# Artificial Intelligence 

Roman Barták

Department of Theoretical Computer Science and Mathematical Logic

Uninformed (blind) search algorithms can find an (optimal) solution to the problem, but they are usually not very efficient.

- BFS, DFS, ID, BiS

Informed (heuristic) search algorithms can find solutions more efficiently thanks to exploiting problem-specific knowledge.

- How to use heuristics in search?
- BestFS, A*, IDA*, RBFS, SMA*
- How to construct heuristics?
- relaxation, pattern databases

Recall that we are looking for (the shortest) path from the initial state to some goal state.
Which information can help the search algorithm?

- For example, the length of path to some goal state.
- However such information is usually not available (if it is available then we do not need to do search). Usually some evaluation function $\mathbf{f}(\mathbf{n})$ is used to evaluate "quality" of node $\mathbf{n}$ based on the length of path to the goal.
- best-first search
- The node with the smallest value of $f(n)$ is used for expansion.
- There are search algorithms with different views of $\mathbf{f}(\mathbf{n})$. Usually the part of $\mathbf{f}(\mathbf{n})$ is a heuristic function $\mathbf{h ( n )}$ estimating the length of the shortest (cheapest) path to the goal state.
- Heuristic functions are the most common form of additional information given to search algorithms
- We will assume that $\mathbf{h ( n )}=\mathbf{0} \Leftrightarrow \mathbf{n}$ is goal.


## Greedy best-first search

Let us try to expand first the node that is closest to some goal state, i.e. $\mathbf{f}(\mathbf{n})=\mathbf{h ( n )}$.

- greedy best-first search algorithm

Example (path Arad $\rightarrow$ Bucharest):

- We have a table of direct distances from any city to Bucharest.
- Note: this information was not part of the original problem formulation!

| Arad | 366 | Mehadia | 241 |
| :--- | ---: | :--- | ---: |
| Bucharest | 0 | Neamt | 234 |
| Craiova | 160 | Oradea | 380 |
| Drobeta | 242 | Pitesti | 100 |
| Eforie | 161 | Rimnicu Vilcea | 193 |
| Fagaras | 176 | Sibiu | 253 |
| Giurgiu | 77 | Timisoara | 329 |
| Hirsova | 151 | Urziceni | 80 |
| Iasi | 226 | Vaslui | 199 |
| Lugoj | 244 | Zerind | 374 |



Is it the shortest path?

We already know that the greedy algorithm may not find the optimal path.

## Can we at least guarantee finding some path?

- If we expand first the node with the smallest cost then the (tree search) algorithm may not find any solution.
Example: path Iasi $\rightarrow$ Fagaras
- Go to Neamt, then back to Iasi, Neamt, ...
- We need to detect repeated visits in cities!
- Time complexity $\mathbf{O}\left(\mathbf{b}^{m}\right)$, where $m$ is the maximal depth
- Memory complexity $\mathbf{O}\left(\mathbf{b}^{m}\right)$
- A good heuristic function can significantly decrease the practical complexity.



## Let us now try to use $\mathbf{f}(\mathbf{n})=\mathbf{g}(\mathbf{n})+\mathbf{h ( n )}$

- recall that $\mathbf{g ( n )}$ is the cost of path from root to $\mathbf{n}$
- probably the most popular heuristic search algorithm
- $\mathbf{f}(\mathbf{n})$ represents the cost of path through $\mathbf{n}$
- the algorithm does not extend already long paths



## What about completeness and optimality of A*?

First a few definitions:

- admissible heuristic $\mathbf{h ( n )}$
- $\mathbf{h ( n )}$ $\leq$ "the cost of the cheapest path from $\mathbf{n}$ to goal"
- an optimistic view (the algorithm assumes a better cost than the real cost)
- function $\mathbf{f}(\mathbf{n})$ in $A^{*}$ is a lower estimate of the cost of path through $\mathbf{n}$
- monotonous (consistent) heuristic $\mathbf{h ( n )}$
- let $\mathbf{n}^{\prime}$ be a successor of $\mathbf{n}$ via action $\mathbf{a}$ and $\mathbf{c}\left(\mathbf{n}, \mathbf{a}, \mathbf{n}^{\prime}\right)$ be the transition cost
- $h(n) \leq c\left(n, a, n^{\prime}\right)+h\left(n^{\prime}\right)$
- this is a form of triangle inequality

Monotonous heuristic is admissible.
let $n_{1}, n_{2}, \ldots, n_{k}$ be the optimal path from $\mathbf{n}_{\mathbf{1}}$ to goal $\mathbf{n}_{\mathbf{k}}$, then $h\left(n_{i}\right)-h\left(n_{i+1}\right) \leq c\left(n_{i}, a_{i}, n_{i+1}\right)$, via monotony $h\left(n_{1}\right) \leq \sum_{i=1, ., k-1} c\left(n_{i}, a_{i}, n_{i+1}\right)$, after ,"sum"


For a monotonous heuristic the values of $f(n)$ are non-decreasing over any path.
Let $\mathbf{n}^{\mathbf{\prime}}$ be a successor of $\mathbf{n}$, i.e. $g\left(n^{\prime}\right)=g(n)+c\left(n, a, n^{\prime}\right)$, then $\left.\mathbf{f}^{\prime} \mathbf{n}^{\prime}\right)=\mathrm{g}\left(\mathrm{n}^{\prime}\right)+\mathrm{h}\left(\mathrm{n}^{\prime}\right)=\mathrm{g}(\mathrm{n})+\mathrm{c}\left(\mathrm{n}, \mathrm{a}, \mathrm{n}^{\prime}\right)+\mathrm{h}\left(\mathrm{n}^{\prime}\right) \geq \mathrm{g}(\mathrm{n})+\mathrm{h}(\mathrm{n})=\mathbf{f}(\mathbf{n})$

## If $h(n)$ is an admissible heuristic then the algorithm A* in TREE-SEARCH is optimal. $^{*}$

- in other words - the first expanded goal is optimal
- Let $\mathrm{G}_{2}$ be sub-optimal goal from the fringe and $\mathrm{C}^{*}$ be the optimal cost
- $f\left(G_{2}\right)=g\left(G_{2}\right)+h\left(G_{2}\right)=g\left(G_{2}\right)>C^{*}$, because $h\left(G_{2}\right)=0$
- Let $\mathbf{n}$ be a node from the fringe and being on the optimal path
- $f(n)=g(n)+h(n) \leq C^{*}$, via admissibility of $h(n)$
- together
- $\mathrm{f}(\mathrm{n}) \leq \mathrm{C}^{*}<\mathrm{f}\left(\mathrm{G}_{2}\right)$,
i.e., the algorithm must expand $\mathbf{n}$ before $\mathrm{G}_{2}$ and this way it finds the optimal path.



## If $h(n)$ is a monotonous heuristic then the algorithm A* in GRAPH-SEARCH is optimal.

- Possible problem: reaching the same state for the second time using a better path - classical GRAPHSEARCH ignores this second path!
- Possible solution: selection of the better of the two paths leading to the closed node (extra bookkeeping) or using monotonous heuristic.
- for monotonous heuristics, the values of $f(n)$ are non-decreasing over any path
- A* selects for expansion the node with the smallest value of $f(n)$, i.e., the values $f(m)$ of other open nodes $\mathbf{m}$ are not smaller, i.e., among all "open" paths to $\mathbf{n}$ there cannot be a shorter path than the path just selected (no path can shorten)
- hence, the first closed goal node is optimal

For non-decreasing function $f(n)$ we can draw contours in the state graph (the nodes inside a given contour have f-costs less than or equal to the contour value.

- for $h(n)=0$ we obtain circles around the start
- for more accurate $h(n)$ we use, the bands will stretch toward the goal state and become more narrowly focused around the optimal path.

- A* expands all nodes such that $\mathbf{f}(\mathbf{n})<\mathbf{C}^{*}$ on the contour
- A* can expand some nodes such that $\mathbf{f}(\mathbf{n})=\mathbf{C}^{*}$
- the nodes $\mathbf{n}$ such that $\mathbf{f}(\mathbf{n})>\mathbf{C}$ * are never expanded
- the algorithm A* is optimally efficient for any given consistent heuristic

Time complexity:
A* can expand an exponential number of nodes

- this can be avoided if $\left|h(n)-h^{*}(n)\right| \leq O\left(\log h^{*}(n)\right)$, where $h^{*}(n)$ is the cost of optimal path from $n$ to goal


## Space complexity:

A* keeps in memory all expanded nodes
A* usually runs out of space long before it runs out of time

## A simple way to decrease memory consumption is iterative deepening. <br> Algorithm IDA*

```
function IDA*(problem) returns a solution sequence
    inputs: problem, a problem
    static: f-limit, the current f}\mathrm{ -CosT limit
        root, a node
    root \leftarrow~MAKE-NODE(INITLAL-STATE[problem])
    f-limit }\leftarrowf\mathrm{ -COST(root)
    loop do
        solution,f-limit \leftarrow DFS-CONTOUR(root,f-limit)
        if solution is non-null then return solution
        if}\mathrm{ -limit = 
```

function DFS-CONTOUR(node, f-limit) returns a solution sequence and a new $f$-COST limit inputs: node, a node
$f$-limit, the current $f$-CosT limit
static: next-f, the $f$-CosT limit for the next contour, initially $\infty$
if $f$ - $\operatorname{CosT}[$ node $]>f$-limit then return null, $f$ - $\operatorname{CosT}[$ node $]$
if GOAL-TEST[problem](STATE%5Bnode%5D) then return node, f-limit
for each node $s$ in SUCCESSORS(node) do
solution, new-f $\leftarrow$ DFS-CONTOUR $(s$, -limit $)$
if solution is non-null then return solution, $f$-limit
next-f $\leftarrow \operatorname{MiN}($ next-f, new-f); end
return null, next- $f$

- the search limit is defined using the cost $\mathbf{f ( n )}$ instead of depth
- for the next iteration we use the smallest value $\mathbf{f}(\mathbf{n})$ of node $\mathbf{n}$ that exceeded the limit in the last iteration
- frequently used algorithm

Let us try to mimic standard best-first search, but using only linear space

- the algorithm stops exploration if there is an alternative path with better cost f(n)
- when the algorithm goes back to node $\mathbf{n}$, it replaces the value $f(n)$ using the cost of successors (remembers the best leaf in the forgotten subtree)
If $h(n)$ is an admissible heuristic then the algorithm is optimal.
- Space complexity O(bd)
- Time complexity is still exponential (suffers from excessive node re-generation)

```
function ReCurSIVE-BEST-FIRST-SEARCH(problem) returns a solution, or failure
    RBFS(problem, MAKE-NODE(InITIAL-STATE[problem]), }\infty\mathrm{ )
function RBFS(problem, node, f_limit) returns a solution, or failure and a new f}f\mathrm{ -cost limit
    if Goal-TEST[problem](STATE[node]) then return node
    successors }\leftarrow\mathrm{ EXPAND(node, problem)
    if successors is empty then return failure, }
    for each s in successors do
        f[s]\leftarrow\operatorname{max}(g(s)+h(s),f[node])
    repeat
        best }\leftarrow\mathrm{ the lowest f}\mathrm{ -value node in successors
        if f[best] > f_limit then return failure, f[best]
        alternative }\leftarrow\mathrm{ the second-lowest f}\mathrm{ -value among successors
        result, f[best]}\leftarrow\operatorname{RBFS}(\mathrm{ problem, best, min(f_limit, alternative))
        if result }\not=\mathrm{ failure then return result

\section*{1. After expansion of Arad, Sibia, Rimnicu Vilcea}

3. The path through Fagaras is now worse, go back Pitesti

2. The path from Rimnicu Vilcea now seems too expensive, go back to the closest neighbour Fagaras a more accurate cost is stored for Rinmicu Vilcea


\section*{IDA* and RBFS do not exploit available memory! This is a pity as the already expanded nodes are reexpanded again (waste of time) Let us try to modify classical A*}
```

function SMA*(problem) returns a solution sequence
inputs: problem, a problem
static: Queue, a queue of nodes ordered by f}f\mathrm{ -cost

```
    Queue \(\leftarrow\) Make-Queve \(\{\) Make-Node(Initial-State[problem])\})
    loop do
        if Queue is empty then return failure
        \(n \leftarrow\) deepest least- f -cost node in Queue
        if \(\operatorname{GoAL}-\operatorname{TeST}(n)\) then return success
        \(s \leftarrow \operatorname{NEXT}-\operatorname{SUCCESSOR}(n)\)
        if \(s\) is not a goal and is at maximum depth then
            \(\mathrm{f}(s) \leftarrow \infty\)
    else
            \(\mathrm{f}(s) \leftarrow \operatorname{MAX}(\mathrm{f}(n), \mathrm{g}(s)+\mathrm{h}(s))\)
    if all of \(n\) 's successors have been generated then
            update \(n\) 's \(f\)-cost and those of its ancestors if necessary
    if SUCCESSORS \((n)\) all in memory then remove \(n\) from Queue
    if memory is full then
            delete shallowest, highest-f-cost node in Queue
            remove it from its parent's successor list
            insert its parent on Queue if necessary
        insert \(s\) on Queue
    end
- when memory is full, drop the worst leaf node - the node with the highest \(f\)-value (if there are more such nodes then drop the shallowest node)
- similarly to RBFS back up the value of the forgotten node to its parent

- Assume memory for three nodes only.
- If there is enough memory to store an optimal path then SMA* finds an optimal solution.
- Otherwise it finds the best path with available memory.
- If the cost of J would be 19, then this is optimal goal, but the path to it can not be stored in memory!

\section*{Weighted A* (satisficing search)}

A* still expands a lot of nodes (to guarantee optimality). If we are willing to accept suboptimal solutions (good enough or satisficing solutions), we can explore fewer nodes. How? We allow inadmissible heuristics.

\section*{Weighted A* \(^{*}\)}
\[
f(n)=g(n)+W \times h(n) \text {, for some } W>1
\]

Finds solutions with the cost between \(\mathrm{C}^{*}\) and \(\mathrm{W} \times \mathrm{C}^{*}\) (in practice, the cost is closer to \(\mathrm{C}^{*}\) than to \(\mathrm{W} \times \mathrm{C}^{*}\) ).


\section*{How to find admissible heuristics?}

\section*{Example: 8-puzzle}
- 22 steps to goal in average
- branching factor around 3


Start State


Goal State
- (complete) search tree: \(3^{22} \approx 3,1 \times 10^{10}\) nodes
- the number of reachable states is only \(9!/ 2=181440\)
- for 15 -puzzle there are \(10^{13}\) states
- we need some heuristic, preferable admissible
\(-h_{1}=\) „the number of misplaced tiles"
\[
=8
\]
- \(\mathrm{h}_{2}=\) „the sum of the distances of the tiles from the goal positions"
\[
=3+1+2+2+2+3+3+2=18
\]
a so called Manhattan heuristic
- the optimal solution needs 26 steps

\section*{How to characterize the quality of a heuristic?}

Effective branching factor b*
- Let the algorithm need \(\mathbf{N}\) nodes to find a solution in depth \(\mathbf{d}\)
- \(\mathbf{b}^{*}\) is a branching factor of a uniform tree of depth \(\mathbf{d}\) containing \(\mathbf{N + 1}\) nodes
\[
N+1=1+b^{*}+\left(b^{*}\right)^{2}+\ldots+\left(b^{*}\right)^{d}
\]

\section*{Example:}
- 8-puzzle
- the average over 100 instances for each of various solution lengths
\begin{tabular}{|rrrrrrr|}
\hline \multicolumn{4}{c}{ Search Cost (nodes generated) } & \multicolumn{3}{c|}{ Effective Branching Factor } \\
\(d\) & BFS & \(\mathrm{A}^{*}\left(h_{1}\right)\) & \(\mathrm{A}^{*}\left(h_{2}\right)\) & BFS & \(\mathrm{A}^{*}\left(h_{1}\right)\) & \(\mathrm{A}^{*}\left(h_{2}\right)\) \\
\hline 6 & 128 & 24 & 19 & 2.01 & 1.42 & 1.34 \\
8 & 368 & 48 & 31 & 1.91 & 1.40 & 1.30 \\
10 & 1033 & 116 & 48 & 1.85 & 1.43 & 1.27 \\
12 & 2672 & 279 & 84 & 1.80 & 1.45 & 1.28 \\
14 & 6783 & 678 & 174 & 1.77 & 1.47 & 1.31 \\
16 & 17270 & 1683 & 364 & 1.74 & 1.48 & 1.32 \\
18 & 41558 & 4102 & 751 & 1.72 & 1.49 & 1.34 \\
20 & 91493 & 9905 & 1318 & 1.69 & 1.50 & 1.34 \\
22 & 175921 & 22955 & 2548 & 1.66 & 1.50 & 1.34 \\
24 & 290082 & 53039 & 5733 & 1.62 & 1.50 & 1.36 \\
26 & 395355 & 110372 & 10080 & 1.58 & 1.50 & 1.35 \\
28 & 463234 & 202565 & 22055 & 1.53 & 1.49 & 1.36 \\
\hline
\end{tabular}

\section*{Is \(\mathbf{h}_{\mathbf{2}}\) (from 8-puzzle) always better than \(\mathbf{h}_{\mathbf{1}}\) and how to recognize it?}
- notice that \(\forall n h_{2}(n) \geq h_{1}(n)\)
- we say that \(h_{2}\) dominates \(h_{1}\)
- A* with \(h_{2}\) never expands more nodes than A* with \(h_{1}\)
- A* expands all nodes such that \(f(n)<C^{*}\), so \(h(n)<C^{*}-g(n)\)
- In particular if it expands a node using \(h_{2}\), then the same node must be expanded using \(h_{1}\)

It is always better to use a heuristic function giving higher values provided that
- the limit \(\mathbf{C *}^{*} \mathbf{- g ( n )}\) is not exceeded (then the heuristic would not be admissible)
- the computation time is not too long

\section*{Can an agent construct admissible heuristics for any problem?}

\section*{Yes, via problem relaxation!}
- relaxation is a simplification of the problem such that the solution of the original problem is also a solution of the relaxed problem (even if not necessarily optimal)
- we need to be able to solve the relaxed problem fast
- the cost of optimal solution to a relaxed problem is a lower bound for the solution to the original problem and hence it is an admissible (and monotonous) heuristic for the original problem

\section*{Example (8-puzzle)}
- a tile can move from square \(A\) to square \(B\) if:
- \(A\) is horizontally or vertically adjacent to \(B\)
- \(B\) is blank
- possible relaxations (omitting some constraints to move a tile):
- a tile can move from square \(A\) to square \(B\) if \(A\) is adjacent to \(B\) (Manhattan distance)
- a tile can move from square \(A\) to square \(B\) if \(B\) is blank
- a tile can move from square \(A\) to square \(B\) (heuristic \(h_{1}\) )

Another approach to admissible heuristics is using a pattern database
- based on solution of specific sub-problems (patterns)
- by searching back from the goal and recording the cost of each new pattern encountered
- heuristic is defined by taking


Start State


Goal State the worst cost of a pattern that matches the current state
- Beware! The "sum" of costs of matching patterns needs not be admissible (the steps for solving one pattern may be used when solving another pattern).

If there are more heuristics, we can always use the maximum value from them (such a heuristic dominates each of the used heuristics).

© 2020 Roman Barták
Department of Theoretical Computer Science and Mathematical Logic
bartak@ktiml.mff.cuni.cz```

