

# AI Planning

## 16. Combining Heuristic Functions Part II - LP Heuristics

### “Beyond Optimal Cost Partitioning”

Álvaro Torralba, Cosmina Croitoru



Winter Term 2018/2019

Thanks to Prof. Malte Helmert and Dr. Gabrielle Röger for slide sources

# Agenda

- 1 Introduction
- 2 Post-hoc Optimization Heuristic
- 3 Operator-counting Framework
- 4 Potential Heuristics
- 5 Conclusion

# Our Agenda for This Chapter

- 2 **Post-hoc optimization:** Combining heuristics with a subset of relevant actions.
- 3 **Operator Counting Heuristics:** Interpreting heuristics as constraints on the set of actions that must be applied.
- 4 **Potential Heuristics:** Deriving heuristics via linear programming.

# Combining Estimates from Abstraction Heuristics

We already know two approaches to derive heuristic estimates from a collection of abstractions:

(Reminder **Chapter 12**): Canonical Heuristic Function  

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

(Reminder **Chapter 16**): Optimal Cost Partitioning: uses a **state-specific LP** to find the **best possible cost partitioning**, and sums up the heuristic estimates.

→ is **very expensive** to compute (recomputing the PDBs in every state).

Is there a way to combine multiple abstraction heuristics that **does not need to re-compute the distances** in every state and is **more informative than the canonical heuristic**?

## Example Task (1)

### Example (Example Task)

- FDR task  $\Pi = \langle V, I, O, \gamma \rangle$  with
- $V = \{A, B, C\}$  with  $\text{dom}(v) = \{0, 1, 2, 3, 4\}$  for all  $v \in V$
  - $I = \{A \mapsto 0, B \mapsto 0, C \mapsto 0\}$
  - $O = \{inc_x^v \mid v \in V, x \in \{0, 1, 2\}\} \cup \{jump^v \mid v \in V\}$ 
    - $inc_x^v = \langle v = x, v := x + 1, 1 \rangle$
    - $jump^v = \langle \bigwedge_{v' \in V: v' \neq v} v' = 4, v := 3, 1 \rangle$
  - $\gamma = A = 3 \wedge B = 3 \wedge C = 3$
- Each optimal plan consists of three increment actions for each variable  $\rightsquigarrow h^*(I) = 9$
  - Each action affects only one variable.

## Post-hoc Optimization Heuristic: Idea

- Consider the example task:
- **type- $v$  action**: action modifying variable  $v$
  - $h^{\{A,B\}} = 6$   
 $\Rightarrow$  any plan contains **at least 6 actions of type  $A$  or  $B$** .
  - $h^{\{A,C\}} = 6$   
 $\Rightarrow$  any plan contains **at least 6 actions of type  $A$  or  $C$** .
  - $h^{\{B,C\}} = 6$   
 $\Rightarrow$  any plan contains **at least 6 actions of type  $B$  or  $C$** .
  - $\Rightarrow$  **at least 9 actions** in any plan

Can we generalize this kind of reasoning?

## Example Task (2)

- In projections on single variables we can reach the goal with a *jump* action:  $h^{\{A\}}(I) = h^{\{B\}}(I) = h^{\{C\}}(I) = 1$ .
- In projections on more variables, we need for each variable three applications of increment actions to reach the abstract goal from the abstract initial state:  $h^{\{A,B\}}(I) = h^{\{A,C\}}(I) = h^{\{B,C\}}(I) = 6$

### Example (Canonical Heuristic)

$$\mathcal{C} = \{\{A\}, \{B\}, \{C\}, \{A, B\}, \{A, C\}, \{B, C\}\}$$

$$h^{\mathcal{C}}(s) = \max\{h^{\{A\}}(s) + h^{\{B\}}(s) + h^{\{C\}}(s), h^{\{A\}}(s) + h^{\{B,C\}}(s), h^{\{B\}}(s) + h^{\{A,C\}}(s), h^{\{C\}}(s) + h^{\{A,B\}}(s)\}$$

$$h^{\mathcal{C}}(I) = 7$$

## Post-hoc Optimization Heuristic: Linear Program

For pattern collection  $\mathcal{C}$ :

### Variables

$X_a$  for each action  $a \in A$

### Objective

Minimize  $\sum_{a \in A} X_a$

### Subject to

$$\sum_{a \in A: a \text{ affects } P} X_a \geq h^P(s) \quad \text{for all patterns } P \in \mathcal{C}$$

$$X_a \geq 0 \quad \text{for all } a \in A$$

$\rightarrow$ For each new state, just change the bounds  $h^P(s)$ .

## Post-hoc Optimization Heuristic: Admissibility

**Theorem (Admissibility):** The post-hoc optimization heuristic is **admissible**.

Proof: Let  $\Pi$  be a planning task and  $\mathcal{C}$  be a pattern collection. Let  $\pi$  be an optimal plan for state  $s$  and let  $c_\pi(A')$  be the cost incurred by actions from  $A' \subseteq A$  in  $\pi$ .

Setting each  $X_{[a]}$  to  $c_\pi([a])$  is a feasible variable assignment: Constraints  $X_{[a]} \geq 0$  are satisfied. For each  $P \in \mathcal{C}$ ,  $\pi$  is a solution in the abstract transition system and the sum in the corresponding constraint equals the cost of the “true” abstract state transitions (i.e., not accounting for self-loops). As  $h^P(s)$  corresponds to the cost of an optimal solution in the abstraction, the inequality holds.

For this assignment the objective function has value  $h^*(s)$  (cost of  $\pi$ ), so the objective value of the LP is admissible.

## Relation to Canonical Heuristic

**Theorem:** The post-hoc optimization heuristic  $h_c^{PhO}$  dominates the canonical heuristic  $h^c$  for the same pattern collection  $\mathcal{C}$ .

**Proof:** Consider the dual  $D$  of the LP solved by  $h_c^{PhO}$  in state  $s$  for a given pattern collection  $\mathcal{C}$ : Maximize  $\sum_{P \in \mathcal{C}} Y_P h^P(s)$  subject to

$$\sum_{P \in \mathcal{C}: o \text{ affects } P} Y_P \leq 1 \quad \text{for all } [o] \in O/\sim$$

$$Y_P \geq 0 \quad \text{for all } P \in \mathcal{C}.$$

We compute a state-specific cost partitioning that can only scale the action costs within each heuristic by a factor  $Y_i$ .

If we restrict the variables in  $D$  to integers, the objective value is the canonical heuristic value  $h^c(s)$ .

→For the canonical heuristic, we need to find all maximal cliques, which is an NP-hard problem.

→The post-hoc optimization heuristic dominates the canonical heuristic and can be computed in polynomial time.

## Reminder: Optimal Cost Partitioning for Landmarks

### Variables

$\text{Count}_o$  for each action  $o$

### Objective

Minimize  $\sum_o \text{Count}_o \cdot c(o)$

### Subject to

$$\sum_{o \in L} \text{Count}_o \geq 1 \text{ for all landmarks } L$$

$$\text{Count}_o \geq 0 \text{ for all actions } o$$

Numbers of action occurrences in any plan satisfy constraints. Minimizing the total plan cost gives an admissible estimate.

Can we apply this idea more generally?

## Operator Counting

### Operator-counting Constraints

- linear constraints whose variables denote number of occurrences of a given action
- must be satisfied by every plan

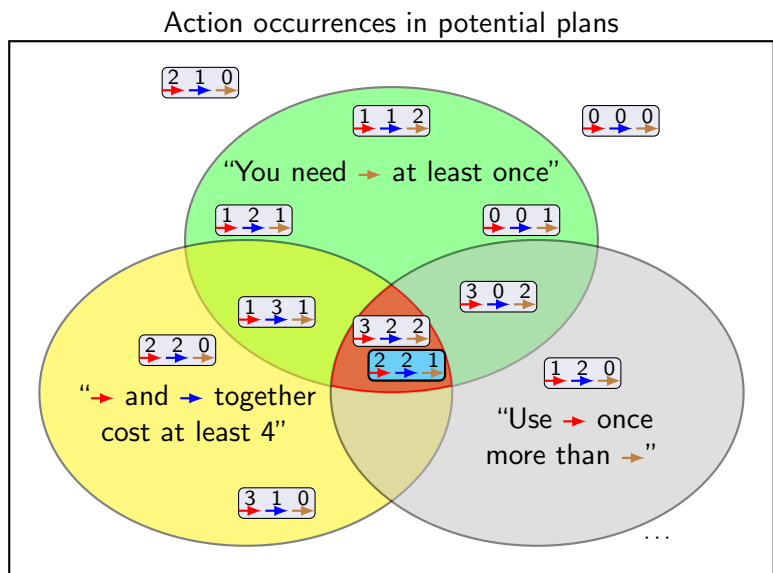
### Examples:

- $\text{Count}_{o_1} + \text{Count}_{o_2} \geq 1$  “must use  $o_1$  or  $o_2$  at least once”
- $\text{Count}_{o_1} - \text{Count}_{o_3} \leq 0$  “cannot use  $o_1$  more often than  $o_3$ ”

### Motivation:

- declarative way to represent knowledge about solutions
- allows reasoning about solutions to derive heuristic estimates

# Operator Counting Heuristics



Álvaro Torralba, Cosmina Croitoru AI Planning Chapter 16: Combining Heuristic Functions Part II - LP Heuristics 16/35

# Operator-counting Constraint

**Definition (Operator-counting Constraints)** Let  $\Pi$  be an FDR task with actions  $A$  and let  $s$  be a state. Let  $\mathcal{V}$  be the set of integer variables  $Count_a$  for each  $a \in A$ . A linear inequality over  $\mathcal{V}$  is called an *operator-counting constraint* for  $s$  if for every plan  $\pi$  for  $s$  setting each  $Count_a$  to the number of occurrences of  $a$  in  $\pi$  is a feasible variable assignment.

→ An operator-counting constraint must be true in every plan!.

What constraints we know already?

- Landmark constraint for landmark  $L$ :  $\sum_{o \in L} Count_o \geq 1$
- Post-hoc Optimization constraint for heuristic  $h$

$$\sum_{o \text{ is relevant for } h} Count_o \cdot c(o) \geq h(s)$$

Álvaro Torralba, Cosmina Croitoru AI Planning Chapter 16: Combining Heuristic Functions Part II - LP Heuristics 17/35

# Flow Heuristic Constraints

**Flow heuristic** Use one flow constraint per fact  $(v, d)$

$$\sum_{(v, d) \text{ produced by } a} Count_a - \sum_{(v, d) \text{ consumed by } a} Count_a \geq [G[v] = d] - [s[v] = d]$$

where:

- $(v, d)$  consumed by  $a$  if  $(v, d) \in pre_a$  and  $(v, d') \in eff_a$  for some  $d' \neq d$
- $(v, d)$  produced by  $a$  if  $(v, d) \in eff_a$  and  $((v, d) \notin pre_a)$
- $[G[v] = d] = 1$  if  $G[v] = d$  and 0 otherwise
- $[s[v] = d] = 1$  if  $s[v] = d$  and 0 otherwise

Álvaro Torralba, Cosmina Croitoru AI Planning Chapter 16: Combining Heuristic Functions Part II - LP Heuristics 18/35

# Operator-counting IP/LP Heuristic

**Definition (Operator-counting IP/LP Heuristic)** The operator-counting integer program  $IP_C$  for a set  $C$  of operator-counting constraints for state  $s$  is

$$\text{Minimize } \sum_{a \in A} Count_a \cdot c(a) \text{ subject to}$$

$$C \text{ and } Count_a \geq 0 \text{ for all } a \in A,$$

The *IP heuristic*  $h_C^{IP}$  is the objective value of  $IP_C$ , the *LP heuristic*  $h_C^{LP}$  is the objective value of its LP-relaxation. If the IP/LP is infeasible, the heuristic estimate is  $\infty$ .

Álvaro Torralba, Cosmina Croitoru AI Planning Chapter 16: Combining Heuristic Functions Part II - LP Heuristics 19/35

## Admissibility

### Theorem (Operator-counting Heuristics are Admissible)

The IP and the LP heuristic are **admissible**.


#### Proof.

Let  $C$  be a set of operator-counting constraints for state  $s$  and  $\pi$  be an optimal plan for  $s$ . The number of action occurrences of  $\pi$  are a feasible solution for  $C$ . As the IP/LP minimizes the total plan cost, the objective value cannot exceed the cost of  $\pi$  and is therefore an admissible estimate. □

## Combining Heuristics

### Combination of two heuristics

- Use both operator-counting constraints
- Combination always **dominates individual heuristics**
- **Positive interaction** between constraints



### Combination often better than best individual heuristic

## Dominance

**Theorem** Let  $C$  and  $C'$  be operator-counting constraints for  $s$  and let  $C \subseteq C'$ . Then  $IP_C \leq IP_{C'}$  and  $LP_C \leq LP_{C'}$ .

Proof: Every feasible solution of  $C'$  is also feasible for  $C$ . As the LP/IP is a minimization problem, the objective value subject to  $C$  can therefore not be larger than the one subject to  $C'$ .

Adding more constraints can only improve the heuristic estimate.

**Theorem:** Combining **operator-counting heuristics** in one LP is equivalent to computing their **optimal general cost partitioning**.

**Proof idea:** The linear programs are each others duals.

Even **better combination** of heuristics with **IP heuristic**

- Considers that actions cannot be used 1.5 times
- But computation is **no longer polynomial**

## Motivation

- Operator-counting heuristics solve an LP to compute the heuristic estimate **for a single state**.
- Can we also define an **entire heuristic function** solving only one LP?
- **Axiomatic approach** for defining heuristics:
  - What should a heuristic look like mathematically?
  - Which properties should it have?
- Define a **space of interesting heuristics**.
- Use **optimization** to pick a good representative.

# Potential Heuristics

## Potential Heuristics: Idea

Heuristic design as an optimization problem:

- Define simple numerical **state features**  $f_1, \dots, f_n$ .
- Consider heuristics that are **linear combinations** of features:

$$h(s) = w_1 f_1(s) + \dots + w_n f_n(s)$$

with weights (**potentials**)  $w_i \in \mathbb{R}$

- Find potentials for which  $h$  is admissible and well-informed.

### Motivation:

- **declarative approach** to heuristic design
- heuristic **very fast to compute** if features are

# Potential Heuristics

## Definition (feature)

A (state) **feature** for a planning task is a numerical function defined on the states of the task:  $f : S \rightarrow \mathbb{R}$ .

## Definition (potential heuristic)

A **potential heuristic** for a set of features  $\mathcal{F} = \{f_1, \dots, f_n\}$  is a heuristic function  $h$  defined as a **linear combination** of the features:

$$h(s) = w_1 f_1(s) + \dots + w_n f_n(s)$$

with weights (**potentials**)  $w_i \in \mathbb{R}$ .

↔ cf. **evaluation functions** for board games like chess

# Atomic Potential Heuristics

**Atomic features** test if some atom is true in a state:

## Definition (atomic feature)

Let  $X = x$  be an atom of a FDR planning task.

The **atomic feature**  $f_{X=x}$  is defined as:

$$f_{X=x}(s) = \begin{cases} 1 & \text{if variable } X \text{ has value } x \text{ in state } s \\ 0 & \text{otherwise} \end{cases}$$

- We only consider **atomic** potential heuristics, which are based on the set of all atomic features.
- **Example** for a task with state variables  $X$  and  $Y$ :

$$h(s) = 3f_{X=a} + \frac{1}{2}f_{X=b} - 2f_{X=c} + \frac{5}{2}f_{Y=d}$$

# How to Set the Weights?

We want to find **good** atomic potential heuristics:

- admissible
- consistent
- well-informed

How to achieve this? **Linear programming to the rescue!**

## Admissible and Consistent Potential Heuristics

Constraints on potentials **characterize** (= are necessary and sufficient for) admissible and consistent atomic potential heuristics:

Goal-awareness

$$\sum_{\text{goal atoms } a} w_a = 0$$

Consistency

$$\sum_{\substack{a \\ \text{consumed} \\ \text{by } o}} w_a - \sum_{\substack{a \\ \text{produced} \\ \text{by } o}} w_a \leq c(o) \quad \text{for all operators } o$$

Remarks:

- assumes transition normal form (not a limitation)
- goal-aware and consistent = admissible and consistent

## Summary

- **Post-hoc optimization heuristic** explores the middle ground between canonical heuristic and optimal cost partitioning.
- For the same pattern collection the post-hoc optimization heuristic **dominates the canonical heuristic**.
- The computation can be done in **polynomial time**.
- Many heuristics can be formulated in terms of **operator-counting constraints**.
- The operator-counting heuristic framework allows to **combine the constraints** and to reason on the entire encoded declarative knowledge.
- The heuristic estimate for the combined constraints **can be better than the one of the best ingredient heuristic** but never worse.

## Well-Informed Potential Heuristics

How to find a **well-informed** potential heuristic?

↪ encode **quality metric** in the **objective function** and use LP solver to find a heuristic maximizing it

Examples:

- maximize **heuristic value of a given state** (e.g., initial state)
- maximize average heuristic value of **all states** (including unreachable ones)
- maximize average heuristic value of some **sample states**
- minimize **estimated search effort**

## Summary

- The combination into one operator-counting heuristic corresponds to the computation of the **optimal general cost partitioning for the ingredient heuristics**.
- General cost partitioning, operator-counting constraints and potential heuristics are **facets of the same phenomenon**.
- Study of each reinforces understanding of the others.

# Transition Normal Form

**Definition (Transition Normal Form)** An FDR planning task  $\Pi = (V, A, c, I, G)$  is in transition normal form (TNF) if for all  $a \in A$ ,  $V[pre(a)] = V[eff(a)]$ , and  $V[G] = V$ .

→An operator in TNF must mention the same variables in the precondition and effect, and a goal in TNF must mention all variables (= specify exactly one goal state).

Any FDR task can be transformed into TNF:

- ① For every variable  $v$ , add a new auxiliary value  $u$  to its domain.
- ② For every variable  $v$  and value  $d \in D_v \setminus \{u\}$ , add a new operator to change the value of  $v$  from  $d$  to  $u$  with no precondition at no cost.
- ③ For all actions, if  $v = d$  is a precondition with no effect on  $v$ , just add the effect  $v := d$ .
- ④ For all actions  $a$  and all variables  $v \in V[eff(a)] \setminus V[pre(a)]$ , add the precondition  $v = u$  to  $pre(a)$ .
- ⑤ For every variable  $v \in V \setminus V[G]$ , add the condition  $v = u$  to  $G$ .

# Reading