

Accuracy of Admissible Heuristic Functions in Selected Planning Domains

Malte Helmert and Robert Mattmüller

Albert-Ludwigs-Universität Freiburg, Germany
{helmert,mattmuel}@informatik.uni-freiburg.de

Abstract

The efficiency of optimal planning algorithms based on heuristic search crucially depends on the accuracy of the heuristic function used to guide the search. Often, we are interested in domain-independent heuristics for planning. In order to assess the limitations of domain-independent heuristic planning, we analyze the (in)accuracy of common domain-independent planning heuristics in the IPC benchmark domains. For a selection of these domains, we analytically investigate the accuracy of the h^+ heuristic, the h^m family of heuristics, and certain (additive) pattern database heuristics, compared to the perfect heuristic h^* . Whereas h^+ and additive pattern database heuristics usually return cost estimates proportional to the true cost, non-additive h^m and non-additive pattern-database heuristics can yield results underestimating the true cost by arbitrarily large factors.

Introduction

Heuristic search with A* and similar algorithms remains the most popular method for optimal sequential planning, with significant effort spent on perfecting old heuristic estimators (Haslum et al. 2007) or deriving new ones (Helmert, Haslum, and Hoffmann 2007). While methods not based on state-space search have achieved remarkable success in addressing related problems, such as optimal parallel planning (Kautz and Selman 1996; 1999), the state of the art in optimal sequential planning is still defined by heuristic search almost exclusively. Symbolic state-space exploration (Edelkamp and Helmert 2001) is the only non-classical approach that sometimes outperforms heuristic search.

Considering the important role of admissible heuristics for optimal sequential planning, or search in general, the question arises how to evaluate the quality of a given heuristic. A popular method is to run a search algorithm against some benchmark tasks and count the number of node expansions. The fewer nodes an algorithm expands, the better.

While experiments of this kind are certainly useful, there are some questions they cannot address. In particular, their results can almost exclusively be interpreted with *relative*, i. e., *comparative* statements: “Heuristic h expands fewer nodes than heuristic h' for benchmark suite X .” Unless experiments show polynomial scaling behavior on a family of

benchmark tasks of growing size, which they very rarely do, the data usually does not lend itself to *absolute* statements of the type “Heuristic h is well-suited for solving tasks from benchmark suite X .” In this contribution, we address this issue by providing absolute quality results for certain popular planning heuristics on some popular benchmark domains taken from the first four International Planning Competitions (McDermott 2000; Bacchus 2001; Long and Fox 2003; Hoffmann and Edelkamp 2005), in the form of comparisons to the perfect heuristic function h^* .

Planning Domains

We consider the planning domains GRIPPER, LOGISTICS, BLOCKSWORLD, MICONIC-STRIPS, MICONIC-SIMPLE-ADL, SCHEDULE and SATELLITE. Familiarity with the domains is assumed. For an in-depth treatment, we refer to the literature (Helmert 2008).

Heuristics

We compare the accuracy of h^+ , h^m , non-additive and additive pattern database (PDB) heuristics relative to the perfect heuristic h^* . An admissible heuristic h maps states s to optimistic estimates of the true cost of reaching a goal state from s . Whenever the planning task to which s belongs, including the available operators and the goal description, is clear from the context, we will simply write $h(s)$ without explicitly mentioning operators or goal.

The h^* Heuristic The *perfect heuristic* h^* assigns to each state s the length of a shortest plan from s to a goal state. Computing h^* for the initial state of a planning task is PSPACE-equivalent in general (Bylander 1994), but can be easier for fixed domains. For the domains we consider, the problem is NP-equivalent, with the exception of GRIPPER and SCHEDULE, where it is polynomial (Helmert 2008).

The h^+ Heuristic The h^+ heuristic (McDermott 1996; Bonet and Geffner 2001; Hoffmann 2005) assigns to each state s of a task \mathcal{T} the length of a shortest plan leading from s to a goal state in the relaxed task \mathcal{T}^+ . Evaluating h^+ is NP-equivalent in general (Bylander 1994), but easier for many of the domains we consider.

The h^m Family of Heuristics The h^m , $m = 1, 2, \dots$, family of heuristics (Haslum and Geffner 2000) is based on

the relaxation where the cost of reaching a state with n satisfied atoms is approximated by the highest cost of reaching a subset with at most m satisfied atoms. It can be computed in polynomial time, with the degree of the polynomial depending on m . Recently, the h^m family has been extended by introducing *additive h^m heuristics* (Haslum, Bonet, and Geffner 2005). The results we present for the h^m family do not apply to their additive counterpart.

PDB Heuristics PDB heuristics (Culberson and Schaeffer 1998) are computed by projecting the state space onto a subset of the state variables and computing the shortest plan in the resulting abstract state space. As these abstractions are solution preserving, abstract solutions are never longer than corresponding concrete ones, making the heuristic admissible. In the case of PDB heuristics we will argue about multi-valued state variable representations instead of binary ones, because this is what all existing PDB planners use.

Possible PDB heuristics for a task are parametrized by the set of variables retained in the abstraction (the *pattern*). The maximum number M of such state variables should be small since the number of states in the abstract task as well as the time to compute and the memory to store a PDB for M variables is exponential in M . PDB heuristics can be computed in polynomial time in the task size n only if $M \in O(\log n)$. We will in the following adhere to this restriction.

Additive PDB Heuristics Most search algorithms that use PDB heuristics consider more than one pattern for deriving their heuristic estimates. Two patterns are considered *additive* (Haslum et al. 2007) if no operator modifies variables from both patterns. If patterns are additive, then *summing up* the entries from the PDB is still guaranteed to yield a value not greater than the true cost, i. e., the heuristic obtained by summation is admissible.

An additive PDB heuristic assigns to each state such a sum of additive pattern values. As in the case of regular PDBs, we must restrict the size of *individual* patterns by $O(\log n)$ for task size n for polynomial-time computability.

Note that it is common to not just consider *one* sum of individual pattern heuristics, but maximize over several such sums. For example, Haslum et al. (2007) define the *canonical heuristic function* for a set of patterns as a maximum over all additive sums of pattern heuristics. In this work, we limit ourselves to additive PDB heuristics without maximization over several sums, leaving this extension for future work.

Accuracy

By definition, the h^* heuristic is exact, i. e., it returns the length of a shortest plan. For the other heuristics, we give domain-dependent worst-case bounds on the accuracy relative to h^* and show that these bounds are tight in the sense that there are families of increasing-size tasks for which the accuracy tends towards the respective bound.

Formally, given a planning domain \mathcal{D} and heuristic h , we want to find the asymptotic accuracy of h in \mathcal{D} . More precisely, we want to find a value $\alpha \in \mathbb{R}$ such that (a) for all legal initial states s in all tasks \mathcal{T} of domain \mathcal{D} , $h(s) \geq \alpha h^*(s) + o(h^*(s))$, and (b) there is a family of tasks $(\mathcal{T}_n)_{n \in \mathbb{N}}$ in \mathcal{D} and solvable initial states $(s_n)_{n \in \mathbb{N}}$ such

that s_n belongs to \mathcal{T}_n , $h^*(s_n) \rightarrow \infty$ for $n \rightarrow \infty$ and $h(s_n) \leq \alpha h^*(s_n) + o(h^*(s_n))$.

In other words, h is *never* worse than αh^* (plus a sublinear term), and it can become as bad as αh^* (plus a sublinear term) for arbitrarily large inputs. These conditions can only be satisfied for a single value α ; moreover, there must be a value $\alpha \in [0, 1]$ which satisfies them for a given heuristic and domain. We denote this constant by $\alpha(h, \mathcal{D})$. Moreover, by $\alpha(h, s)$ we denote the ratio $h(s)/h^*(s)$ (we only consider non-goal states; otherwise $\alpha(h, s)$ is undefined). We simply write $\alpha(s)$ when the heuristic is clear from the context.

Results

Accuracy of h^+

First, we investigate the accuracy of the h^+ heuristic. As a general observation, $h^+(s)$, if not infinite, is always bounded by a linear term in the number of propositional state variables of the task. Therefore, for all domains \mathcal{D} where there exists a family of tasks $(\mathcal{T}_n)_{n \in \mathbb{N}}$ of growing size where optimal solution lengths grow super-linearly with task size (formally, $h^*(s_n) = \omega(\|\mathcal{T}_n\|)$ for states s_n of \mathcal{T}_n), we must have $\alpha(h^+, \mathcal{D}) = 0$. However, none of the benchmark domains we study has this property – they can all be solved by linear-sized plans.

GRIPPER Let s_n be a state with $n > 0$ balls still at their start location. We assume the robot is at the balls' start location, both grippers are empty and n is even. Otherwise, the h^* and h^+ values can increase by a small constant. The optimal plan consists of n *pick-up* and *drop* actions each, $n/2$ forth moves (holding two balls) and $n/2 - 1$ back moves (with empty grippers). Thus, $h^*(s_n) = 3n - 1$. In the relaxed task, the robot can be at both locations simultaneously after having performed its first *move* action. Together with the $2n$ *pick-up* and *drop* actions still necessary, this gives an optimal relaxed plan length of $h^+(s_n) = 2n + 1$. Thus, $\alpha(s_n) = (2n + 1)/(3n - 1) > 2/3$. On the other hand, for the family $(s_n)_{n \in \mathbb{N}}$ of tasks described above, $\lim_{n \rightarrow \infty} \alpha(s_n) = 2/3$. Thus, $\alpha(h^+, \text{GRIPPER}) = 2/3$.

LOGISTICS Consider a state s_n with n cities, no airports, and for each city c_i , two locations ℓ_{i1} and ℓ_{i2} , one truck t_i located at ℓ_{i1} , and one package with origin ℓ_{i2} and destination ℓ_{i1} . An optimal plan consists of $4n$ actions, namely two *drive*, one *pick-up* and one *drop* action in each city. An optimal relaxed plan only needs three actions per city since the second *drive* action can be omitted because the atom $at(t_i, \ell_{i1})$ holding initially is never deleted. Thus, we get $\alpha(h^+, \text{LOGISTICS}) \leq 3/4$.

For the other direction, we show that $h^+ \geq 3/4 h^*$ for all LOGISTICS tasks. Of the results presented in this paper, this one has the most complicated proof, which we can only sketch here for space reasons. The *delivery graph* of a LOGISTICS task is the directed multigraph which contains an arc from location u to location v if some package needs to be moved from u to v . We can show that it is sufficient to consider the case where the delivery graph is weakly connected and satisfies certain additional properties, such as all vertices having an indegree and outdegree of at least 1, and

there is only one airplane. For that case, we consider a particular three-phase plan: first, collect packages in each city and deliver them to the airport, if necessary; second, solve the (MICONIC-STRIPS-like) task for the airplane optimally; third, distribute packages in the cities. We can show that the length L of this plan is at most $4/3h^+(s)$. Since obviously $L \geq h^*(s)$, this implies that $h^+ \geq 3/4h^*$. Together with the upper bound, we get $\alpha(h^+, \text{LOGISTICS}) = 3/4$.

BLOCKSWORLD Let us first consider the family $(s_n)_{n \in \mathbb{N}}$ of states where s_n consists of one stack B_1, \dots, B_{n+1} (from top to bottom) and another singleton stack B_{n+2} and the goal is to produce the stack B_1, \dots, B_n, B_{n+2} , with the position of B_{n+1} not being important. One optimal plan for s_n consists of $n - 1$ pairs of *unstack* and *put-down* actions to put the blocks B_1, \dots, B_{n-1} on the table, one *unstack* and *stack* action to move B_n from B_{n+1} to B_{n+2} , and $n - 1$ pairs of *pick-up* and *stack* actions to rebuild the stack on top of block B_n . Thus, $h^*(s_n) = 4n - 2$. In order to solve the relaxed task, we have to *unstack* blocks B_1, \dots, B_n in order to add the *holding*(B_n) proposition. We can *unstack* all blocks without emptying the hand as the proposition *handempty* is true initially and is never falsified. To satisfy the goal, we then only need to stack B_n on B_{n+2} . The rest of the stack does not need to be rebuilt, as all propositions *on*(B_i, B_{i+1}) for $i < n$ still hold. This means that $\lim_{n \rightarrow \infty} \alpha(s_n) = \lim_{n \rightarrow \infty} (n + 1)/(4n - 2) = 1/4$.

For the other direction of the proof, we will show that $h^+(s) \geq 1/4h^*(s) + d$ for a constant $d \in \mathbb{N}$ in all non-goal BLOCKSWORLD states. Consider a block B . If B is not touched in an optimal plan, it is not touched in an optimal relaxed plan either. If it is touched in an optimal plan, there are at most four actions moving B around, as it is always sufficient to pick-up/unstack and put-down/stack B once or twice. At least one of those up to four actions is also needed in an optimal relaxed plan. Therefore, $h^+(s) \geq 1/4h^*(s)$, and thus $\alpha(h^+, \text{BLOCKSWORLD}) = 1/4$.

MICONIC-STRIPS Consider a state s_n with $2n + 1$ locations, where location 0 is the elevator start location. There are $2n$ passengers, one waiting at each location $i = 1, \dots, 2n$. The destination of the passenger at location i is $i + 1$ if i is odd and $i - 1$ if i is even. The optimal solution contains precisely seven actions per passenger pair, i. e., $h^*(s_n) = 7n$. In the relaxed task, a shortest plan is essentially the same as in the original task, the only exception being the fact that we can save one move action per location pair (the second move action back to the location visited twice). Therefore, $h^+(s_n) = 6n$ and $\lim_{n \rightarrow \infty} \alpha(s_n) = 6/7$.

To see that $h^+(s) \geq 6/7h^*(s) + d$ for all MICONIC-STRIPS states s and a suitable constant d , note that the only actions that can potentially be saved in a relaxed plan are *move* actions when a location has to be visited twice. There is never a need to visit a location ℓ more than twice. A location may only need to be visited twice if it is part of a cycle in the graph defined by the “passenger edges”. The shortest such cycle consists of just two locations. This is exactly the case in the task family described above. In that case, one in seven actions can be saved in a relaxed plan. If some passenger origin or destination locations are shared,

the ratio of actions that can be skipped in a relaxed plan only decreases. The same holds if the length of the “passenger cycles” increases. Therefore, at least 6 out of 7 actions from an optimal plan are also needed in an optimal relaxed plan. Consequently, $\alpha(h^+, \text{MICONIC-STRIPS}) = 6/7$.

MICONIC-SIMPLE-ADL Assume that there are n floors and each floor is origin or destination of at least one passenger. Let K^* be the size of a minimum feedback vertex set for the dependency graph given by the passengers’ origin-destination arcs. Then, in an optimal plan, there are $n + K^*$ or $n + K^* - 1$ *move* and $n + K^*$ *stop* actions. In a corresponding relaxed plan, the $n + K^*$ *stop* actions are still necessary, but $n - 1$ *move* actions are sufficient as the *lift-at*(f) propositions once added are never removed. This gives an accuracy bounded by $(2n + K^* - 1)/(2n + 2K^*) \geq 1 - (n + 1)/(2n + 2n) \rightarrow 3/4$. This bound is reached if the dependency graph is complete, which can be achieved for tasks of any size by including $n(n - 1)$ passengers in the task that need to move from each floor to each other floor. We thus get $\alpha(h^+, \text{MICONIC-SIMPLE-ADL}) = 3/4$.

SCHEDULE The objective in the SCHEDULE domain is to change certain attributes of (some of) the parts to work on, which can include their *shape*, *surface condition* and *color*. Parts also have a *temperature*, but there are no goals defined on that. We consider states s_n defined as follows: All machines and parts are currently available, and there are n parts which are currently polished, colored blue, and have no particular shape. The goal is to have all these parts polished, blue, and in cylindrical shape. Observe that in the SCHEDULE domain there are two types of actions, namely transformation actions and time-step actions, but in a relaxed plan the latter are never necessary. In particular, a relaxed plan for s_n simply applies the *do-lathe* operator to each part, so $h^+(s_n) = n$. An actual optimal plan must however apply a sequence of three operators for each part, including the *do-lathe* operator, the *do-polish* operator, and either *do-spray-paint* or *do-immersion-paint*. The latter are necessary because the *do-lathe* operator, which is necessary for obtaining the cylindrical shape, removes the polish and color from the part. (There are certain ordering constraints between these operators, but we do not discuss them in detail.) Moreover, $n + 1$ time-step actions are necessary: n time-step actions are needed until all parts have been lathed and the last lathed part is available again, and then one more time-step action is needed to polish and paint the last part (a time-step action is not needed at the end of the plan). Thus, $\lim_{n \rightarrow \infty} \alpha(s_n) = \lim_{n \rightarrow \infty} n/(4n + 1) = 1/4$.

We now show that the h^+ heuristic is never worse than $1/4h^* + d$. For this, consider a state s where n parts need to be processed (i. e., have a defined goal that is not already satisfied). Clearly, $h^+(s) \geq n$. To obtain an upper bound on $h^*(s)$, we describe a particular plan for s : First, classify the parts into different groups according to their current and goal attributes. The set \mathcal{C} of possible classes is constant, i. e., fixed for all SCHEDULE tasks. For each class $C \in \mathcal{C}$, we can find a sequence of at most 3 transformation actions that creates the goal attributes for the parts in C . (This is always possible.) If there are m_C objects in the class, all parts in

C can then be transformed into their goal states by using at most $3m_C$ transformation actions and at most $m_C + 2$ time-step actions, after which all machines are available again. Thus, altogether we need no more than $\sum_{C \in \mathcal{C}} (4m_C + 2) = 4n + 2|\mathcal{C}| = 4n + d$ actions, where $d := 2|\mathcal{C}|$. For our analysis, the additive constant does not matter, so we get $\alpha(h^+, \text{SCHEDULE}) = 1/4$.

SATELLITE Only *pointing*, *power-avail*, *calibrated* and *power-on* propositions appear in delete lists. Consider the family of states $(s_n)_{n \in \mathbb{N}}$ where s_n is described by one satellite with n cameras, each only supporting one mode which is different for each camera, all cameras having the same calibration target (different from the current pointing direction) and the same image target (different from the calibration target and the current pointing direction) in their respective modes. In an optimal plan, six actions per image are necessary, namely powering the i -th camera on, turning towards the calibration target, calibrating, turning towards the image target, taking the image, and switching the i -th camera off. (The last step can be omitted for the last image.) In an optimal relaxed plan, only two turn actions and no *switch-off* actions are necessary. Otherwise, it is identical to the non-relaxed plan. The *power-on*, *calibrate*, and *take-image* actions are still necessary. Therefore, in the limit, optimal relaxed plans are merely half as long as non-relaxed optimal plans, and $\lim_{n \rightarrow \infty} \alpha(s_n) = 1/2$.

For the other direction, i. e., $h^+(s) \geq 1/2 h^*(s)$ for all states s , we first observe that there exists an optimal relaxed plan that only uses a camera after it has been powered on and calibrated and before the next camera has been powered on and calibrated. (For example, first move all *take-image* actions to the end of the plan, then group them by the used camera, then move the *power-on* and *calibrate* actions before each corresponding group of *take-image* actions.) This optimal relaxed plan can then be transformed into a plan for the original task by adding an appropriate *power-off* action before each *power-on* action in order to ensure power availability, and adding *turn-to* actions before each *calibrate* and *take-image* action. Thus, optimal plans are at most twice as long as optimal relaxed plans. We conclude that $\alpha(h^+, \text{SATELLITE}) = 1/2$.

Synopsis In all considered domains, the h^+ heuristic yields results within a constant factor of the perfect heuristic values. The concrete factors vary from domain to domain.

Accuracy of h^m

For the accuracy of the h^m family of heuristics, we get the same result for all domains, namely $\alpha(h^m, \mathcal{D}) = 0$. The common reason for this is that we can find families of states $(s_n)_{n \in \mathbb{N}}$ such that $h^*(s_n)$ is unbounded as n approaches infinity, whereas there exists a function f such that $h^m(s_n) \leq f(m)$ for all $m \in \mathbb{N}$. For each fixed $m \in \mathbb{N}$, $f(m)$ is fixed and $h^*(s_n)$ can get arbitrarily large, so the accuracy of the h^m heuristic tends towards zero.

Note that $h^m(s_n)$ can always be bounded from above by the cost of reaching an m -elementary subgoal set from s_n , maximized over all such subgoals. Thus, it suffices to show that for the states we consider, every m -elementary subgoal

set can be reached in a number of steps that only depends on m , but not on n . (We are interested in behavior in the limit, so we can always safely assume $n \geq m$.)

The unboundedness of $h^*(s_n)$ for growing n is trivial for the states we present, so we will not mention it explicitly.

GRIPPER From any GRIPPER state, it is possible to move m balls to the goal room with at most $4m + 1$ actions (drop a ball to free a gripper if necessary, then move, pick up, move, drop for each ball).

LOGISTICS Each subgoal set of size m can be reached in at most $12m$ steps, up to 12 for each goal (move truck, pick up, move truck, drop; move airplane, pick up, move airplane, drop; move truck, pick up, move truck, drop).

BLOCKSWORLD Consider a family of states s_n with n blocks on the table to be stacked into a single tower. Then each m -elementary subgoal can be reached in $2m$ steps.

MICONIC-STRIPS and MICONIC-SIMPLE-ADL Consider states with n passengers to be served with n different origin floors, different from the current elevator location. All subgoals of size m , i. e., transporting m passengers to their goals, can be reached in $4m$ steps (for each passenger: move to origin, pick up/stop, move to destination, drop/stop).

SCHEDULE Consider states with n parts to be processed. Any subset of m parts can be processed in $6m$ steps (each part needs up to three transformations as argued earlier, and each can be preceded by a time-step action to guarantee availability of the machines).

SATELLITE Consider states s_n with n remaining image goals. Any subset of m goals can be satisfied in $6m$ steps (for each objective: power off an instrument if necessary, power on an instrument, turn to calibration target, calibrate, turn to objective, take image).

Synopsis In all considered domains, there are families of increasing-size tasks with h^* values roughly proportional to n and h^m values roughly proportional to m . As h^* increases indefinitely for growing n , whereas h^m never exceeds a domain-dependent constant which is a function of m , the accuracy of h^m tends to zero in all domains.

Accuracy of PDB Heuristics

The accuracy of the h^{PDB} family depends on the choice of the pattern. We assume that the PDB has a size limit polynomial in the input size n , say $O(n^m)$. Then a pattern may contain no more than $O(\log n^m) = O(m \log n) = O(\log n)$ state variables in order to respect the limit. This implies that $h^{\text{PDB}}(s)$ cannot be greater than the cost of reaching the most expensive subgoal set of cardinality $O(\log n)$.

As the proofs in the previous section showed, in all considered domains there exist families of states $(s_n)_{n \in \mathbb{N}}$ for which h^* grows linearly with n but m -elementary subgoal costs only grow linearly with m . Since $O(\log n)/\Omega(n) \rightarrow 0$ for $n \rightarrow \infty$, we obtain the same results for h^{PDB} as for h^m : for all domains, $\alpha(h^{\text{PDB}}, \mathcal{D}) = 0$. Note, however, that while for h^m , the heuristic estimates were always bounded by constants for a fixed value of m , for PDB heuristics with suitably chosen patterns they are bounded by values that grow

logarithmically with n . In that sense, PDB heuristics are more powerful than h^m in these domains.

Accuracy of Additive PDB Heuristics

As in the previous section, we need to bound the size of *individual* patterns by $O(\log n)$ for tasks of size n . However, for additive PDB heuristics, multiple such patterns are considered, and their values summed. We remark that since patterns in these domains need to be disjoint to be additive, there cannot be more than n pairwise additive patterns for a task with n state variables, so there is no need to impose a bound on the number of patterns to ensure polynomial time and space requirements.

GRIPPER We can restrict attention to states s_n with n balls in the initial room, the robot in the initial room, and no balls currently carried. Violating these restrictions only changes overall costs by a constant.

Using an additive PDB with singleton patterns for each ball location variable, $h_{\text{add}}^{\text{PDB}}(s_n) = 2n$, whereas the optimal cost is $3n - 1$ if n is even and $3n$ if n is odd. In either case, the accuracy tends towards $2/3$, so we get $\alpha(h_{\text{add}}^{\text{PDB}}, \text{GRIPPER}) \geq 2/3$.

On the other hand, this value cannot be asymptotically improved for large tasks with any pattern collection: At most one pattern, which can represent only $O(\log n)$ many balls, can include the robot location and hence estimate the cost for the balls in the pattern accurately. All other patterns must ignore the movement costs for the robot. We thus obtain $h_{\text{add}}^{\text{PDB}}(s_n) = 2n + O(\log n)$ compared to $h^*(s_n) = 3n$, so that $\alpha(s_n) = (2n + O(\log n))/3n$ approaches $2/3$ in the limit. Therefore, we have $\alpha(h_{\text{add}}^{\text{PDB}}, \text{GRIPPER}) = 2/3$.

LOGISTICS Again, by using singleton patterns for each goal variable, we obtain a heuristic which accurately captures all *pick-up* and *drop* actions needed to reach the goal. An optimal solution to a LOGISTICS task never requires more movements than *pick-up* and *drop* actions combined, so the accuracy is at least $1/2$.

To prove a lower bound, consider the family $(s_n)_{n \in \mathbb{N}}$ of tasks defined as follows. There is a single truck in a city with $2n + 1$ locations. There are n packages, all to be moved within this city. All initial and goal locations of the packages are different from each other and from the initial truck location. The optimal cost for this task is clearly $h^*(s_n) = 4n$. Similar to the GRIPPER proof, only one pattern in the pattern collection, which includes no more than $O(\log n)$ package variables, can include the variable representing the truck location. For all packages not included in this pattern, only the pick-up and drop actions are counted by $h_{\text{add}}^{\text{PDB}}(s_n)$. Thus, $\alpha(s_n) = (2n + O(\log n))/4n$, approaching $1/2$ in the limit. We thus get $\alpha(h_{\text{add}}^{\text{PDB}}, \text{LOGISTICS}) = 1/2$.

BLOCKSWORLD Consider the state s_n consisting of a stack B_1, \dots, B_n, B_{n+1} to be transformed into the stack B_1, \dots, B_{n+1}, B_n . In other words, the two bottommost blocks must be swapped. The true cost is $4n$.

There are only two variables whose current values differ between s_n and the goal, and hence all but two patterns in the additive PDB heuristic must assign heuristic values of

0 (since, relative to these patterns, the goal has already been reached). Thus, at most two patterns contribute to $h_{\text{add}}^{\text{PDB}}$, and again their magnitude is bounded by $O(\log n)$ since they can only incorporate $O(\log n)$ variables and BLOCKSWORLD tasks of size k can be solved with $O(k)$ steps.

Therefore, additive PDB heuristics for BLOCKSWORLD can get arbitrarily inaccurate. We obtain an accuracy of $\alpha(h_{\text{add}}^{\text{PDB}}, \text{BLOCKSWORLD}) = 0$.

MICONIC-STRIPS This domain is a special case of LOGISTICS, so from the result for that domain, we can conclude $\alpha(h_{\text{add}}^{\text{PDB}}, \text{MICONIC-STRIPS}) \geq 1/2$. On the other hand, the family of tasks used for showing the LOGISTICS upper bound are also MICONIC-STRIPS tasks, so $\alpha(h_{\text{add}}^{\text{PDB}}, \text{MICONIC-STRIPS}) = 1/2$.

MICONIC-SIMPLE-ADL We assume an encoding with one elevator location variable and one variable for each passenger indicating its status (at origin, boarded, or served). Let s_n be a state with n passengers, all at their origin floor, which is the same for all passengers and identical to the current elevator floor. The destination floors are different for each passenger. The optimal plan length is $2n + 1$. Since stopping at the start floor affects all passenger variables, no two passengers can appear in different patterns or the additivity requirement would be violated. Hence, only one pattern can contain passengers and thus contribute to the $h_{\text{add}}^{\text{PDB}}$ value. Since the size of the pattern is bounded by $O(\log n)$, we have $h_{\text{add}}^{\text{PDB}}(s_n) = O(\log n)$ and hence $\alpha(h_{\text{add}}^{\text{PDB}}, \text{MICONIC-SIMPLE-ADL}) = 0$.

SCHEDULE Consider the pattern collection where there is one pattern for each part, containing all attribute variables describing that part (i. e., its color, surface condition, shape, temperature, and has-hole variables). These patterns are all additive and only comprise a constant number of variables each. With this collection, an additive PDB heuristic captures the costs of all necessary transformation actions correctly, but ignores all time-step actions. In general, there are never more time-step than other actions (because two time-step actions in sequence can be collapsed) and there never is a time-step action at the end of an optimal plan, so $h_{\text{add}}^{\text{PDB}}(s) \geq 1/2 h^*(s)$ for all states s .

On the other hand, there exist state families $(s_n)_{n \in \mathbb{N}}$ for which this bound is tight. In particular, consider states where n parts need to be polished, the polisher and the parts are available, and there are no further goals. An optimal solution requires $2n - 1$ actions, whereas $h_{\text{add}}^{\text{PDB}}(s) = n$ for the pattern collection we described. By using a different pattern collection, the result cannot be improved significantly; similar to the previous proofs, we could only capture the interaction between the polish and time-step actions for logarithmically many parts without exceeding the pattern size bound. Consequently, $\alpha(h_{\text{add}}^{\text{PDB}}, \text{SCHEDULE}) = 1/2$.

SATELLITE As argued in the section on h^m , the optimal solution length from any SATELLITE state with n goals is at most $6n$. An additive PDB heuristic with singleton patterns for the goal variables achieves a heuristic estimate of n , so $\alpha(h_{\text{add}}^{\text{PDB}}, \text{SATELLITE}) \geq 1/6$.

Domain	h^+	h^m	h^{PDB}	$h_{\text{add}}^{\text{PDB}}$
GRIPPER	2/3	0	0	2/3
LOGISTICS	3/4	0	0	1/2
BLOCKSWORLD	1/4	0	0	0
MICONIC-STRIPS	6/7	0	0	1/2
MICONIC-SIMPLE-ADL	3/4	0	0	0
SCHEDULE	1/4	0	0	1/2
SATELLITE	1/2	0	0	1/6

Figure 1: Accuracy of the h^+ , h^m , and (additive) PDB heuristics in selected planning domains.

For the other direction, consider a state s_n with a single satellite and two pointing directions, A and B . There are n image modes, and the satellite has n^2 instruments supporting one mode each (n for each mode). A is the calibration target of all instruments, and B is the initial pointing direction of the satellite. The goal is to take images of B in each mode. An optimal plan has length $6n - 1$. The key observation is that a pattern for a given image goal cannot include variables for each instrument that can take that image, because there are $n > O(\log n)$ such instruments. Hence, in the pattern database abstraction, images can be taken by an uncalibrated, powered-off instrument, so that only the *take-image* action is counted for each goal. We get $\alpha(s_n) = n/6n = 1/6$, and hence $\alpha(h_{\text{add}}^{\text{PDB}}, \text{SATELLITE}) = 1/6$.

Synopsis Except for the BLOCKSWORLD and MICONIC-SIMPLE-ADL domains, where all additive PDB heuristic can perform arbitrarily badly, additive PDB heuristics with suitably chosen patterns do not perform worse than a constant factor times the perfect heuristic value.

Conclusion

As far as we are aware, we have presented the first detailed analysis of domain-specific accuracy results for a number of popular planning heuristics (summarized in Fig. 1). One (maybe unsurprising, but still interesting) observation is that the h^m heuristic family and plain PDB heuristics become arbitrarily inaccurate as task size increases.

With the exception of the SCHEDULE domain, we generally got the best results for the h^+ heuristic. However, in interpreting this result, one should not forget that this heuristic is NP-hard to compute in general, which makes additive PDB heuristics attractive.

For these, we remark that the problem of optimizing the pattern collections – a task we performed by hand in our analysis – is of critical importance for obtaining good performance. On the other hand, with the sole exception of the SCHEDULE domain, all our results already hold when using singleton patterns exclusively. Moreover, current techniques for the automatic selection of pattern collections (Haslum et al. 2007) have the potential to improve on this baseline.

Finally, we point out some open issues. One important omission is that we did not discuss additive h^m heuristics and restricted the discussion of additive PDBs to the case of a single sum. Extending our analysis beyond these restrictions is an interesting and important avenue for future work.

Acknowledgments

Carmel Domshlak and Michael Katz suggested the proof that additive pattern database cannot achieve a better bound than $1/6$ in the SATELLITE domain. They also provided some very helpful comments on an earlier version of this paper.

This work was partly supported by the German Research Council (DFG) as part of the Transregional Collaborative Research Center “Automatic Verification and Analysis of Complex Systems” (SFB/TR 14 AVACS). See <http://www.avacs.org/> for more information.

References

- Bacchus, F. 2001. The AIPS’00 planning competition. *AI Magazine* 22(3):47–56.
- Bonet, B., and Geffner, H. 2001. Planning as heuristic search. *AIJ* 129(1):5–33.
- Bylander, T. 1994. The computational complexity of propositional STRIPS planning. *AIJ* 69(1–2):165–204.
- Culberson, J. C., and Schaeffer, J. 1998. Pattern databases. *Computational Intelligence* 14(3):318–334.
- Edelkamp, S., and Helmert, M. 2001. The model checking integrated planning system (MIPS). *AI Magazine* 22(3):67–71.
- Haslum, P., and Geffner, H. 2000. Admissible heuristics for optimal planning. In *Proc. AIPS 2000*, 140–149.
- Haslum, P.; Botea, A.; Helmert, M.; Bonet, B.; and Koenig, S. 2007. Domain-independent construction of pattern database heuristics for cost-optimal planning. In *Proc. AAAI 2007*, 1007–1012.
- Haslum, P.; Bonet, B.; and Geffner, H. 2005. New admissible heuristics for domain-independent planning. In *Proc. AAAI 2005*, 1163–1168.
- Helmert, M.; Haslum, P.; and Hoffmann, J. 2007. Flexible abstraction heuristics for optimal sequential planning. In *Proc. ICAPS 2007*, 176–183.
- Helmert, M. 2008. *Understanding Planning Tasks – Domain Complexity and Heuristic Decomposition*, volume 4929 of *LNAI*. Springer-Verlag.
- Hoffmann, J., and Edelkamp, S. 2005. The deterministic part of IPC-4: An overview. *JAIR* 24:519–579.
- Hoffmann, J. 2005. Where ‘ignoring delete lists’ works: Local search topology in planning benchmarks. *JAIR* 24:685–758.
- Kautz, H., and Selman, B. 1996. Pushing the envelope: Planning, propositional logic, and stochastic search. In *Proc. AAAI 1996*, 1194–1201.
- Kautz, H., and Selman, B. 1999. Unifying SAT-based and graph-based planning. In *Proc. IJCAI 1999*, 318–325.
- Long, D., and Fox, M. 2003. The 3rd International Planning Competition: Results and analysis. *JAIR* 20:1–59.
- McDermott, D. 1996. A heuristic estimator for means-ends analysis in planning. In *Proc. AIPS 1996*, 142–149.
- McDermott, D. 2000. The 1998 AI Planning Systems competition. *AI Magazine* 21(2):35–55.