Heuristics and Symmetries in Classical Planning: Additional Proofs

Alexander Shleyfman
Technion, Haifa, Israel
alesh@technion.ac.il

Michael Katz
IBM Haifa Research Lab, Israel
katzm@il.ibm.com

Malte Helmert and **Silvan Sievers** and **Martin Wehrle**
University of Basel, Switzerland
{malte.helmert,silvan.sievers,martin.wehrle}@unibas.ch

Technical Report CS-2014-006

Symmetries of the State Transition Graph – Theorem 1 of the Main Paper

Theorem 1 *Let Π be a planning task, let s be one of its states, let π be a sequence of operators of Π, and let σ be a symmetry of $\mathcal{T}_\Pi$. Then π is a plan for s iff $\sigma(\pi)$ is a plan for $\sigma(s)$, and the two plans have the same cost.*

Proof: We only need to show one direction, since the inverse element σ^{-1} is also a symmetry. Let $\pi = o_1 \cdot \ldots \cdot o_k$ be a plan for s and let $s_0, \ldots, s_k$ be the sequence of states such that $s_0 = s$ and for $1 \leq i \leq k$, $s_i = s_{i-1}[\![o_i]\!]$. Let $\mathcal{T}_\Pi = \langle S, E\rangle$. Then we have $\langle s_{i-1}, s_i; o_i\rangle \in E$ for all $1 \leq i \leq k$ and therefore, by the definition of symmetries we have $\langle \sigma(s_{i-1}), \sigma(s_i); \sigma(o_i)\rangle \in E$ and $C(\sigma(o_i)) = C(o_i)$. Thus, $\sigma(\pi)$ is a sequence of actions that is applicable in $\sigma(s)$ and ends with the state $\sigma(s_k)$. Since s_k is a goal state, by the definition of symmetries $\sigma(s_k)$ is a goal state. Furthermore, $C(\sigma(\pi)) = \sum_{i=1}^{k} C(\sigma(o_i)) = \sum_{i=1}^{k} C(o_i) = C(\pi)$. □

Structural Symmetries – Theorem 2 of the Main Paper

For convenience, we repeat here the definition of structural symmetries.

Definition 1 *Let $\Pi = \langle P, O, I, G, C\rangle$ be a* STRIPS *planning task. A permutation σ of Π is a **structural symmetry** if*
- $\sigma(P) = P$
- $\sigma(O) = O$, *and for all $o \in O$:*
  - $pre(\sigma(o)) = \sigma(pre(o))$
  - $add(\sigma(o)) = \sigma(add(o))$
  - $del(\sigma(o)) = \sigma(del(o))$
  - $C(\sigma(o)) = C(o)$
- $\sigma(G) = G$

Lemma 1 *Let Π be a planning task. The set of structural symmetries of Π with the permutation composition operation is a group.*

Proof: Let Σ be the set of structural symmetries of $\Pi = \langle P, O, I, G, C\rangle$. To show that a set of permutations of a finite set forms a group under composition, it is sufficient to show that it is nonempty and that it is closed under composition (if $\sigma_1 \in \Sigma$ and $\sigma_2 \in \Sigma$, then $\sigma_1 \circ \sigma_2 \in \Sigma$). For the former, it is easy to verify that the identity function id satisfies the conditions of Definition 1.

For the latter, let $\sigma_1, \sigma_2 \in \Sigma$, and consider $\sigma := \sigma_1 \circ \sigma_2$. We have $\sigma(P) = \sigma_1(\sigma_2(P)) = \sigma_1(P) = P$, where we use, in sequence: the definition of σ, the fact that σ_2 is a structural symmetry, and the fact that σ_1 is a structural symmetry. In the same way, we obtain $\sigma(O) = O$ and $\sigma(G) = G$. Let $o \in O$. We have $pre(\sigma(o)) = pre(\sigma_1(\sigma_2(o)) = \sigma_1(pre(\sigma_2(o))) = \sigma_1(\sigma_2(pre(o))) = \sigma(pre(o))$, where we use, in sequence: the definition of σ, the fact that σ_1 is a structural symmetry, the fact that σ_2 is a structural symmetry, and the definition of σ. In the same way, we obtain $add(\sigma(o)) = \sigma(add(o))$ and $del(\sigma(o)) = \sigma(del(o))$. Finally, we get $C(\sigma(o)) = C(\sigma_1(\sigma_2(o))) = C(\sigma_2(o)) = C(o)$, where we use, in sequence: the definition of σ, the fact that σ_1 is a structural symmetry, and the fact that σ_2 is a structural symmetry. □

Lemma 2 *Let Π be a planning task and σ be a structural symmetry of Π. Let α be the mapping of $\mathcal{T}_\Pi = \langle S, E\rangle$ induced by σ, that is $\alpha(s) = \sigma(s)$ and $\alpha(\langle s, s'; o\rangle) = \langle \sigma(s), \sigma(s'); \sigma(o)\rangle$. Then α is an automorphism of $\mathcal{T}_\Pi$.*

Proof: We must show for all states s, s' and operators o:
1. $\langle s, s'; o\rangle \in E$ iff $\langle \sigma(s), \sigma(s'); \sigma(o)\rangle \in E$
2. $C(\sigma(o)) = C(o)$
3. s is a goal state iff $\sigma(s)$ is a goal state

Property 1: Let $\langle s, s'; o\rangle \in E$. Because o is applicable in s, we have $pre(o) \subseteq s$ and hence $\sigma(pre(o)) \subseteq \sigma(s)$. Because σ is a structural symmetry, we have $\sigma(pre(o)) = pre(\sigma(o))$ and hence $pre(\sigma(o)) \subseteq \sigma(s)$, showing that $\sigma(o)$ is applicable in $\sigma(s)$.

We have $s' = (s \setminus del(o)) \cup add(o)$, and hence

$$\begin{aligned}\sigma(s') &= \sigma((s \setminus del(o)) \cup add(o)) \\ &= \sigma(s \setminus del(o)) \cup \sigma(add(o)) \\ &= (\sigma(s) \setminus \sigma(del(o))) \cup \sigma(add(o)) \qquad (*) \\ &= (\sigma(s) \setminus del(\sigma(o))) \cup add(\sigma(o)), \qquad (**)\end{aligned}$$

which establishes that $\sigma(s')$ is the successor of $\sigma(s)$ under operator $\sigma(o)$. Here, (*) uses the fact that σ is injective, and (**) uses the fact that σ is a structural symmetry.

Together, this shows the "⇒" part of property 1. For the opposite direction, let $\langle \sigma(s), \sigma(s'); \sigma(o)\rangle \in E$ for a structural symmetry σ. From Lemma 1, σ^{-1} is also a structural symmetry, and applying the "⇒" result, we get $\langle \sigma^{-1}(\sigma(s)), \sigma^{-1}(\sigma(s')); \sigma^{-1}(\sigma(o))\rangle \in E$, and hence $\langle s, s'; o\rangle \in E$, proving the result.

Property 2: This is true by definition of structural symmetries.

Property 3: If s is a goal state, then $G \subseteq s$, where G is the set of goal facts. This implies $\sigma(G) \subseteq \sigma(s)$ and hence $G \subseteq \sigma(s)$ (because $\sigma(G) = G$ by definition of structural symmetries), and hence $\sigma(s)$ is a goal state. The opposite direction follows in the same way as with Property 1 by considering the inverse symmetry σ^{-1}. □

We now seek to relate structural symmetries to PDG symmetries. To this end, we next repeat the definition of PDGs.

Definition 2 *Let $\Pi = \langle P, O, I, G, C\rangle$ be a* STRIPS *planning task. The* ***problem description graph*** *(PDG) of Π is the colored digraph $\langle N, E\rangle$ with nodes*

$$N = \bigcup_{p \in P} \{v_p, v_p^{\mathrm{T}}, v_p^{\mathrm{F}}\} \cup \{v_o \mid o \in O\},$$

node colors

$$col(v) = \begin{cases} 1 & \text{if } v = v_p^{\mathrm{T}}, p \in G \\ 2 + C(o) & \text{if } v = v_o, o \in O \\ 0 & \text{otherwise} \end{cases}$$

and edges

$$E = \bigcup_{p \in P} \{\langle v_p, v_p^{\mathrm{T}}\rangle, \langle v_p, v_p^{\mathrm{F}}\rangle\} \cup \bigcup_{o \in O} \left(E_o^{pre} \cup E_o^{add} \cup E_o^{del}\right),$$

where

$$E_o^{pre} = \{\langle v_p^{\mathrm{T}}, v_o\rangle \mid p \in pre(o)\},$$

$$E_o^{add} = \{\langle v_o, v_p^{\mathrm{T}}\rangle \mid p \in add(o)\},$$

$$E_o^{del} = \{\langle v_o, v_p^{\mathrm{F}}\rangle \mid p \in del(o)\}.$$

Next, we define the notion of *PDG symmetries*.

Definition 3 *Let Π a planning task. A* ***PDG symmetry*** *of Π is an automorphism of the problem description graph $\langle N, E\rangle$ of Π, i.e., a function $\alpha : N \to N$ such that:*

1. *α is bijective*
2. *$col(\alpha(n)) = col(n)$ for all $n \in N$*
3. *$\langle n, n'\rangle \in E$ iff $\langle \alpha(n), \alpha(n')\rangle \in E$*

PDG symmetries induce transition graph symmetries. To formally define these, first observe that every PDG symmetry must map each node of the form v_p ($p \in P$) to a variable of the same form (because these nodes are the only nodes with color 0 and outgoing arcs leading to nodes with color 0 or 1, and automorphisms must preserve such connectivity properties). Having established this, it is easy to show that if $\alpha(v_p) = v_{p'}$, then one of the two possibilities must hold:

1. $\alpha(v_p^{\mathrm{T}}) = \alpha(v_{p'}^{\mathrm{T}})$ and $\alpha(v_p^{\mathrm{F}}) = \alpha(v_{p'}^{\mathrm{F}})$, or
2. $\alpha(v_p^{\mathrm{T}}) = \alpha(v_{p'}^{\mathrm{F}})$ and $\alpha(v_p^{\mathrm{F}}) = \alpha(v_{p'}^{\mathrm{T}})$.

This means that PDG symmetries induce a permutation on the propositions of the task, and they additionally have the ability of swapping between truth and falsity of the mapped-to propositions. In other words, they can be viewed as permutations of the *literals* of a planning task where literals of the same variable are always mapped to literals of the same variable, but true and false literals may be reversed by the mapping.

Definition 4 *Let α be a PDG symmetry of a planning task $\Pi = \langle P, O, I, G, C\rangle$ with states S.*

The ***transition graph symmetry*** *induced by α is the function σ defined on $S \cup O$ as follows:*

$$\sigma(s) = \{p \mid \exists p' \in P : (p' \in s \wedge \alpha(v_{p'}^{\mathrm{T}}) = v_p^{\mathrm{T}} \vee p' \notin s \wedge \alpha(v_{p'}^{\mathrm{F}}) = v_p^{\mathrm{T}})\}$$

$$\sigma(o) = o' \text{ whenever } \alpha(v_o) = v_{o'} \text{ for operators } o, o' \in O$$

It is not difficult to verify that σ is well-defined and satisfies the criteria of a transition graph symmetry. We do not show a formal proof because this follows from the earlier work on symmetries by Pochter, Zohar, and Rosenschein (2011).

Unlike structural symmetries, which operate on the level of *propositions*, PDG symmetries can be viewed as operating on *variable/value* assignments. As a consequence of this, if s is a state and σ is a transition graph symmetry induced by a PDG symmetries, it is possible that $|\sigma(s)| \neq |s|$, which can never happen with structural symmetries. However, we will see in Theorem 2 that this does not give PDG symmetries significantly more expressiveness than structural symmetries.

We first show that structural symmetries can be viewed as a special case of PDG symmetries.

Lemma 3 *Let σ be a structural symmetry of a planning task $\Pi = \langle P, O, I, G, C\rangle$, and let $\langle N, E\rangle$ be the PDG of Π.*

Define $\alpha : N \to N$ as follows:

1. *$\alpha(v_p) = v_{\sigma(p)}$ for all $p \in P$*
2. *$\alpha(v_p^{\mathrm{T}}) = v_{\sigma(p)}^{\mathrm{T}}$ for all $p \in P$*
3. *$\alpha(v_p^{\mathrm{F}}) = v_{\sigma(p)}^{\mathrm{F}}$ for all $p \in P$*
4. *$\alpha(v_o) = v_{\sigma(o)}$ for all $o \in O$*

Then α is a PDG symmetry that induces the same transition graph symmetry as σ.

Proof: We first prove that α is a PDG symmetry by verifying the three properties of Definition 2.

Property 1 (bijectivity of α): follows easily from $\sigma(P) = P$ and $\sigma(O) = O$

Property 2 (α preserves colors): for nodes of the form v_o with $o \in O$, we have $col(\alpha(v_o)) = col(v_{\sigma(o)}) = 2 + C(\sigma(o)) = 2 + C(o) = col(v_o)$, where we use that $C(\sigma(o)) = C(o)$ for structural symmetries σ. All other nodes have color 0 or 1 and are mapped to nodes that have color 0 or 1, so it is sufficient to show that $col(v) = 1$ iff $col(\alpha(v)) = 1$. Consider a node v with $col(v) = 1$. Then $v = v_p^{\mathrm{T}}$ for some $p \in G$. We obtain $\alpha(v) = v_{\sigma(p)}^{\mathrm{T}}$

with $\sigma(p) \in G$ (because $p \in G$ and $\sigma(G) = G$), and hence $col(\alpha(v)) = 1$. The opposite direction (showing that $col(\alpha(v)) = 1$ implies $col(v) = 1$) is analogous.

For Property 3 (α preserves arcs), it is sufficient to show the implication "If $e = \langle n, n' \rangle \in E$, then $\alpha(e) := \langle \alpha(n), \alpha(n') \rangle \in E$." The opposite direction then follows from a counting argument and the fact that α is a bijection. So consider the different kinds of arcs $e \in E$:

- If $e = \langle v_p, v_p^{\mathrm{T}} \rangle$, then $\alpha(e) = \langle \alpha(v_p), \alpha(v_p^{\mathrm{T}}) \rangle = \langle v_{\sigma(p)}, v_{\sigma(p)}^{\mathrm{T}} \rangle \in E$.
- The case of arcs of the form $\langle v_p, v_p^{\mathrm{F}} \rangle$ is analogous.
- If $e \in E_o^{pre}$ for some operator o, then $e = \langle v_p^{\mathrm{T}}, v_o \rangle$ for some proposition $p \in pre(o)$. Then $\sigma(p) \in \sigma(pre(o)) = pre(\sigma(o))$ (where the equality holds because σ is a structural symmetry), and hence $\alpha(e) = \langle \alpha(v_p^{\mathrm{T}}), \alpha(v_o) \rangle = \langle v_{\sigma(p)}^{\mathrm{T}}, v_{\sigma(o)} \rangle \in E_{\sigma(o)}^{pre}$.
- The case of arcs in E_o^{add} and of arcs in E_o^{del} is analogous.

It remains to show that σ and α induce the same transition system symmetry, i.e., that $\sigma(s) = \sigma'(s)$ for all states s and $\sigma(o) = \sigma'(o)$ for all operators o, where σ' is the transition system symmetry induced by α. This is obvious for operators o, so consider a state s.

By definition, we have $\sigma'(s) = \{p \mid \exists p' \in P : (p' \in s \wedge \alpha(v_{p'}^{\mathrm{T}}) = v_p^{\mathrm{T}} \vee p' \notin s \wedge \alpha(v_{p'}^{\mathrm{F}}) = v_p^{\mathrm{T}})\}$. The second case can never occur for the given PDG symmetry α because it never maps "negative literals" to "true literals", so for the given α, we obtain:

$$\begin{aligned}
\sigma'(s) &= \{p \mid \exists p' \in P : p' \in s \wedge \alpha(v_{p'}^{\mathrm{T}}) = v_p^{\mathrm{T}}\} \\
&= \{p \mid \exists p' \in P : p' \in s \wedge v_{\sigma(p')}^{\mathrm{T}} = v_p^{\mathrm{T}}\} \\
&= \{p \mid \exists p' \in P : p' \in s \wedge \sigma(p') = p\} \\
&= \{\sigma(p') \mid \exists p' \in P : p' \in s\} \\
&= \{\sigma(p') \mid p' \in s\} \\
&= \sigma(s).
\end{aligned}$$

□

We now show that the opposite conversion – from PDG symmetries to structural symmetries – is also possible if each proposition of the planning task occur as a precondition of some operator or as a goal.

Lemma 4 *Let $\Pi = \langle P, O, I, G, C \rangle$ be a planning task with $\bigcup_{o \in O} pre(o) \cup G = P$. Let α be a PDG symmetry of Π, and let $\langle N, E \rangle$ be the PDG of Π.*

Define $\sigma : P \cup O \to P \cup O$ as follows:

1. *$\sigma(p) = p'$ for all $p, p' \in P$ with $\alpha(v_p) = v_{p'}$.*
2. *$\sigma(o) = o'$ for all $o, o' \in O$ with $\alpha(v_o) = v_{o'}$*

Then σ is a structural symmetry that induces the same transition graph symmetry as α.

Proof: We need to establish that σ is well-defined, that it is a structural symmetry, and that it induces the same transition graph symmetry as α.

We first show that σ is well-defined:

- $N^P = \{v_p \mid p \in P\}$ are the only nodes with color 0 and no incoming edges in the PDG, and hence we must have $\alpha(N^P) = N^P$. This ensures that σ is well-defined on P. It also establishes $\sigma(P) = P$.
- Similarly, $N^O = \{v_o \mid o \in O\}$ are the only nodes with color strictly larger than 1, and hence we must have $\alpha(N^O) = N^O$. This ensures that σ is well-defined on O. It also establishes $\sigma(O) = O$. Moreover, we have $C(\sigma(o)) = col(v_{\sigma(o)}) - 2 = col(\alpha(v_o)) - 2 = col(v_o) - 2 = C(o)$ for all operators o.

We now show that σ satisfies the remaining properties of structural symmetries. From our previous observation regarding the way that PDG symmetries map a group of nodes referring to one variable to a group of nodes referring to one variable, for all $p \in P$ we must have $\alpha(v_p^{\mathrm{T}}) = v_{\sigma(p)}^{\mathrm{T}}$ or $\alpha(v_p^{\mathrm{T}}) = v_{\sigma(p)}^{\mathrm{F}}$. Under the restriction to Π in the statement of the lemma, the first case must always apply:

- Each $p \in P$ belongs to G or be a precondition of some operator. First consider $p \in G$. Then we must have $\alpha(v_p^{\mathrm{T}}) = v_{\sigma(p)}^{\mathrm{T}}$ because α must preserve the color of v_p^{T}, which is 1.
- Now consider p such that $p \in pre(o)$ for some operator o. Then we must have $\alpha(v_p^{\mathrm{T}}) = v_{\sigma(p)}^{\mathrm{T}}$ because v_p^{T} has an outgoing arc in the PDG (to v_o) and hence must be mapped to another node with an outgoing arc. Nodes of the form $v_{p'}^{\mathrm{F}}$ do not have outgoing arcs.

We have thus shown that α only maps "positive literals" to "positive literals" and consequently must map "negative literals" to "negative literal". The proof of the remaining properties of σ and the proof that α and σ define the same transition system symmetry is now a mechanical exercise along the same lines as the proof of the previous lemma. □

Theorem 2 *Let Π be a planning task. Then:*

1. *If σ is a structural symmetry of Π, then σ (viewed as a function on the states and operators of Π) is a transition graph symmetry of $\mathcal{T}_\Pi$.*
2. *The structural symmetries form a subgroup of $Aut(\mathcal{T}_\Pi)$.*
3. *Every structural symmetry of Π corresponds to a PDG symmetry of Π in the sense that they induce the same transition graph symmetry.*
4. *If each proposition of Π occurs as an operator precondition or in the goal, then every PDG symmetry of Π corresponds to a structural symmetry of Π in the sense that they induce the same transition graph symmetry.*

Proof: Statement 1 is due to Lemma 2. Statement 2 is due to Lemmas 2 and 1. Statement 3 is due to Lemma 3. Statement 4 is due to Lemma 4. □

Delete Relaxation – Corner Cases and Theorem 6 of the Main Paper

After the proof that the optimal delete relaxation heuristic h^+ is invariant under symmetry, the main paper notes (footnote 2):

To keep the presentation short, we gloss over some details here: the case of zero-cost actions and the case where the equations minimize or maximize over empty sets are discussed in the technical report.

We now discuss these cases.

Empty Sets and Infinities

The equations defining $h^{\max}$ contain two instances of minimization and maximization: Eq. (5) maximizes over the costs of achieving all facts in a given set. If the set is empty, this maximum can simply be defined to be 0. The more complicated instance is Eq. (3), which includes a minimization over the cost estimate for of all operators that achieve a given fact. If no operator exists that achieves a given fact, this is ill-defined.

A simple solution to avoid treating this case specially is to consider a slightly modified planning task Π' which is identical to the given planning task Π except that we add an additional operator o^{all}, without preconditions, which adds all facts of the task.

We set the cost of o^{all} large enough that it cannot affect the $h^{\max}$ estimate of any state with finite $h^{\max}$ value. For example, we can set $C(o^{\text{all}})$ to 1 plus the sum of operator costs of all other operators.[1] Then $h^{\max}$ in Π and Π' are identical for all states with finite $h^{\max}$ estimates in Π. Moreover, we can ignore the case of states with infinite $h^{\max}$ value because $h^{\max}(s) = \infty$ iff $h^{+}(s) = \infty$ for all states s, and we have already established earlier that h^{+} is invariant under structural symmetry.

It is easy to see that Π and Π' have essentially the same structural symmetries: every structural symmetry in Π' must map o^{all} to itself (because no other operator has the same cost), and it is easy to verify that σ is a structural symmetry of Π iff σ' defined as $\sigma'(o^{\text{all}}) = o^{\text{all}}$ and $\sigma'(x) = \sigma(x)$ for all $x \neq o^{\text{all}}$ is a structural symmetry of Π'.

Hence, we can in the following assume that we are working on the modified task Π' instead of Π. This means that the minimization in Eq. (3) is always over a non-empty set and that all numbers involved in the computation of $h^{\max}$ are finite.

Zero-Cost Operators

In the presence of zero-cost operators, the equations defining $h^{\max}$ and other delete relaxation heuristics can have multiple solutions, which complicates the analysis of the heuristic.

Let $\Pi = \langle P, O, I, G, C \rangle$ be a planning task. For $\epsilon > 0$, let Π_ϵ be the planning task obtained from Π by replacing the cost of 0-cost operators with ϵ. Thus, Π_ϵ has no 0-cost operators. Let $h^{\max}_\epsilon$ be the maximum heuristic for Π_ϵ. Then $h^{\max}_\epsilon$ is invariant under structural symmetry because we have established this property for $h^{\max}$ on tasks without 0-cost operators.

If we fix the planning task $\Pi = \langle P, O, I, G, C \rangle$ and one of its states s, we can vie $h^{\max}_\epsilon(s)$ as a function of s. From the well-known definition of $h^{\max}$ in terms of relaxed planning graphs, it is easy to see that $h^{\max}_\epsilon(s)$ can be expressed

[1]The same basic idea works for other delete relaxation heuristics like h^{add}, but we might need to assign a higher cost to o^{all}.

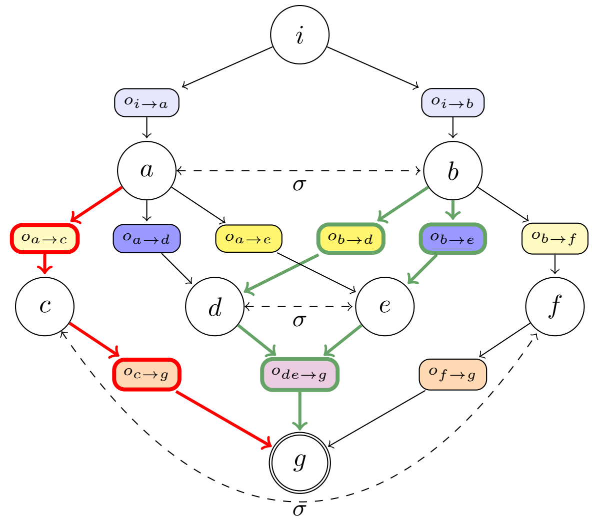

Figure 1: An example STRIPS planning task.

as a finite-size[2] expression whose components are all of the following type:

- maximum over a number of subexpressions,
- summation over a number of subexpressions, or
- the cost of a given operator.

If we take the operator costs as the variables in this expression, all ingredients are continuous functions of their components, and the composition of continuous functions is continuous.

Hence, the $h^{\max}$ value for a given state is a continuous function of the operator costs, and we get

$$h^{\max}(s) = \lim_{\epsilon \to 0} h^{\max}_\epsilon(s).$$

Therefore, showing that $h^{\max}$ is invariant under structural symmetry for operators with positive cost is sufficient for showing that the same also holds in the general case of non-negative cost.

FF Heuristic

Theorem 6 *1. There exist tie-breaking policies for which FF/$h^{\max}$ is not invariant under structural symmetry.*

2. There exist tie-breaking policies for which FF/h^{add} is not invariant under structural symmetry.

3. Let h^{FF} be a randomized variant of the FF heuristic where supporters are selected w.r.t. a heuristic that is invariant under structural symmetry (like $h^{\max}$ or h^{add}) and ties are broken uniformly randomly. This heuristic is invariant under structural symmetry in the sense that for all states s and structural symmetries σ, $h^{\text{FF}}(s)$ and $h^{\text{FF}}(\sigma(s))$ are identically distributed random variables.

[2]There must be a constant L such that the $h^{\max}$ costs in the relaxed planning graph "level off" after at most L layers, no matter how the operator costs are chosen. For example, setting L to the number of facts is always sufficient since at least one new fact must receive its final $h^{\max}$ value at each new layer where the relaxed planning graph has not leveled off yet.

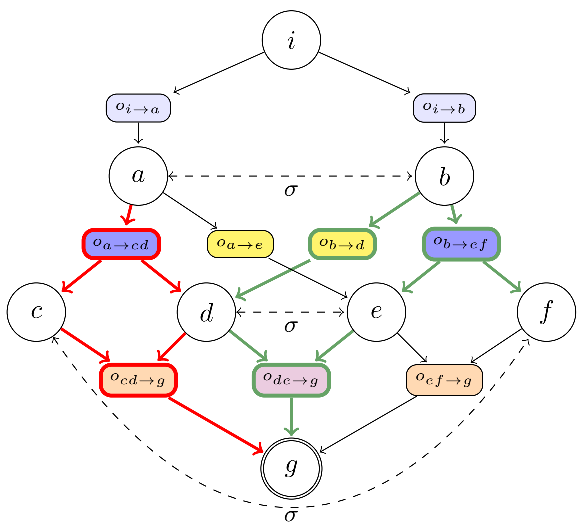

Figure 2: An example STRIPS planning task.

Proof: Statement 1. is shown by the example planning task Π_1 illustrated in Figure 1. The propositions are marked by circle nodes and operators by rounded nodes with preconditions/add effects as incoming/outgoing arcs. Formally, $\Pi = \langle P, O, I, G, C\rangle$ where $P = \{i, a, b, c, d, e, f, g\}$, with $I = \{i\}$ and $G = \{g\}$. Operators $O = \{o_{i\to a}, o_{i\to b}, o_{a\to c}, o_{a\to d}, o_{a\to e}, o_{b\to d}, o_{b\to e}, o_{b\to f}, o_{c\to g}, o_{de\to g}, o_{f\to g}\}$ have unit cost. Operators are denoted by their preconditions and add effects, $o_{pre\to add}$, and delete effects are empty for all operators. Let $s = \{i, a\}$. Note that both $o_{c\to g}$ and $o_{de\to g}$ can be best-supporters $o_g(s)$ of g, since $h^{\max}(c; s) = h^{\max}(\{d, e\}; s) = 1$. However, if $o_g(s) = o_{c\to g}$ is chosen, then $h^{\mathrm{FF}}(s) = C(o_{a\to c}) + C(o_{c\to g}) = 2$. Alternatively, if $o_g(s) = o_{de\to g}$, then $h^{\mathrm{FF}}(s) = C(o_{a\to d}) + C(o_{a\to e}) + C(o_{de\to g}) = 3$.

Our example task Π has one structural symmetry σ of order two, with mapping of propositions depicted by dotted edges, self loops left out. Mapping of operators by σ is depicted by colors, having operators with the same color mapped into each other. Note that for each $o_{x\to y} \in O$, we have $\sigma(o_{x\to y}) = o_{\sigma(x)\to\sigma(y)}$, and thus σ is indeed a structural symmetry of Π.

Now, let $s' = \sigma(s) = \{i, b\}$. Note that, similarly to s, we have $o_g(s') \in \{o_{f\to g}, o_{de\to g}\}$. Also, similarly to s, if $o_g(s') = o_{f\to g}$, then $h^{\mathrm{FF}}(s') = 2$, and if $o_g(s') = o_{de\to g}$, then $h^{\mathrm{FF}}(s') = 3$. Since nothing prevents a tie breaking where $o_g(s') \neq \sigma(o_g(s))$, there exists a tie breaking such that $h^{\mathrm{FF}}(s) \neq h^{\mathrm{FF}}(s')$. One such tie breaking and the resulting heuristic calculation for s and s' are depicted in Figure 1 by red and green colors, respectively.

Statement 2. is shown by the example planning task Π_2 illustrated in Figure 2. The notation we use here is similar to the notation of Statement 1. As in the previous example the operators $O = \{o_{i\to a}, o_{i\to b}, o_{a\to cd}, o_{a\to d}, o_{a\to ef}, o_{b\to d}, o_{b\to e}, o_{b\to f}, o_{cd\to g}, o_{de\to g}, o_{ef\to g}\}$ have unit cost.

Let $s = \{i, a\}$. Note that both $o_{cd\to g}$ and $o_{de\to g}$ can be best-supporters $o_g(s)$ of g, since $h^{\mathrm{add}}(c; s) + h^{\mathrm{add}}(d; s) = h^{\mathrm{add}}(d; s) + h^{\mathrm{add}}(e; s) = 2$. If $o_g(s) = o_{cd\to g}$ is chosen, then $h^{\mathrm{FF}}(s) = C(o_{a\to cd}) + C(o_{cd\to g}) = 2$. However, if $o_g(s) = o_{de\to g}$, then $h^{\mathrm{FF}}(s) = C(o_{a\to cd}) + C(o_{a\to e}) + C(o_{de\to g}) = 3$.

The symmetrical part of the example for the state $s' = \sigma(s) = \{i, b\}$ is done almost identically to Statement 1.

For statement 3., we must define *supp*, *plan* and h^{FF} as random variables, modeling the random tie-breaking between supporters with the same *opcost* values in Eq. 3 of the main paper. For this purpose, for every state s, we define a finite probability space with carrier set $\Omega(s)$, where each atomic event $\omega \in \Omega(s)$ corresponds to one of the possible choices for the best supporter function in state s. More precisely, let $\Pi = \langle P, O, I, G, C\rangle$. Then

$$\Omega(s) = \{\mathit{supp} : P \to O \mid \forall p \in P : \\ \mathit{opcost}(\mathit{supp}(p), s) = \min_{o\in O: p\in \mathit{add}(o)} \mathit{opcost}(o, s)\}$$

consists of all possible best supporter functions. We set $P(s, \omega) = \frac{1}{|\Omega(s)|}$ for all states s and all $\omega \in \Omega(s)$. In other words, for every state s, each best supporter function that minimizes *opcost* in s is assigned the same probability. This is equivalent to saying that each individual best supporter decision for a given fact $p \in P$ is made uniformly randomly. (The probability space $\Omega(s)$ can be viewed as a product of individual probability spaces that uniformly choose between the best supporters of each fact.)

The definition of the FF heuristic value depends on *propcost*, *opcost* and *setcost*, which remain non-probabilistic in this setting because they only depend on heuristic computations such as h^{add} or $h^{\max}$ (for the computation of supporters) that are not affected by tie-breaking. It also depends on *supp* and *plan*, which are now random variables, i.e., also depend on the random choice of $\omega \in \Omega(s)$. We reflect this by making *supp*, *plan* and h^{FF} itself functions of ω in addition to the arguments they previously depended on, obtaining:

$$\mathit{supp}(p, s, \omega) = \omega(p) \quad \text{if } p \notin s \tag{1}$$

$$\mathit{plan}(p, s, \omega) = \emptyset \quad \text{if } p \in s \tag{2}$$

$$\mathit{plan}(p, s, \omega) = \{\mathit{supp}(p, s, \omega)\} \cup \bigcup_{q\in \mathit{pre}(\mathit{supp}(p,s,\omega))} \mathit{plan}(q, s, \omega) \quad \text{if } p \notin s \tag{3}$$

$$h^{\mathrm{FF}}(s, \omega) = \sum_{o\in\bigcup_{q\in G} \mathit{plan}(q,s,\omega)} C(o) \tag{4}$$

(Note how the definition of $\mathit{supp}(p, s, \omega)$ captures that ω is the randomly selected best supporter function for state s, so $\omega(p)$ is the randomly chosen best supporter for fact p.)

For a given best supporter function $\omega : P \to O$ and structural symmetry σ, let $\sigma(\omega) : P \to O$ be the function that maps $\sigma(p)$ to $\sigma(o)$ whenever $\omega(p) = o$. Due to the symmetry results for h^{add} shown in the main paper, we have that whenever ω is a best supporter function for a given state s, $\sigma(\omega)$ is a best supporter function for $\sigma(s)$. Therefore, we get $\sigma(\Omega(s)) = \Omega(\sigma(s))$. Among other things, this shows that $\Omega(s)$ and $\Omega(\sigma(s))$ have the same size, and hence the probability of choosing best supporter function ω in state s is the same as the probability of choosing $\sigma(\omega)$ in $\sigma(s)$.

We now show how *supp* and *plan* interact with structural symmetries σ. For *supp*, we obtain:

$$\begin{aligned} \textit{supp}(\sigma(p), \sigma(s), \sigma(\omega)) &= \sigma(\omega)(\sigma(p)) \\ &= \sigma(\omega(p)) \\ &= \sigma(\textit{supp}(p, s, \omega)). \end{aligned}$$

Next, we define $\textit{plan}^\sigma$ as follows:

$$\textit{plan}^\sigma(p, s, \omega) = \sigma^{-1}(\textit{plan}(\sigma(p), \sigma(s), \sigma(\omega)))$$

For all $p \in s$, we have:

$$\begin{aligned} \textit{plan}^\sigma(p, s, \omega) &= \sigma^{-1}(\textit{plan}(\sigma(p), \sigma(s), \sigma(\omega))) \\ &= \emptyset, \end{aligned}$$

where the last equality holds because $\sigma(p) \in \sigma(s)$.

For all $p \notin s$, we have $\sigma(p) \notin \sigma(s)$, and we obtain:

$$\begin{aligned} &\sigma(\textit{plan}^\sigma(p, s, \omega)) \\ &= \sigma(\sigma^{-1}(\textit{plan}(\sigma(p), \sigma(s), \sigma(\omega)))) \\ &= \textit{plan}(\sigma(p), \sigma(s), \sigma(\omega)) \\ &= \{\textit{supp}(\sigma(p), \sigma(s), \sigma(\omega))\} \cup \bigcup_{q \in \textit{pre}(\textit{supp}(\sigma(p), \sigma(s), \sigma(\omega)))} \textit{plan}(q, \sigma(s), \sigma(\omega)) \\ &= \{\sigma(\textit{supp}(p, s, \omega))\} \cup \bigcup_{\sigma(q') \in \textit{pre}(\textit{supp}(\sigma(p), \sigma(s), \sigma(\omega)))} \textit{plan}(\sigma(q'), \sigma(s), \sigma(\omega)) \\ &= \{\sigma(\textit{supp}(p, s, \omega))\} \cup \bigcup_{\sigma(q') \in \textit{pre}(\sigma(\textit{supp}(p, s, \omega)))} \textit{plan}(\sigma(q'), \sigma(s), \sigma(\omega)) \\ &= \{\sigma(\textit{supp}(p, s, \omega))\} \cup \bigcup_{\sigma(q') \in \sigma(\textit{pre}(\textit{supp}(p, s, \omega)))} \textit{plan}(\sigma(q'), \sigma(s), \sigma(\omega)) \\ &= \{\sigma(\textit{supp}(p, s, \omega))\} \cup \bigcup_{q' \in \textit{pre}(\textit{supp}(p, s, \omega))} \textit{plan}(\sigma(q'), \sigma(s), \sigma(\omega)) \\ &= \{\sigma(\textit{supp}(p, s, \omega))\} \cup \bigcup_{q' \in \textit{pre}(\textit{supp}(p, s, \omega))} \sigma(\textit{plan}^\sigma(q', s, \omega)) \\ &= \sigma(\{\textit{supp}(p, s, \omega)\} \cup \bigcup_{q' \in \textit{pre}(\textit{supp}(p, s, \omega))} \textit{plan}^\sigma(q', s, \omega)) \end{aligned}$$

and therefore

$$\begin{aligned} &\textit{plan}^\sigma(p, s, \omega) \\ &= \{\textit{supp}(p, s, \omega)\} \cup \bigcup_{q' \in \textit{pre}(\textit{supp}(p, s, \omega))} \textit{plan}^\sigma(q', s, \omega). \end{aligned}$$

This shows that $\textit{plan}^\sigma$ satisfies the same recursive equations as *plan*. Because these equations uniquely define *plan*, we must thus have $\textit{plan}^\sigma = \textit{plan}$, and we can conclude:

$$\begin{aligned} \sigma(\textit{plan}(p, s, \omega)) &= \sigma(\textit{plan}^\sigma(p, s, \omega)) \\ &= \sigma(\sigma^{-1}(\textit{plan}(\sigma(p), \sigma(s), \sigma(\omega)))) \\ &= \textit{plan}(\sigma(p), \sigma(s), \sigma(\omega)). \end{aligned}$$

For the random variable denoting the FF heuristic, we thus obtain:

$$\begin{aligned} h^{\mathrm{FF}}(\sigma(s), \sigma(\omega)) &= \sum_{o \in \bigcup_{q \in G} \textit{plan}(q, \sigma(s), \sigma(\omega))} C(o) \\ &= \sum_{o \in \bigcup_{q \in G} \textit{plan}(\sigma(q), \sigma(s), \sigma(\omega))} C(o) \\ &= \sum_{o \in \bigcup_{q \in G} \sigma(\textit{plan}(q, s, \omega))} C(o) \\ &= \sum_{o' \in \bigcup_{q \in G} \textit{plan}(q, s, \omega)} C(\sigma(o')) \\ &= \sum_{o' \in \bigcup_{q \in G} \textit{plan}(q, s, \omega)} C(o') \\ &= h^{\mathrm{FF}}(s, \omega). \end{aligned}$$

Taking into account the bijection between $\omega \in \Omega(s)$ and $\sigma(\omega) \in \Omega(\sigma(s))$ discussed previously, this shows that $h^{\mathrm{FF}}(s)$ and $h^{\mathrm{FF}}(\sigma(s))$, seen as random variables, are identically distributed. This concludes the proof. □

Landmarks – Theorem 11 of the Main Paper

Theorem 11 *Let $\mathcal{L}$ be a landmark generation method that is invariant under structural symmetry, and let h be a heuristic such that $h(s)$ derives a heuristic estimate from $\mathcal{L}(s)$ using one of the following techniques:*

1. *counting landmarks (Richter and Westphal 2010)*
2. *summing the minimal operator costs of each landmark (Richter, Helmert, and Westphal 2008)*
3. *optimal cost partitioning (Karpas and Domshlak 2009)*
4. *uniform cost partitioning with or without special treatment of action landmarks (Karpas and Domshlak 2009)*
5. *hitting sets (Bonet and Helmert 2010)*

Then h is invariant under structural symmetry.

Proof: All parts except 4. are shown in the main paper, so we now discuss the remaining part 4. Let $\Pi = \langle P, O, I, G, C \rangle$ be the planning task under consideration, and let $\mathcal{L}$ be a landmark generation method that is invariant under structural symmetry. Let h_1 be the variant of the landmark heuristic with uniform cost partitioning that treats action landmarks specially, and let h_2 be the one that does not.

First, we define the set of action landmarks that are treated specially, which we denote by $O_{\mathrm{AL}}(s)$:

$$\begin{aligned} O_{\mathrm{AL}}(s) &= \{o \in O \mid \{o\} \in \mathcal{L}(s)\} && \text{if } h = h_1 \\ O_{\mathrm{AL}}(s) &= \emptyset && \text{if } h = h_2. \end{aligned}$$

Denote by $\mathcal{L}'(s)$ all landmarks that participate in the cost partitioning process:[3]

$$\mathcal{L}'(s) = \{L \in \mathcal{L}(s) \mid L \cap O_{\mathrm{AL}}(s) = \emptyset\}$$

[3]Informally, in h_2, these are simply all landmarks. In h_1, these are all landmarks which do not contain an operator that defines an action landmark. The stated definition in terms of $O_{\mathrm{AL}}(s)$ works in either case.

Denote by $O_\mathrm{L}(s)$ all operators that participate in at least one landmark in $\mathcal{L}'(s)$:

$$O_\mathrm{L}(s) = \bigcup_{L \in \mathcal{L}'(s)} L$$

We assign a "per-landmark cost share" to each such operator in state s, which is the operator cost divided by the number of landmarks for which it participates in the cost partitioning:

$$\mathit{share}(o, s) = \frac{C(o)}{|\{L \in \mathcal{L}'(s) \mid o \in L\}|} \quad \text{for all } o \in O_\mathrm{L}(s).$$

The overall heuristic value is then the sum of two parts:

$$h(s) = h_\mathrm{AL}(s) + h_\mathrm{L}(s),$$

where $h_\mathrm{AL}(s)$ is the cost contributed by action landmarks that are treated specially:

$$h_\mathrm{AL}(s) = \sum_{o \in O_\mathrm{AL}(s)} C(o)$$

and $h_\mathrm{L}(s)$ is the cost contributed by other landmarks:

$$h_\mathrm{L}(s) = \sum_{L \in \mathcal{L}'(s)} \min_{o \in L} \mathit{share}(o, s)$$

We now show that h is invariant under structural symmetry, so let σ be a structural symmetry. The proof proceeds in stages, showing in sequence:

$$O_\mathrm{AL}(\sigma(s)) = \sigma(O_\mathrm{AL}(s)) \quad (5)$$
$$\mathcal{L}'(\sigma(s)) = \sigma(\mathcal{L}'(s)) \quad (6)$$
$$O_\mathrm{L}(\sigma(s)) = \sigma(O_\mathrm{L}(s)) \quad (7)$$
$$\mathit{share}(\sigma(o), \sigma(s)) = \mathit{share}(o, s) \quad (8)$$
$$h_\mathrm{AL}(\sigma(s)) = h_\mathrm{AL}(s) \quad (9)$$
$$h_\mathrm{L}(\sigma(s)) = h_\mathrm{L}(s) \quad (10)$$
$$h(\sigma(s)) = h(s) \quad (11)$$

The last of these results then concludes the proof.

Regarding Eq. 5: if $h = h_2$, this is trivial, as both sides are equal to $\emptyset$. If $h = h_1$, we get

$$\begin{aligned}
O_\mathrm{AL}(\sigma(s)) &= \{o \in O \mid \{o\} \in \mathcal{L}(\sigma(s))\} \\
&= \{o \in O \mid \{o\} \in \sigma(\mathcal{L}(s))\} \\
&= \{\sigma(o') \mid o' \in O, \{\sigma(o')\} \in \sigma(\mathcal{L}(s))\} \\
&= \{\sigma(o') \mid o' \in O, \{o'\} \in \mathcal{L}(s)\} \\
&= \sigma(\{o' \in O \mid \{o'\} \in \mathcal{L}(s)\}) \\
&= \sigma(O_\mathrm{AL}(s)).
\end{aligned}$$

Regarding Eq. 6:

$$\begin{aligned}
\mathcal{L}'(\sigma(s)) &= \{L \in \mathcal{L}(\sigma(s)) \mid L \cap O_\mathrm{AL}(\sigma(s)) = \emptyset\} \\
&= \{L \in \sigma(\mathcal{L}(s)) \mid L \cap \sigma(O_\mathrm{AL}(s)) = \emptyset\} \\
&= \{\sigma(L') \mid L' \in \mathcal{L}(s), \sigma(L') \cap \sigma(O_\mathrm{AL}(s)) = \emptyset\} \\
&= \{\sigma(L') \mid L' \in \mathcal{L}(s), \sigma(L' \cap O_\mathrm{AL}(s)) = \emptyset\} \\
&= \{\sigma(L') \mid L' \in \mathcal{L}(s), L' \cap O_\mathrm{AL}(s) = \emptyset\} \\
&= \sigma(\{L' \in \mathcal{L}(s) \mid L' \cap O_\mathrm{AL}(s) = \emptyset\}) \\
&= \sigma(\mathcal{L}'(s))
\end{aligned}$$

Regarding Eq. 7:

$$\begin{aligned}
O_\mathrm{L}(\sigma(s)) &= \bigcup_{L \in \mathcal{L}'(\sigma(s))} L \\
&= \bigcup_{L \in \sigma(\mathcal{L}'(s))} L \\
&= \bigcup_{L' \in \mathcal{L}'(s)} \sigma(L') \\
&= \sigma\Bigg(\bigcup_{L' \in \mathcal{L}'(s)} L'\Bigg) \\
&= \sigma(O_\mathrm{L}(s))
\end{aligned}$$

Regarding Eq. 8, for all $o \in O_\mathrm{L}(s)$ (or equivalently: $\sigma(o) \in \sigma(O_\mathrm{L}(s)) = O_\mathrm{L}(\sigma(s))$):

$$\begin{aligned}
\mathit{share}(\sigma(o), \sigma(s)) &= \frac{C(\sigma(o))}{|\{L \in \mathcal{L}'(\sigma(s)) \mid \sigma(o) \in L\}|} \\
&= \frac{C(o)}{|\{L \in \sigma(\mathcal{L}'(s)) \mid \sigma(o) \in L\}|} \\
&= \frac{C(o)}{|\{\sigma(L') \in \sigma(\mathcal{L}'(s)) \mid \sigma(o) \in \sigma(L')\}|} \\
&= \frac{C(o)}{|\{L' \in \mathcal{L}'(s) \mid o \in L'\}|} \\
&= \mathit{share}(o, s)
\end{aligned}$$

Regarding Eq. 9:

$$\begin{aligned}
h_\mathrm{AL}(\sigma(s)) &= \sum_{o \in O_\mathrm{AL}(\sigma(s))} C(o) \\
&= \sum_{o \in \sigma(O_\mathrm{AL}(s))} C(o) \\
&= \sum_{o' \in O_\mathrm{AL}(s)} C(\sigma(o')) \\
&= \sum_{o' \in O_\mathrm{AL}(s)} C(o') \\
&= h_\mathrm{AL}(s)
\end{aligned}$$

Regarding Eq. 10:

$$\begin{aligned}
h_\mathrm{L}(\sigma(s)) &= \sum_{L \in \mathcal{L}'(\sigma(s))} \min_{o \in L} \mathit{share}(o, \sigma(s)) \\
&= \sum_{L \in \sigma(\mathcal{L}'(s))} \min_{o \in L} \mathit{share}(o, \sigma(s)) \\
&= \sum_{L' \in \mathcal{L}'(s)} \min_{o \in \sigma(L')} \mathit{share}(o, \sigma(s)) \\
&= \sum_{L' \in \mathcal{L}'(s)} \min_{o' \in L'} \mathit{share}(\sigma(o'), \sigma(s)) \\
&= \sum_{L' \in \mathcal{L}'(s)} \min_{o' \in L'} \mathit{share}(o', s) \\
&= h_\mathrm{L}(s)
\end{aligned}$$

Regarding Eq. 11:

$$\begin{aligned} h(\sigma(s)) &= h_{\mathrm{AL}}(\sigma(s)) + h_{\mathrm{L}}(\sigma(s)) \\ &= h_{\mathrm{AL}}(s) + h_{\mathrm{L}}(s) \\ &= h(s) \end{aligned}$$

□

References

Bonet, B., and Helmert, M. 2010. Strengthening landmark heuristics via hitting sets. In *Proc. ECAI 2010*, 329–334.

Karpas, E., and Domshlak, C. 2009. Cost-optimal planning with landmarks. In *Proc. IJCAI 2009*, 1728–1733.

Pochter, N.; Zohar, A.; and Rosenschein, J. S. 2011. Exploiting problem symmetries in state-based planners. In *Proc. AAAI 2011*, 1004–1009.

Richter, S., and Westphal, M. 2010. The LAMA planner: Guiding cost-based anytime planning with landmarks. *JAIR* 39:127–177.

Richter, S.; Helmert, M.; and Westphal, M. 2008. Landmarks revisited. In *Proc. AAAI 2008*, 975–982.