

Principles of AI Planning

11. Planning as search: pattern database heuristics

Albert-Ludwigs-Universität Freiburg



Bernhard Nebel and Robert Mattmüller
December 18th, 2019



Pattern databases heuristics

- PDB heuristics
- Projections
- Examples
- Overview
- Implementing PDBs
- Additivity
- Pattern selection
- Summary

December 18th, 2019

B. Nebel, R. Mattmüller – AI Planning

2 / 52

Pattern database heuristics



- The most commonly used abstraction heuristics in search and planning are **pattern database (PDB) heuristics**.
- PDB heuristics were originally introduced for the **15-puzzle** (Culberson & Schaeffer, 1996) and for **Rubik's cube** (Korf, 1997).
- The first use for **domain-independent planning** is due to Edelkamp (2001).
- Since then, much research has focused on the theoretical properties of pattern databases, how to use pattern databases more effectively, how to find good patterns, etc.

- PDB heuristics
- Projections
- Examples
- Overview
- Implementing PDBs
- Additivity
- Pattern selection
- Summary

December 18th, 2019

B. Nebel, R. Mattmüller – AI Planning

3 / 52

Pattern database heuristics informally



Pattern databases: informally

A pattern database heuristic for a planning task is an abstraction heuristic where

- some aspects of the task are represented in the abstraction **with perfect precision**, while
- all other aspects of the task are **not represented at all**.

Example (15-puzzle)

- Choose a subset T of tiles (the **pattern**).
- Faithfully represent the locations of T in the abstraction.
- Assume that all other tiles and the blank can be anywhere in the abstraction.

- PDB heuristics
- Projections
- Examples
- Overview
- Implementing PDBs
- Additivity
- Pattern selection
- Summary

December 18th, 2019

B. Nebel, R. Mattmüller – AI Planning

4 / 52

Formally, pattern database heuristics are induced abstractions of a particular class of homomorphisms called **projections**.

Definition (projections)

Let Π be an FDR planning task with variable set V and state set S . Let $P \subseteq V$, and let S' be the set of states over P .

The **projection** $\pi_P : S \rightarrow S'$ is defined as $\pi_P(s) := s|_P$ (with $s|_P(v) := s(v)$ for all $v \in P$).

We call P the **pattern** of the projection π_P .

In other words, π_P maps two states s_1 and s_2 to the same abstract state iff they agree on all variables in P .

- PDB heuristics
- Projections
- Examples
- Overview
- Implementing PDBs
- Additivity
- Pattern selection
- Summary

Abstraction heuristics for projections are called **pattern database (PDB) heuristics**.

Definition (pattern database heuristic)


The abstraction heuristic induced by π_P is called a **pattern database heuristic** or **PDB heuristic**.

We write h^P as a short-hand for h^{π_P} .

Why are they called **pattern database heuristics**?

- Heuristic values for PDB heuristics are traditionally stored in a 1-dimensional table (array) called a **pattern database (PDB)**. Hence the name “PDB heuristic”.

- PDB heuristics
- Projections
- Examples
- Overview
- Implementing PDBs
- Additivity
- Pattern selection
- Summary




- PDB heuristics
- Projections
- Examples
- Overview
- Implementing PDBs
- Additivity
- Pattern selection
- Summary

Logistics problem with one package, two trucks, two locations:

- state variable **package**: $\{L, R, A, B\}$
- state variable **truck A**: $\{L, R\}$
- state variable **truck B**: $\{L, R\}$

Abstraction induced by $\pi_{\{\text{package}\}}$:



$$h^{\{\text{package}\}}(\text{LRR}) = 2$$

- PDB heuristics
- Projections
- Examples
- Overview
- Implementing PDBs
- Additivity
- Pattern selection
- Summary

Precomputation step



Let Π be a planning task and P a pattern.

Let $\mathcal{T} = \mathcal{T}(\Pi)$ and $\mathcal{T}' = \mathcal{T}^{\pi_P}$.

- We want to compute a graph representation of \mathcal{T}' .
- \mathcal{T}' is defined through a homomorphism of \mathcal{T} .
 - For example, each concrete transition induces an abstract transition.
- However, we cannot **compute** \mathcal{T}' by iterating over all transitions of \mathcal{T} .
 - This would take time $\Omega(\|\mathcal{T}\|)$.
 - This is prohibitively large (or else we could solve the task using breadth-first search or similar techniques).
- Hence, we need a way of computing \mathcal{T}' in time which is **polynomial only in $\|\Pi\|$ and $\|\mathcal{T}'\|$** .

PDB
heuristics
Implementing
PDBs
Precomputation
Lookup
Additivity
Pattern
selection
Summary

Syntactic projections



Definition (syntactic projection)

Let $\Pi = \langle V, I, O, \gamma \rangle$ be an FDR planning task, and let $P \subseteq V$ be a subset of its variables.

The **syntactic projection** $\Pi|_P$ of Π to P is the FDR planning task $\langle P, I|_P, \{o|_P \mid o \in O\}, \gamma|_P \rangle$, where

- $\varphi|_P$ for formula φ is defined as the formula obtained from φ by replacing all atoms $(v = d)$ with $v \notin P$ by \top , and
- $o|_P$ for operator o is defined by replacing all formulas φ occurring in the precondition or effect conditions of o with $\varphi|_P$ and all atomic effects $(v := d)$ with $v \notin P$ with the empty effect \top .

Put simply, $\Pi|_P$ throws away all information not pertaining to variables in P .

PDB
heuristics
Implementing
PDBs
Precomputation
Lookup
Additivity
Pattern
selection
Summary

Trivially inapplicable operators



Definition (trivially inapplicable operator)

An operator $\langle \chi, e \rangle$ of a SAS⁺ task is called **trivially inapplicable** if

- χ contains the atoms $(v = d)$ and $(v = d')$ for some variable v and values $d \neq d'$, or
- e contains the effects $(v := d)$ and $(v := d')$ for some variable v and values $d \neq d'$.

Notes:

- Trivially inapplicable operators are never applicable and can thus be safely omitted from the task.
- Trivially inapplicable operators can be detected in linear time.

PDB
heuristics
Implementing
PDBs
Precomputation
Lookup
Additivity
Pattern
selection
Summary

Trivially unsolvable SAS⁺ tasks



Definition (trivially unsolvable SAS⁺ tasks)

A SAS⁺ task $\Pi = \langle V, I, O, \gamma \rangle$ is called **trivially unsolvable** if γ contains the atoms $(v = d)$ and $(v = d')$ for some variable v and values $d \neq d'$.

Notes:

- Trivially unsolvable SAS⁺ tasks have no goal states, and are hence unsolvable.
- Trivially unsolvable SAS⁺ tasks can be detected in linear time.

PDB
heuristics
Implementing
PDBs
Precomputation
Lookup
Additivity
Pattern
selection
Summary

Theorem (syntactic projections vs. projections)

Let Π be a SAS⁺ task that is not trivially unsolvable and has no trivially inapplicable operators, and let P be a pattern for Π .

Then $\mathcal{T}(\Pi|_P) \stackrel{G}{\sim} \mathcal{T}(\Pi)^{\pi_P}$. □

Note: The restrictions to SAS⁺ tasks and to tasks without trivially inapplicable operators are necessary.

Using the equivalence theorem, we can compute pattern databases for (not trivially unsolvable) SAS⁺ tasks Π and patterns P :

Computing pattern databases

def compute-PDB(Π, P):

Remove trivially inapplicable operators from Π .

Compute $\Pi' := \Pi|_P$.

Compute $\mathcal{T}' := \mathcal{T}(\Pi')$.

Perform a backward breadth-first search from the goal states of \mathcal{T}' to compute all abstract goal distances.

$PDB :=$ a table containing all goal distances in \mathcal{T}'

return PDB

The algorithm runs **in polynomial time and space** in terms of $\|\Pi\| + |PDB|$.

- Most practical implementations of PDB heuristics are limited to SAS⁺ tasks (or modest generalizations).
- One way to avoid the issues with general FDR tasks is to convert them to equivalent SAS⁺ tasks.
- However, most direct conversions can exponentially increase the task size in the worst case.

↪ We will only consider SAS⁺ tasks in this chapter.

- During search, the PDB is the only piece of information necessary to represent h^P . (It is not necessary to store the abstract transition system itself at this point.)
- Hence, the space requirements for PDBs during search are linear in the number of abstract states S' : there is one table entry for each abstract state.
- During search, $h^P(s)$ is computed by mapping $\pi_P(s)$ to a natural number in the range $\{0, \dots, |S'| - 1\}$ using a **perfect hash function**, then looking up the table entry for that number.

- The space requirements for a pattern database grow **exponentially** with the **number of state variables** in the pattern.
- This places severe limits on the usefulness of single PDB heuristics h^P for larger planning task.
- To overcome this limitation, planners using pattern databases work with **collections of multiple patterns**.
- When using two patterns P_1 and P_2 , it is always possible to use the **maximum** of h^{P_1} and h^{P_2} as an admissible and consistent heuristic estimate.
- However, when possible, it is much preferable to use the **sum** of h^{P_1} and h^{P_2} as a heuristic estimate, since $h^{P_1} + h^{P_2} \geq \max\{h^{P_1}, h^{P_2}\}$.

Theorem (additive pattern sets)

Let P_1, \dots, P_k be patterns for an FDR planning task Π .

If there exists no operator that has an effect on a variable $v_i \in P_i$ and on a variable $v_j \in P_j$ for some $i \neq j$, then $\sum_{i=1}^k h^{P_i}$ is an admissible and consistent heuristic for Π .

Proof.

If there exists no such operator, then no label of $\mathcal{T}(\Pi)$ affects both $\mathcal{T}(\Pi)^{\pi_{P_i}}$ and $\mathcal{T}(\Pi)^{\pi_{P_j}}$ for $i \neq j$. By the theorem on affecting transition labels, this means that any two projections π_{P_i} and π_{P_j} are orthogonal. The claim follows with the theorem on additivity for orthogonal abstraction mappings.

A pattern set $\{P_1, \dots, P_k\}$ which satisfies the criterion of the theorem is called an **additive pattern set** or **additive set**.

Finding additive pattern sets

The theorem on additive pattern sets gives us a simple criterion to decide which pattern heuristics can be admissibly added.

Given a **pattern collection** \mathcal{C} (i. e., a set of patterns), we can use this information as follows:

- 1 Build the **compatibility graph** for \mathcal{C} .
 - Vertices correspond to patterns $P \in \mathcal{C}$.
 - There is an edge between two vertices iff no operator affects both incident patterns.
- 2 Compute **all maximal cliques** of the graph. These correspond to maximal additive subsets of \mathcal{C} .
 - Computing large cliques is an NP-hard problem, and a graph can have exponentially many maximal cliques.
 - However, there are **output-polynomial** algorithms for finding all maximal cliques (Tomita, Tanaka & Takahashi, 2004) which have led to good results in practice.

The canonical heuristic function

Definition (canonical heuristic function)

Let Π be an FDR planning task, and let \mathcal{C} be a pattern collection for Π .

The **canonical heuristic** $h^{\mathcal{C}}$ for pattern collection \mathcal{C} is defined as

$$h^{\mathcal{C}}(s) = \max_{\mathcal{D} \in \text{cliques}(\mathcal{C})} \sum_{P \in \mathcal{D}} h^P(s),$$

where $\text{cliques}(\mathcal{C})$ is the set of all maximal cliques in the compatibility graph for \mathcal{C} .

For all choices of \mathcal{C} , heuristic $h^{\mathcal{C}}$ is admissible and consistent.

How good is the canonical heuristic function?



- The canonical heuristic function is the **best possible** admissible heuristic we can derive from \mathcal{C} using **our additivity criterion**.
- In theory, even better heuristic estimates can be obtained from projection heuristics using a **more general additivity criterion** based on an idea called **cost partitioning**.
 - Optimal polynomial cost partitioning algorithms exist (Katz & Domshlak, 2008a).

PDB
heuristics
Implemen-
ting
PDBs
Additivity
Pattern
selection
Summary

Canonical heuristic function: example



Example

Consider a planning task with state variables $V = \{v_1, v_2, v_3\}$ and the pattern collection $\mathcal{C} = \{P_1, \dots, P_4\}$ with $P_1 = \{v_1, v_2\}$, $P_2 = \{v_1\}$, $P_3 = \{v_2\}$ and $P_4 = \{v_3\}$.

There are operators affecting each individual variable, and the only operators affecting several variables affect v_1 and v_3 .

What are the maximal cliques in the compatibility graph for \mathcal{C} ?

Answer: $\{P_1\}, \{P_2, P_3\}, \{P_3, P_4\}$

What is the canonical heuristic function $h^{\mathcal{C}}$?

Answer:
$$h^{\mathcal{C}} = \max \{h^{P_1}, h^{P_2} + h^{P_3}, h^{P_3} + h^{P_4}\}$$
$$= \max \{h^{\{v_1, v_2\}}, h^{\{v_1\}} + h^{\{v_2\}}, h^{\{v_2\}} + h^{\{v_3\}}\}$$

PDB
heuristics
Implemen-
ting
PDBs
Additivity
Pattern
selection
Summary

Pattern selection



PDB
heuristics
Implemen-
ting
PDBs
Additivity
Pattern
selection
Local search
Search space
Estimating
heuristic quality
Summary

Pattern selection as an optimization problem



Only one question remains to be answered now in order to apply PDBs to planning tasks in practice:

How do we automatically find a good pattern collection?

The idea

Pattern selection can be cast as an **optimization problem**:

- **Given:** a set of **candidate solutions**
(= pattern collections which fit into a given memory limit)
- **Find:** a **best possible** solution, or an approximation
(= pattern collection with high heuristic quality)

PDB
heuristics
Implemen-
ting
PDBs
Additivity
Pattern
selection
Local search
Search space
Estimating
heuristic quality
Summary

Pattern selection as local search



How to solve this optimization problem?

- For problems of interesting size, we cannot hope to find (and prove) a **globally optimal** pattern collection.
 - **Question:** How many candidates are there?
- Instead, we try to find **good** solutions by **local search**.

Two approaches from the literature:

- Edelkamp (2007): using **evolutionary algorithm**
- Haslum et al. (2007): using **hill-climbing**

↪ we present the main ideas of the second approach here

PDB
heuristics
Implementing
PDBs
Additivity
Pattern
selection
Local search
Search space
Estimating
heuristic quality
Summary

Pattern selection as hill-climbing



Reminder: Hill-climbing

$\sigma := \text{make-root-node}(\text{init}())$

forever:

if **is-goal**(state(σ)):

return extract-solution(σ)

$\Sigma' := \{ \text{make-node}(\sigma, o, s) \mid \langle o, s \rangle \in \text{succ}(\text{state}(\sigma)) \}$

$\sigma :=$ an element of Σ' minimizing h (random tie breaking)

Four questions to answer to use this for pattern selection:

- 1 **init:** What is the initial pattern collection?
- 2 **is-goal:** When do we terminate?
- 3 **succ:** Which collections are neighbours of the current collection?
- 4 **h:** How do we rank the quality of pattern collections?

PDB
heuristics
Implementing
PDBs
Additivity
Pattern
selection
Local search
Search space
Estimating
heuristic quality
Summary

Search space



We first discuss the **search space** (init, is-goal, succ).

The basic idea is that we

- start from **small patterns** of only a single variable each,
- grow them by **adding slightly larger patterns**, and
- stop when **heuristic quality no longer improves**.

To motivate the precise definition of our search space, we need a little more theory.

PDB
heuristics
Implementing
PDBs
Additivity
Pattern
selection
Local search
Search space
Estimating
heuristic quality
Summary

Initial pattern collection



Theorem (non-goal patterns are trivial)

Let Π be a SAS⁺ planning task that is not trivially unsolvable, and let P be a pattern for Π such that no variable in P is mentioned in the goal formula of Π . Then $h^P(s) = 0$ for all states s .

Proof.

All states in the abstraction are goal states.

This motivates our first answer:

1. **What is the initial pattern collection?**

The initial pattern collection is

$\{ \{v\} \mid v \text{ is a state variable mentioned in the goal formula} \}$.

PDB
heuristics
Implementing
PDBs
Additivity
Pattern
selection
Local search
Search space
Estimating
heuristic quality
Summary

Our second question has a very simple answer:

2. When do we terminate?

We terminate as soon as the current pattern collection has no successors of better quality.

Note that this also covers the case where there are no successors at all because further growth of the current pattern collection would exceed a memory limit.

Our search neighbourhood is defined through **incremental growth** of the current pattern collection.

A successor is obtained by

- starting from the **current pattern collection** \mathcal{C} ,
- choosing **one of its patterns** $P \in \mathcal{C}$ (without removing it from \mathcal{C} !),
- generating a new pattern by extending P with a single variable ($P' = P \cup \{v\}$), and
- adding P' to \mathcal{C} to form the new pattern collection \mathcal{C}'

However, not all such collections \mathcal{C}' are useful.

Definition (causal graph)

Let $\Pi = \langle V, I, O, \gamma \rangle$ be an FDR planning task.

The **causal graph** of Π , $CG(\Pi)$, is the directed graph with vertex set V and an arc from $u \in V$ to $v \in V$ iff $u \neq v$ and there exists an operator $o \in O$ such that:

- u appears anywhere in o (in precondition, effect conditions or atomic effects), and
- v is modified by an effect of o .

Idea: an arc $\langle u, v \rangle$ in the causal graph indicates that variable u is in some way relevant for modifying the value of v

Definition (causally relevant variables)

Let $\Pi = \langle V, I, O, \gamma \rangle$ be an FDR planning task and let $P \subseteq V$ be a pattern for Π .

We say that $v \in P$ is **causally relevant for P** if $CG(\Pi)$ contains a directed path from v to a variable $v' \in P$ that is mentioned in the goal formula γ .

Note: The definition implies that variables in P mentioned in the goal are always causally relevant for P .

Causally irrelevant variables are useless



- PDB heuristics
- Implementing PDBs
- Additivity
- Pattern selection
- Local search
- Search space
- Estimating heuristic quality
- Summary

Theorem (causally irrelevant variables are useless)

Let $P \subseteq V$ be a pattern for an FDR planning task Π , and let $P' \subseteq P$ consist of all variables that are causally relevant for P . Then $h^{P'}(s) = h^P(s)$ for all states s . \square

Corollary: There is no point in growing a pattern by adding a variable that is causally irrelevant in the resulting pattern.

Causally connected patterns



- PDB heuristics
- Implementing PDBs
- Additivity
- Pattern selection
- Local search
- Search space
- Estimating heuristic quality
- Summary

Definition (causally connected patterns)

Let $\Pi = \langle V, I, O, \gamma \rangle$ be an FDR planning task and let $P \subseteq V$ be a pattern for Π .

We say that P is **causally connected** if the subgraph of $CG(\Pi)$ induced by P is weakly connected (i. e., contains a path from every vertex to every other vertex, ignoring arc directions).

Disconnected patterns are decomposable



- PDB heuristics
- Implementing PDBs
- Additivity
- Pattern selection
- Local search
- Search space
- Estimating heuristic quality
- Summary

Theorem (causally disconnected patterns are decomposable)

Let $P \subseteq V$ be a pattern for a SAS⁺ planning task Π that is not causally connected, and let P_1, P_2 be a partition of P into non-empty subsets such that $CG(\Pi)$ contains no arc between the two sets.

Then $h^{P_1}(s) + h^{P_2}(s) = h^P(s)$ for all states s . \square

Corollary: There is no point in including a causally disconnected pattern in the collection. (Using its connected components instead requires less space and gives identical results.)

Search neighbourhood



- PDB heuristics
- Implementing PDBs
- Additivity
- Pattern selection
- Local search
- Search space
- Estimating heuristic quality
- Summary

We can now put the pieces together to define our search neighbourhood, obtaining the third answer:

3. Which collections are neighbours of the current collection?

The neighbours of \mathcal{C} are all pattern collections $\mathcal{C} \cup \{P'\}$ where

- $P' = P \cup \{v\}$ for some $P \in \mathcal{C}$,
- $P' \notin \mathcal{C}$,
- all variables of P' are causally relevant in P' ,
- P' is causally connected, and
- all pattern databases in $\mathcal{C} \cup \{P'\}$ can be represented within some prespecified space limit

Search neighborhood (ctd.)



- PDB heuristics
- Implementing PDBs
- Additivity
- Pattern selection
- Local search
- Search space
- Estimating heuristic quality
- Summary

Remark: For causal relevance and connectivity, there is a sufficient and necessary criterion which is easy to check:

- v is a predecessor of some $u \in P$ in the causal graph, **or**
- v is a successor of some $u \in P$ in the causal graph and is mentioned in the goal formula.

What is a good pattern collection?



- PDB heuristics
- Implementing PDBs
- Additivity
- Pattern selection
- Local search
- Search space
- Estimating heuristic quality
- Summary

- The last question we need to answer is **how to rank** the quality of pattern collections.
- This is perhaps the most critical point: without a good ranking criterion, pattern collections are chosen blindly.

The first search-based approach to pattern selection (Edelkamp, 2007) used the following strategy:

- only additive sets are used as pattern collections
 - no need for something like the canonical heuristic function
- the quality of a **single pattern** is estimated by its **mean heuristic value** (the higher, the better)
- the quality of a **pattern collection** is estimated by the **sum** of the individual pattern qualities

Discussion of the mean value approach



- PDB heuristics
- Implementing PDBs
- Additivity
- Pattern selection
- Local search
- Search space
- Estimating heuristic quality
- Summary

Pros of the approach:

- mean heuristic values are clearly correlated with search performance \rightsquigarrow the **quality measure makes sense**
- mean heuristic values are quite **easy to calculate**

Cons of the approach:

- cannot reasonably deal with **infinite** heuristic estimates
- difficult to generalize to pattern collections that are **not fully additive**
- there are **better predictors** for search performance than mean heuristic values


So what is a good pattern collection, again?



- PDB heuristics
- Implementing PDBs
- Additivity
- Pattern selection
- Local search
- Search space
- Estimating heuristic quality
- Summary

How can we come up with a better quality measure?

- We are chiefly interested in **minimizing the number of node expansions** for the **canonical heuristic function** during the **actual search phase** of the planner.
 - There is theoretical work on **predicting node expansions** of heuristic search algorithms based on parameters of the heuristic (Korf, Reid & Edelkamp, 2001).
- \rightsquigarrow Try to estimate these parameters, then use their analysis.



- PDB heuristics
- Implementing PDBs
- Additivity
- Pattern selection
- Local search
- Search space
- Estimating heuristic quality
- Summary

With some assumptions and simplifications, we reduce the problem of ranking the pattern collection quality to this:

Measuring degree of improvement

- Generate M states s_1, \dots, s_M through random walks in the search space from the initial state (according to certain parameters not discussed in detail).
- The **degree of improvement** of a pattern collection \mathcal{C}' which is generated as a successor of collection \mathcal{C} is the **number of sample states s_i for which $h^{\mathcal{C}'}(s_i) > h^{\mathcal{C}}(s_i)$.**

- PDB heuristics
- Implementing PDBs
- Additivity
- Pattern selection
- Local search
- Search space
- Estimating heuristic quality
- Summary

- So we need to compute $h^{\mathcal{C}'}(s)$ for some states s and each candidate successor collection \mathcal{C}' .
- We have PDBs for all patterns in \mathcal{C} , but not for the new pattern $P' \in \mathcal{C}'$ (of the form $P \cup \{v\}$ for some $P \in \mathcal{C}$).
- If possible, we want to avoid computing the complete pattern database except for the best successor (where we will need it later anyway).

Idea:

- For SAS⁺ tasks Π , $h^{P'}(s)$ is identical to the **optimal solution length for the syntactic projection $\Pi|_{P'}$.**
- We can use **any optimal planning algorithm** for this.
- In particular, we can use **A*** search using h^P as a heuristic.

- PDB heuristics
- Implementing PDBs
- Additivity
- Pattern selection
- Summary

- **Pattern database (PDB) heuristics** are abstraction heuristics based on **projection** to a subset of variables.
- For SAS⁺ tasks, they can easily be implemented via **syntactic projections** on the task representation.
- PDBs are **lookup tables** that store heuristic values, indexed by **perfect hash values** for projected states.
- PDB values can be looked up **very fast**, in time $O(k)$ for a projection to k variables.

- PDB heuristics
- Implementing PDBs
- Additivity
- Pattern selection
- Summary

- When faced with multiple PDB heuristics (a **pattern collection**), we want to **admissibly add** their values where possible, and **maximize** where addition is inadmissible.
- The **canonical heuristic function** is the **best possible** additive/maximizing combination for a given pattern collection given our additivity criterion.
- One way to **automatically find a good pattern collection** is by performing **search in the space of pattern collections**.