}} \leftarrow s_{\text{start}}$

for all directed edges (u, v) with changed costs

Update $c(u, v)$

UpdateVertex(u)


ComputeShortestPath()

A Generalized Framework for LPA* with inconsistent heuristics

(“A Generalized Framework for Lifelong Planning A* Search” Sven Koenig, Maxim Likhachev)

- Heuristic Search-based Planning (HSP) used in modern symbolic planners
- planners that find minimum-cost plans do not scale to large domains: the need to use more informed but inconsistent heuristics
 - the bigger the cost heuristic, the less states expanded
- experimentally A* works better with inconsistent heuristics, when for f -value tie bigger g is preferred (not smaller)
 - LPA* is not even correct: it can fail to find a finite cost solution
 - no matter whether smaller or bigger g is preferred, LPA* is worse than A* (replanning from scratch) on inconsistent heuristics
- solution: the priority queue of GLPA* does not contain all locally inconsistent states but only those not yet expanded as overconsistent; however, GLPA* updates the priority queue to contain all locally inconsistent states between calls to ComputePlan(); (NotYet in the algorithms)
 - now one can use inconsistent heuristics and reversed g tie breaks

When the Model is Wrong – Dyna+



Dyna- Q + uses a heuristic to enforce exploration (especially useful in non-stationary environments): if a transition has not been tried in n time steps, then planning backups are done as if that transition produced a reward of $r + \kappa\sqrt{n}$, for some small κ . Perhaps a similar idea can be incorporated into above algorithms.

Appendix

D* Lite optimized algorithm

procedure CalculateKey(s):

 return $(\min(g(s), \text{rhs}(s)) + h(s_{\text{start}}, s) + k_m, \min(g(s), \text{rhs}(s)))$

procedure Initialize():

$U \leftarrow \emptyset$

$k_m \leftarrow 0$

 for all $s \in \mathcal{S}$ do $\text{rhs}(s) \leftarrow g(s) \leftarrow \infty$

$\text{rhs}(s_{\text{goal}}) \leftarrow 0$

$U.\text{Insert}(s_{\text{goal}}, (h(s_{\text{start}}, s_{\text{goal}}), 0))$

procedure UpdateVertex(u):

 if $u \neq s_{\text{goal}} \wedge u \in U$ then $\text{rhs}(u) \leftarrow \min_{s' \in \text{Succ}(u)} (c(u, s') + g(s'))$

 if $u \in U \wedge g(u) = \text{rhs}(u)$ then $U.\text{Remove}(u)$

 if $g(u) \neq \text{rhs}(u) \wedge u \in U$ { $\wedge \text{NotYet}(u)$ }

$U.\text{Insert}(u, \text{CalculateKey}(u))$

procedure ComputeShortestPath():

while $U.\text{TopKey}() \dot{<} \text{CalculateKey}(s_{\text{start}}) \vee \text{rhs}(s_{\text{start}}) \neq g(s_{\text{start}})$

$k_{\text{old}} \leftarrow U.\text{TopKey}()$

$u \leftarrow U.\text{Pop}()$

$k_{\text{new}} \leftarrow \text{CalculateKey}(u)$

if $k_{\text{old}} \dot{<} k_{\text{new}}$

$U.\text{Insert}(u, k_{\text{new}})$

else if $g(u) > \text{rhs}(u)$

$g(u) \leftarrow \text{rhs}(u)$

for all $s \in \text{Pred}(u)$

if $s \neq s_{\text{goal}}$

$\text{rhs}(s) \leftarrow \min(\text{rhs}(s), c(s, u) + g(u))$

UpdateVertex(s)

else

$g_{\text{old}} \leftarrow g(u)$

$g(u) \leftarrow \infty$

for all $s \in \text{Pred}(u) \cup \{u\}$

if $\text{rhs}(s) = c(s, u) + g_{\text{old}} \wedge s \neq s_{\text{goal}}$

$\text{rhs}(s) \leftarrow \min(\text{rhs}(s), c(s, u) + g(u))$

UpdateVertex(s)

{insert remaining locally inconsistent states to U }

procedure Main():

$s_{\text{last}} \leftarrow s_{\text{start}}$

Initialize()

ComputeShortestPath()

while $s_{\text{start}} \neq s_{\text{goal}}$

$s_{\text{start}} \leftarrow \arg \min_{s' \in \text{Succ}(s_{\text{start}})} (c(s_{\text{start}}, s') + g(s'))$

Move to s_{start}

Scan graph for any changed edge costs

if any edge cost changed

$k_m \leftarrow k_m + h(s_{\text{last}}, s_{\text{start}})$

$s_{\text{last}} \leftarrow s_{\text{start}}$

for all directed edges (u, v) with changed costs

$c_{\text{old}} \leftarrow c(u, v)$

Update $c(u, v)$

if $c_{\text{old}} > c(u, v)$

if $u \neq s_{\text{goal}}$

$\text{rhs}(u) \leftarrow \min (\text{rhs}(u), c(u, v) + g(v))$

```
else if rhs(u) = cold + g(v)
    if u ≠ sgoal
        rhs(u) ← mins' ∈ Succ(u) (c(u, s') +
            g(s'))
    UpdateVertex(u)
ComputeShortestPath()
```