Basics of AI and Machine Learning

State-Space Search: Heuristics

Jendrik Seipp

Linköping University

Slides modified from Basel AI group, with permission

Introduction

Informed Search Algorithms

- search algorithms considered so far: blind
  because they do not use any aspects of the problem to solve other than its formal definition (state space)
- problem: scalability
  $\rightsquigarrow$ prohibitive time and space requirements already for seemingly simple problems
- idea: try to find (problem-specific) criteria to distinguish good and bad states
  $\rightsquigarrow$ prefer good states

$\rightsquigarrow$ informed ("heuristic") search algorithms

Heuristics

Heuristics

Definition (heuristic)

Let $\mathcal{S}$ be a state space with states S.
A heuristic function or heuristic for $\mathcal{S}$ is a function

$$h : S \to \mathbb{R}_0^+ \cup \{\infty\},$$

mapping each state to a non-negative number (or ∞).

Heuristics: Intuition

idea: $h(s)$ estimates distance (= cost of cheapest path) from s to closest goal state

- heuristics can be arbitrary functions
- intuition: the closer h is to true goal distance, the more efficient the search using h

Representation of Heuristics

In our black box model, heuristics are an additional element of the state space interface:

State Spaces as Black Boxes (Extended)

- init()
- is_goal(s)
- succ(s)
- cost(a)
- h(s): heuristic value for state s
  result: non-negative integer or ∞

Examples

Example: Blocks World

possible heuristic:

count blocks x that currently lie on y
and must lie on $z \neq y$ in the goal
(including case where y or z is the table)

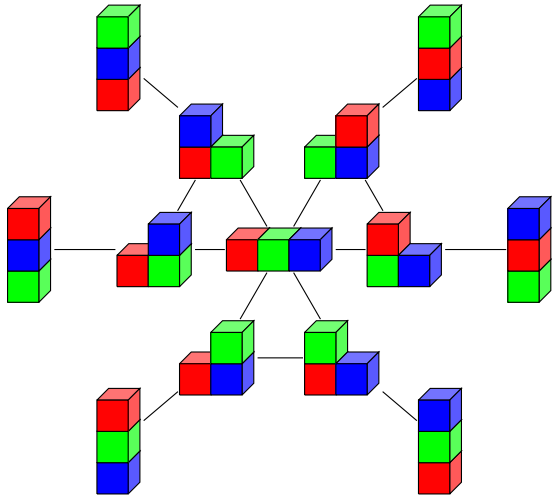

Example: Blocks World

possible heuristic:

count blocks x that currently lie on y
and must lie on $z \neq y$ in the goal
(including case where y or z is the table)

How accurate is this heuristic?

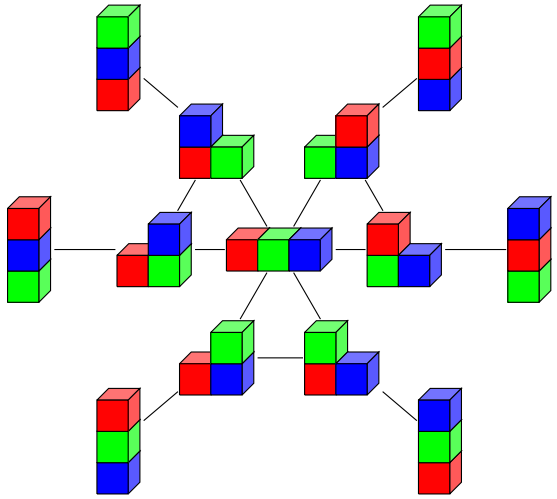

Example: Route Planning in Romania

possible heuristic: straight-line distance to Bucharest

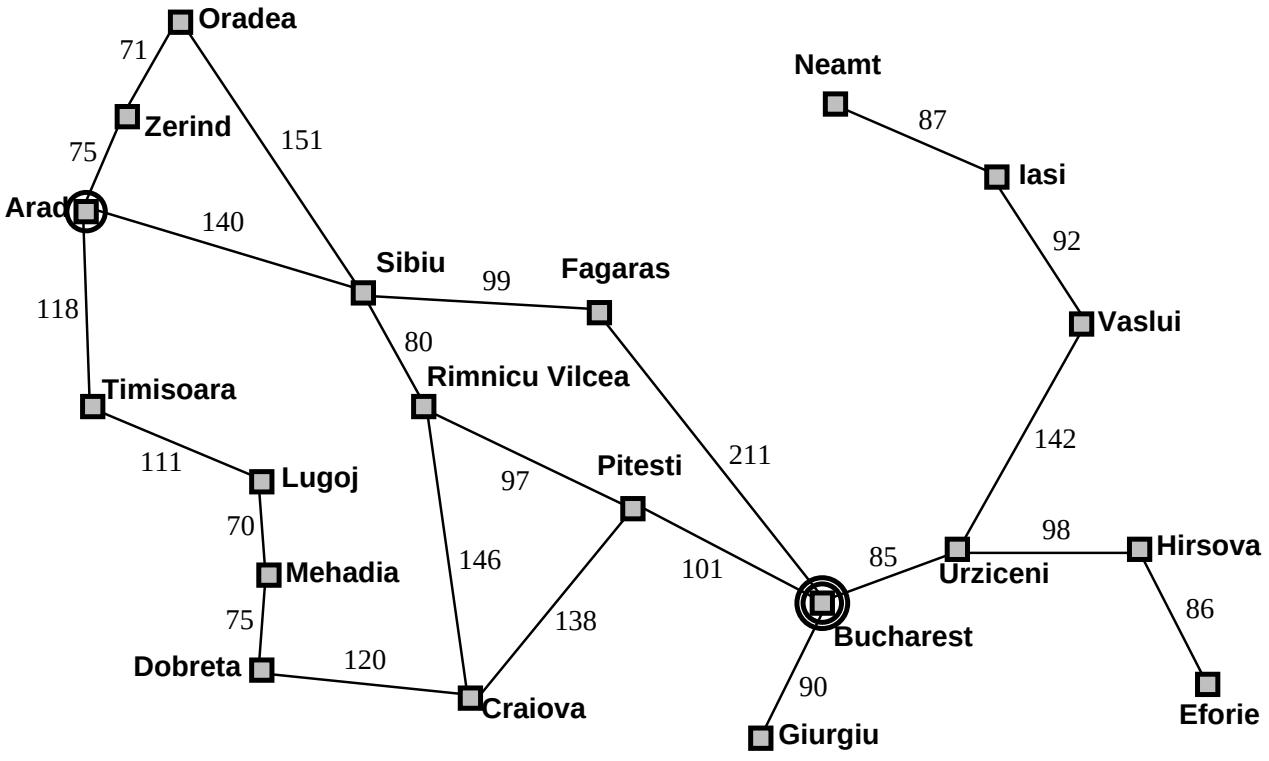

Arad	366
Bucharest	0
Craiova	160
Drobeta	242
Eforie	161
Fagaras	176
Giurgiu	77
Hirsova	151
Iasi	226
Lugoj	244
Mehadia	241
Neamt	234
Oradea	380
Pitesti	100
Rimnicu Vilcea	193
Sibiu	253
Timisoara	329
Urziceni	80
Vaslui	199
Zerind	374

Example: Missionaries and Cannibals

Setting: Missionaries and Cannibals

- Six people must cross a river.
- Their rowing boat can carry one or two people across the river at a time (it is too small for three).
- Three people are missionaries, three are cannibals.
- Missionaries may never stay with a majority of cannibals.

possible heuristic: number of people on the wrong river bank

Example: Missionaries and Cannibals

Setting: Missionaries and Cannibals

- Six people must cross a river.
- Their rowing boat can carry one or two people across the river at a time (it is too small for three).
- Three people are missionaries, three are cannibals.
- Missionaries may never stay with a majority of cannibals.

possible heuristic: number of people on the wrong river bank

$\rightsquigarrow$ with our formulation of states as triples $\langle m, c, b \rangle$:
$h(\langle m, c, b \rangle) = m + c$

Properties of Heuristics

Perfect Heuristic

Definition (perfect heuristic)

Let $\mathcal{S}$ be a state space with states S.

The perfect heuristic for $\mathcal{S}$, written h^*, maps each state $s \in S$ to the cost of an optimal solution for s.

remark: $h^*(s) = \infty$ if no solution for s exists

Properties of Heuristics

Definition (safe, goal-aware, admissible, consistent)

Let $\mathcal{S}$ be a state space with states S.

A heuristic h for $\mathcal{S}$ is called

- safe if $h^*(s) = \infty$ for all $s \in S$ with $h(s) = \infty$
- goal-aware if $h(s) = 0$ for all goal states s
- admissible if $h(s) \leq h^*(s)$ for all states $s \in S$
- consistent if $h(s) \leq cost(a) + h(s')$ for all transitions $s \xrightarrow{a} s'$

Examples

Properties of Heuristics: Examples

Which of our three example heuristics have which properties?

Route Planning in Romania

straight-line distance:

- safe
- goal-aware
- admissible
- consistent

Why?

Properties of Heuristics: Examples

Which of our three example heuristics have which properties?

Blocks World

misplaced blocks:

- safe?
- goal-aware?
- admissible?
- consistent?

Properties of Heuristics: Examples

Which of our three example heuristics have which properties?

Missionaries and Cannibals

people on wrong river bank:

- safe?
- goal-aware?
- admissible?
- consistent?

Connections

Properties of Heuristics: Connections

Theorem (admissible $\Longrightarrow$ safe + goal-aware)

Let h be an admissible heuristic.

Then h is safe and goal-aware.

Theorem (goal-aware + consistent $\Longrightarrow$ admissible)

Let h be a goal-aware and consistent heuristic.

Then h is admissible.

Summary

Summary

- heuristics estimate distance of a state to the goal
- can be used to focus search on promising states

$\rightsquigarrow$ soon: search algorithms that use heuristics

- perfect heuristic h^*: true cost to the goal
- important properties: safe, goal-aware, admissible, consistent
- connections between these properties
  - admissible $\Longrightarrow$ safe and goal-aware
  - goal-aware and consistent $\Longrightarrow$ admissible