Artificial Intelligence

Planning: Abstraction Heuristics

Jendrik Seipp

Linköping University

based on slides by Thomas Keller and Malte Helmert (University of Basel)

Heuristics

Reminder: Heuristics

Definition (heuristic)

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

A heuristic function or heuristic for $\mathcal{S}$ is a function

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

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

Reminder: 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

Reminder: 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 $s \xrightarrow{a} s'$

A Simple Planning Heuristic

The STRIPS planner (Fikes & Nilsson, 1971) uses the number of goals not yet satisfied in a STRIPS planning task as heuristic:

$$h(s) := |G \setminus s|.$$

intuition: fewer unsatisfied goals $\rightsquigarrow$ closer to goal state

$\rightsquigarrow$ STRIPS heuristic

Problems of STRIPS Heuristic

drawback of STRIPS heuristic?

- rather uninformed:
  for state s, if there is no applicable action a in s such that applying a in s satisfies strictly more (or fewer) goals,
  then all successor states have the same heuristic value as s
- very sensitive to reformulation:
  can easily transform any planning task into an equivalent one where $h(s) = 1$ for all non-goal states
- ignores almost the whole task structure:
  the heuristic values do not depend on the actions

$\rightsquigarrow$ we need better methods to design heuristics

Planning Heuristics

General Procedure for Obtaining a Heuristic

Solve a simplified version of the problem.

there are many ideas for domain-independent planning heuristics:

- abstraction $\rightsquigarrow$ now
- delete relaxation $\rightsquigarrow$ later
- landmarks
- critical paths
- network flows
- potential heuristics

Planning Heuristics

Abstraction: Idea

Estimate solution costs by considering a smaller planning task where states are merged into abstract states.

there are many ideas for domain-independent planning heuristics:

- abstraction $\rightsquigarrow$ now
- delete relaxation $\rightsquigarrow$ later
- landmarks
- critical paths
- network flows
- potential heuristics

Example: Logistics Task with One Package, Two Trucks

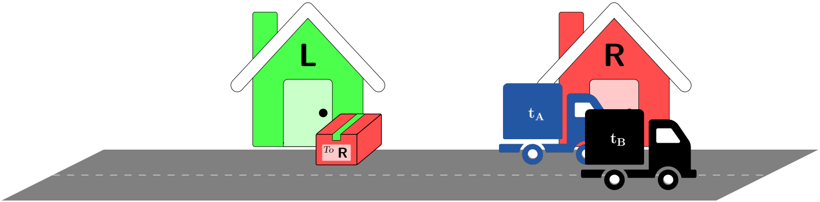

Example: Logistics Task with One Package, Two Trucks

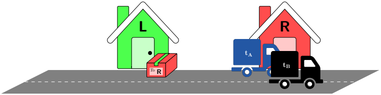

- $V = \{p, t_A, t_B\}$
- $dom(p) = \{L, R, A, B\}$ and $dom(t_A) = dom(t_B) = \{L, R\}$
- $I = \{p \mapsto L, t_A \mapsto R, t_B \mapsto R\}$
- $G = \{p \mapsto R\}$
- $A = \{load_{i,j} \mid i \in \{A, B\}, j \in \{L, R\}\}$
  $\cup \{unload_{i,j} \mid i \in \{A, B\}, j \in \{L, R\}\}$
  $\cup \{move_{i,j,j'} \mid i \in \{A, B\}, j, j' \in \{L, R\}, j \neq j'\}$ with:
  - $load_{i,j}$ has precondition $\{t_i \mapsto j, p \mapsto j\}$, effect $\{p \mapsto i\}$
  - $unload_{i,j}$ has precondition $\{t_i \mapsto j, p \mapsto i\}$, effect $\{p \mapsto j\}$
  - $move_{i,j,j'}$ has precondition $\{t_i \mapsto j\}$, effect $\{t_i \mapsto j'\}$
  - all actions have cost 1

State Space for Example Task

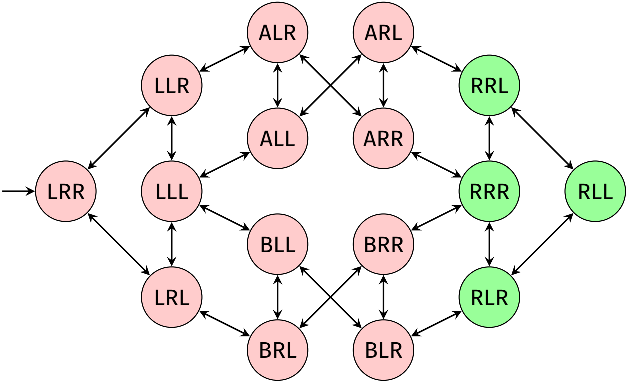

- state $\{p \mapsto i, t_A \mapsto j, t_B \mapsto k\}$ denoted as ijk
- annotations of edges not shown for simplicity
- for example, edge from LLL to ALL has annotation $\mathit{load}_{A,L}$

Abstractions

Abstraction

abstractions drop distinctions between certain states,
but preserve the state space behavior as well as possible.

- an abstraction of a state space $\mathcal{S}$ is defined by an abstraction function α that determines which states can be distinguished in the abstraction
- based on $\mathcal{S}$ and α, we compute the abstract state space $\mathcal{S}^{\alpha}$ which is "similar" to $\mathcal{S}$ but smaller

idea of the abstraction heuristic h^{α}:
use abstract solution costs (solution costs in $\mathcal{S}^{\alpha}$)
as heuristic values for concrete solution costs (solution costs in $\mathcal{S}$)

Induced Abstraction

Definition (induced abstraction)

Let $\mathcal{S} = \langle S, A, cost, T, s_I, S_\star \rangle$ be a state space, and let $\alpha : S \longrightarrow S'$ be a surjective function.

The abstraction of $\mathcal{S}$ induced by α, denoted as $\mathcal{S}^\alpha$, is the state space $\mathcal{S}^\alpha = \langle S', A, cost, T', s'_I, S'_\star \rangle$ with:

- $T' = \{\langle \alpha(s), a, \alpha(t) \rangle \mid \langle s, a, t \rangle \in T\}$
- $s'_I = \alpha(s_I)$
- $S'_\star = \{\alpha(s) \mid s \in S_\star\}$

Abstraction: Example

concrete state space

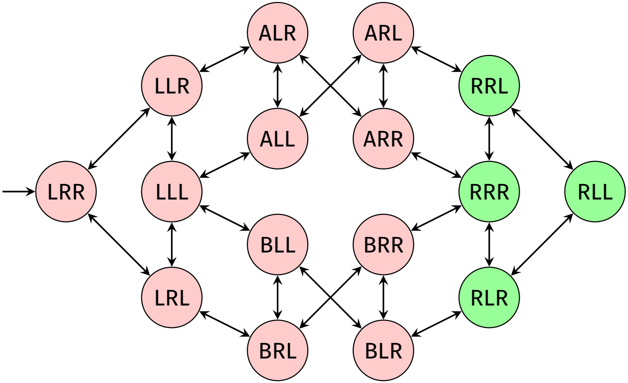

Abstraction: Example

(an) abstract state space

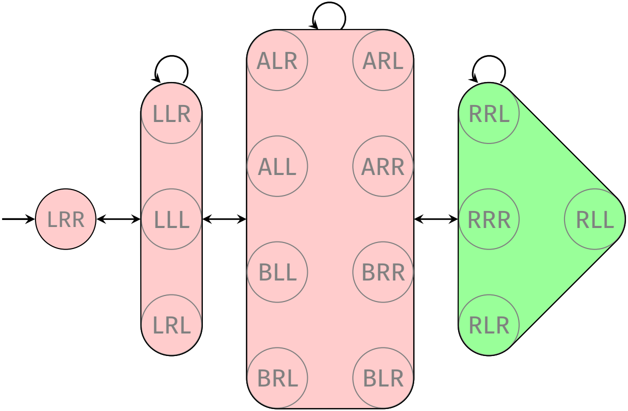

remark: most edges correspond to several (parallel) transitions with different annotations

Abstraction Heuristic: Example

(an) abstract state space

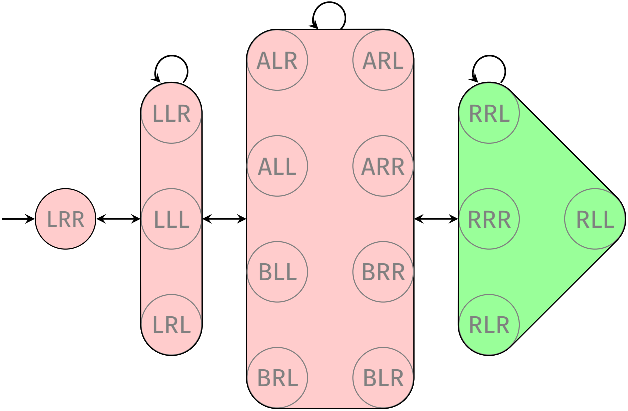

$h^\alpha(LRR) = ?,$
$h^\alpha(LLR) = ?,$
$h^\alpha(LLL) = ?,$
$h^\alpha(LRL) = ?,$
$h^\alpha(ALR) = ?,$
$\dots$

remark: most edges correspond to several (parallel) transitions with different annotations

Abstraction: Example

(an) abstract state space

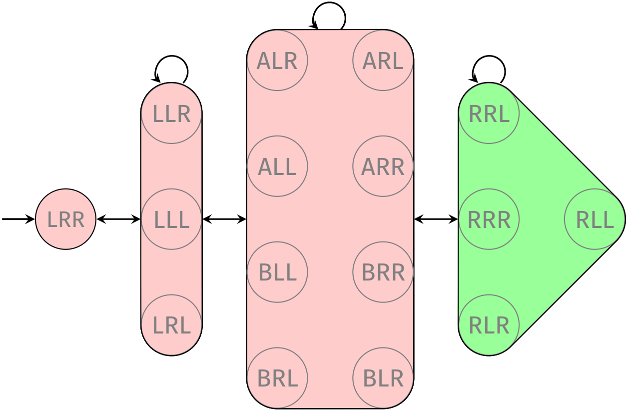

$h^\alpha(LRR) = 3,$
$h^\alpha(LLR) = 2,$
$h^\alpha(LLL) = 2,$
$h^\alpha(LRL) = 2,$
$h^\alpha(ALR) = 1,$
$\dots$

remark: most edges correspond to several (parallel) transitions with different annotations

Abstraction Heuristics: Discussion

- every abstraction heuristic is admissible and consistent
- the choice of the abstraction function α is very important
  - every α yields an admissible and consistent heuristic
  - but most α lead to poor heuristics
- a "good" α must yield an informative heuristic ...
- ...as well as being efficiently computable

How can we find a suitable α?

Usually a Bad Idea: Single-State Abstraction

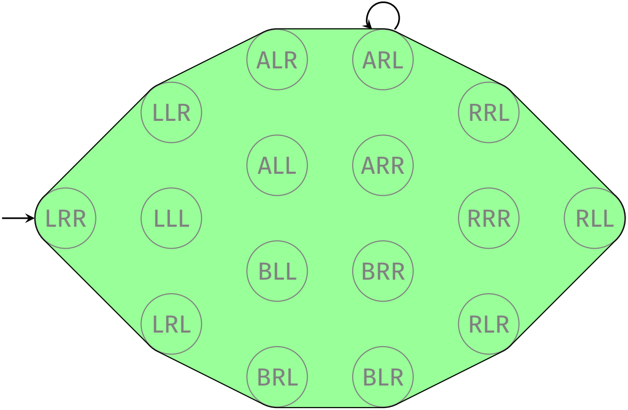

one state abstraction: $\alpha(s) := \text{const}$

Usually a Bad Idea: Single-State Abstraction

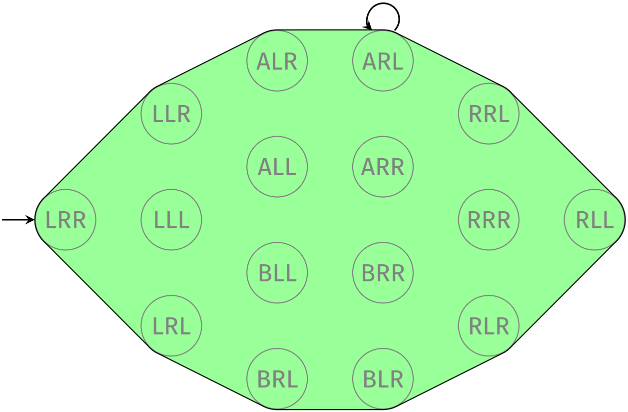

one state abstraction: $\alpha(s) :=$ const

- + compactly representable and α easy to compute
- − very uninformed heuristic

Usually a Bad Idea: Identity Abstraction

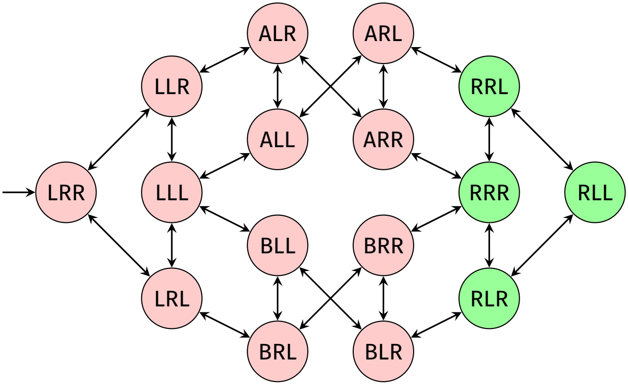

identity abstraction: $\alpha(s) := s$

Usually a Bad Idea: Identity Abstraction

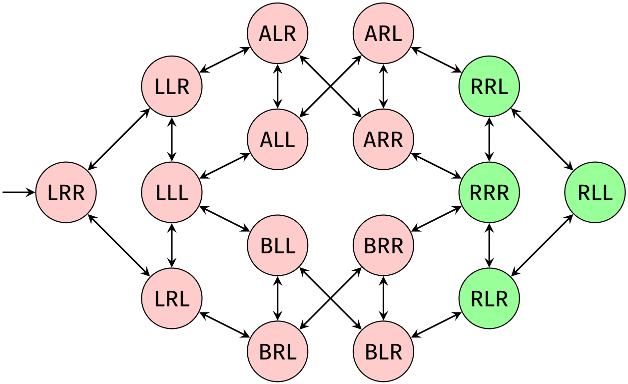

identity abstraction: $\alpha(s) := s$

- \+ perfect heuristic and α easy to compute
- – too many abstract states $\rightsquigarrow$ computation of h^α too hard

Automatic Computation of Suitable Abstractions

Main Problem with Abstraction Heuristics

How to find a good abstraction?

several successful methods:

- pattern databases (PDBs) (Culberson & Schaeffer, 1996)
- domain abstractions (Hernádvölgyi and Holte, 2000)
- merge-and-shrink abstractions (Dräger, Finkbeiner & Podelski, 2006)
- Cartesian abstractions (Seipp & Helmert, 2013)

Pattern Databases

Pattern Databases: Background

- the most common abstraction heuristics are pattern database heuristics
- originally introduced for the 15-puzzle (Culberson & Schaeffer, 1996) and for Rubik's Cube (Korf, 1997)
- introduced for automated planning by Edelkamp (2001)
- for many search problems the best known heuristics
- many research papers studying
  - theoretical properties
  - efficient implementation and application
  - pattern selection
  - ...

Pattern Databases: Projections

a PDB heuristic is an abstraction heuristic where

- some aspects (= state variables) of the task are preserved with perfect precision while
- all other aspects are not preserved at all

formalized as projections on a pattern P, e.g.:

- $s = \{v_1 \mapsto d_1, v_2 \mapsto d_2, v_3 \mapsto d_3\}$
- projection on $P = \{v_1\}$ (= ignore v_2, v_3):
  $\alpha(s) = s|_P = \{v_1 \mapsto d_1\}$
- projection on $P = \{v_1, v_3\}$ (= ignore v_2):
  $\alpha(s) = s|_P = \{v_1 \mapsto d_1, v_3 \mapsto d_3\}$

Pattern Databases: Definition

Definition (pattern database heuristic)

Let P be a subset of the variables of a planning task.

The abstraction heuristic induced by the projection π_P on P is called pattern database heuristic (PDB heuristic) with pattern P.

abbreviated notation: h^P for h^{π_P}

remark: "pattern databases" in analogy to endgame databases (which have been successfully applied in 2-person-games)

Example: Concrete State Space

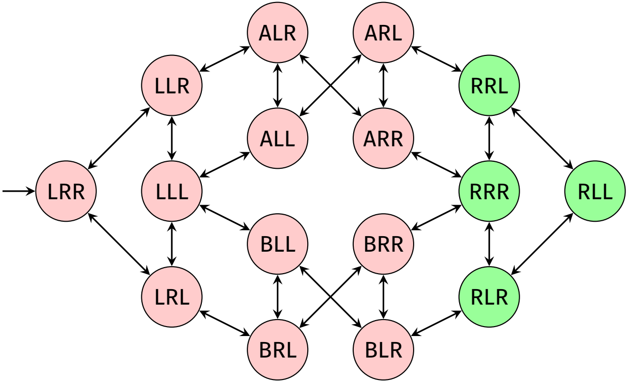

- state variable for package, p: $\{$L, R, A, B$\}$
- state variable for truck A, t_A: $\{$L, R$\}$
- state variable for truck B, t_B: $\{$L, R$\}$

Example: Projection (1)

abstraction $\mathcal{S}^{\pi_{\{p\}}}$ induced by $\pi_{\{p\}}$:

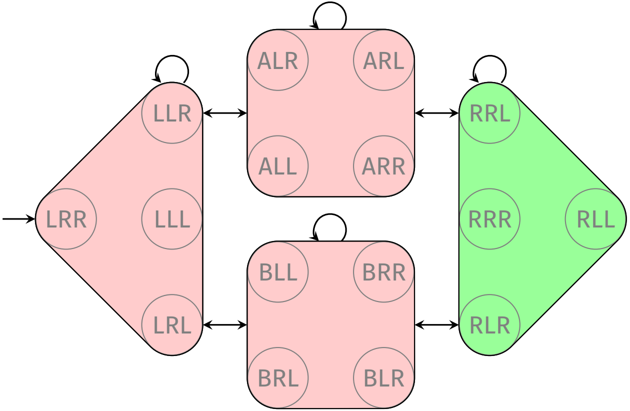

$$h^{\{p\}}(LRR) = 2,$$
$$h^{\{p\}}(LLR) = 2,$$
$$h^{\{p\}}(LLL) = 2,$$
$$h^{\{p\}}(LRL) = 2,$$
$$h^{\{p\}}(ALR) = 1,$$
$$\ldots$$

Example: Projection (1)

abstraction $\mathcal{S}^{\pi_{\{p\}}}$ induced by $\pi_{\{p\}}$:

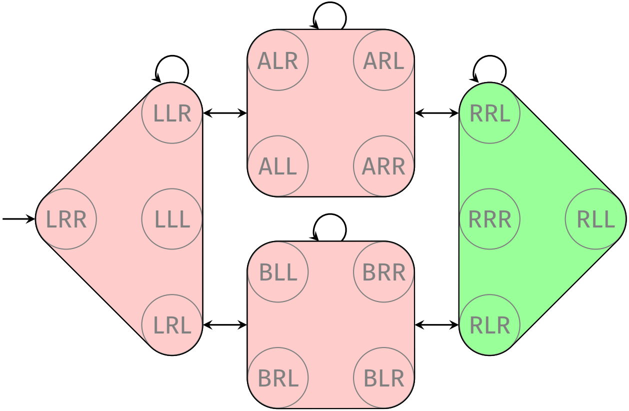

PDB :

$\{p \mapsto L\} = 2,$

$\{p \mapsto A\} = 1,$

$\{p \mapsto B\} = 1,$

$\{p \mapsto R\} = 0$

Example: Projection (2)

abstraction $\mathcal{S}^{\pi_{\{p,t_A\}}}$ induced by $\pi_{\{p,t_A\}}$:

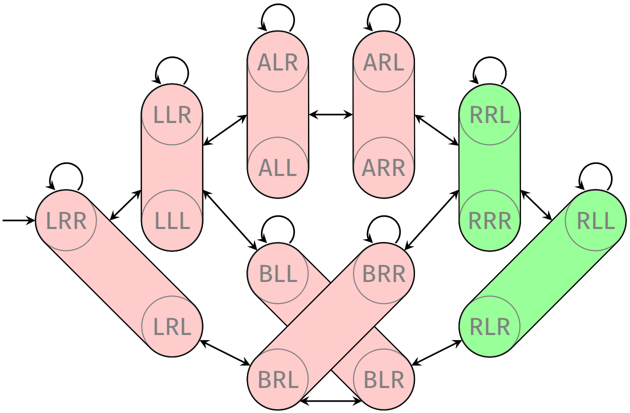

$h^{\{p,t_A\}}(LRR) = 2,$
$h^{\{p,t_A\}}(LLR) = 2,$
$h^{\{p,t_A\}}(LLL) = 2,$
$h^{\{p,t_A\}}(LRL) = 2,$
$h^{\{p,t_A\}}(ALR) = 2,$
$\ldots$

Example: Projection (2)

abstraction $\mathcal{S}^{\pi_{\{p,t_A\}}}$ induced by $\pi_{\{p,t_A\}}$:

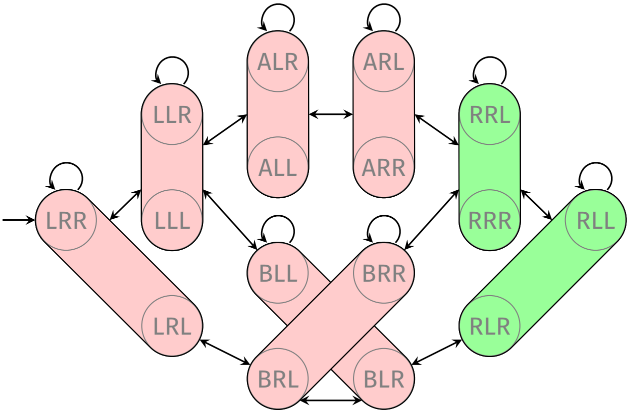

PDB :

$\{p \mapsto L, t_A \mapsto L\} = 2,$

$\{p \mapsto L, t_A \mapsto R\} = 2,$

$\{p \mapsto A, t_A \mapsto L\} = 2,$

$\{p \mapsto A, t_A \mapsto R\} = 1,$

$\{p \mapsto B, t_A \mapsto L\} = 1,$

$\{p \mapsto B, t_A \mapsto R\} = 1,$

$\{p \mapsto R, t_A \mapsto L\} = 0,$

$\{p \mapsto R, t_A \mapsto R\} = 0$

Pattern Databases in Practice

practical aspects which we do not discuss here:

- How to automatically find good patterns?
- How to combine multiple PDB heuristics?
- How to implement PDB heuristics efficiently?
  - good implementations efficiently handle abstract state spaces with 10^7, 10^8 or more abstract states
  - effort independent of the size of the concrete state space
  - usually all heuristic values are precomputed
    $\rightsquigarrow$ space complexity = number of abstract states

Pattern Databases Example

Consider the following state space:

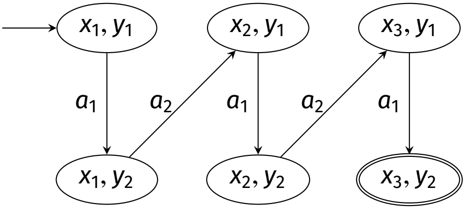

Draw the abstract state space that is induced by the projection $\pi_{\{X\}}$ to the variable X. Give the heuristic value of the pattern database heuristic $h^{\{X\}}$ for each state in the original state space. A state label of x, y corresponds to state $\{X \mapsto x, Y \mapsto y\}$ and all actions have a cost of 1.

Pattern Databases Example: Solution

The abstract state space is

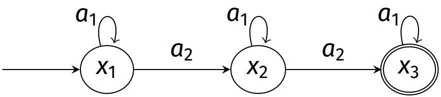

The heuristic function is

$$h^{\{X\}}(x_1y_1) = 2 \qquad h^{\{X\}}(x_2y_1) = 1 \qquad h^{\{X\}}(x_3y_1) = 0$$

$$h^{\{X\}}(x_1y_2) = 2 \qquad h^{\{X\}}(x_2y_2) = 1 \qquad h^{\{X\}}(x_3y_2) = 0$$

Summary

planning formalisms:

- STRIPS: particularly simple, easy to handle for algorithms
  - binary state variables
  - preconditions, add and delete effects, goals: sets of variables
- SAS^+: extension of STRIPS
  - state variables with arbitrary finite domains

abstraction heuristics:

- estimate solution cost by considering a smaller planning task
- Pattern database heuristics are abstraction heuristics based on projections onto state variable subsets (patterns): states are distinguishable iff they differ on the pattern.