LP-based Heuristics for Cost-optimal Classical Planning

4. Potential Heuristics

Florian Pommerening   Gabriele Röger   Malte Helmert

ICAPS 2015 Tutorial

June 7, 2015

Overview

Potential Heuristics

Reminder:

Idea 3: Potential Heuristics

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

Comparison to Previous Parts (1)

What is the same as in operator-counting constraints:

- We again use LPs to compute (admissible) heuristic values (spoiler alert!)

Comparison to Previous Parts (2)

What is different from operator-counting constraints (computationally):

- With potential heuristics, solving one LP defines the entire heuristic function, not just the estimate for a single state.
- Hence we only need one LP solver call, making LP solving much less time-critical.

Comparison to Previous Parts (3)

What is different from operator-counting constraints (conceptually):

- 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

Literature on Potential Heuristics: The Story So Far

Papers studying potential heuristics:

- introduced by Pommerening et al. (AAAI 2015)
  - ⇝ main focus of this presentation
- studied in more depth by Seipp et al. (ICAPS 2015)
  - ⇝ presentation: Thursday, joint ICAPS/SoCS session (last session of conference)
- sufficient to consider transition normal form (Pommerening and Helmert, ICAPS 2015)
  - ⇝ presentation: Tuesday, first afternoon session

Potential Heuristics

Features

Definition (feature)

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

Potential Heuristics

Definition (potential heuristic)

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

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

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

$\rightsquigarrow$ cf. evaluation functions for board games like chess

Atomic Potential Heuristics

Atomic features test if some proposition is true in a state:

Definition (atomic feature)

Let $X = x$ be an atomic proposition of a 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:

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

Finding Good Potential Heuristics

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 (i.e., $h(s) = 0$ for goal states)

$$\sum_{\text{goal facts } f} w_f = 0$$

Consistency

$$\sum_{\substack{f \text{ consumed} \\ \text{by } o}} w_f - \sum_{\substack{f \text{ produced} \\ \text{by } o}} w_f \leq cost(o) \quad \text{for all operators } o$$

Remarks:

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

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

⇝ see Seipp et al. presentation (joint ICAPS/SoCS session)

Connections

Connections

So what does this have to do with what we talked about before?

Connections

So what does this have to do with what we talked about before?

Theorem (Pommerening et al., AAAI 2015)

For state s, let $h^{\mathrm{maxpot}}(s)$ denote the maximal heuristic value of all admissible and consistent atomic potential heuristics in s.

Then $h^{\mathrm{maxpot}}(s) = h^{\mathrm{SEQ}}(s) = h^{\mathrm{gOCP}}(s)$.

- h^{SEQ}: state equation heuristic a.k.a. flow heuristic
- h^{gOCP}: optimal general cost partitioning of atomic projections

proof idea: compare dual of $h^{\mathrm{SEQ}}(s)$ LP
to potential heuristic constraints optimized for state s

What Do We Take From This?

- general cost partitioning, operator-counting constraints and potential heuristics: facets of the same phenomenon
- study of each reinforces understanding of the others
- potential heuristics: fast admissible approximations of h^{SEQ}
- clear path towards generalization beyond h^{SEQ}: use non-atomic features

The End

1. ~~Introduction and Overview~~
2. ~~Cost Partitioning~~
3. ~~Operator Counting~~
4. ~~Potential Heuristics~~

Thank you for your attention!