

Simplifying Planning Tasks with Fact-Level Relevance Analysis

Cameron Allen¹, Anita de Mello Koch², Harsha Kokel³, George Konidaris², Michael Katz³

¹UC Berkeley ²Brown University ³IBM
camallen@berkeley.edu, michaelkatz1@ibm.com

Abstract

Planning problems described in formal languages often contain information that is irrelevant to reaching the specified goal, either by design (to account for a variety of possible goals) or as a result of partially automated construction. Consequently, modern planners prune (some) extraneous information by performing safe relevance analysis of state variables prior to search, improving search performance. We show that much more aggressive pruning is possible by reasoning about relevance not on the level of variables, but rather variable *values*, i.e. *facts*. Our approach iteratively identifies relevant facts whenever they appear in the task goal or in the precondition of a relevant action, and identifies relevant actions whenever they achieve a relevant fact (rather than when they simply modify a relevant variable). This fact-level analysis already reduces task size relative to the more commonly-used variable-level analysis, but it also offers several opportunities for further algorithmic improvements. We formally prove that fact-level relevance analysis, and each subsequent algorithmic improvement we introduce, preserves all shortest optimal plans, without introducing new ones. We then show empirically that our approach significantly reduces task size and consequently planning time across a wide range of planning tasks.

Introduction

A planning agent must either learn or be supplied with a world model rich enough to express any task the agent might encounter during its operational lifetime. Consider a planning agent within the game of Minecraft. Mastering the environment requires hundreds of diverse skills for accomplishing behaviors like navigation, resource gathering, crafting various items, fighting, managing hunger, and decorating.

As the agent’s world model grows in complexity to account for new tasks, so too does the cost of planning with it. If the agent is to efficiently accomplish a particular task, like obtaining a stone pickaxe, it cannot afford to consider all the ways it might bake a cake along the way (see Figure 2, left). Such actions are out of scope for that task and ought to be pruned. However, for another task, where the goal is explicitly to bake a cake, the food-related actions might become essential whereas the pickaxe-related actions become irrelevant. Deciding for a given task which parts of the world model the agent must consider requires either manual human

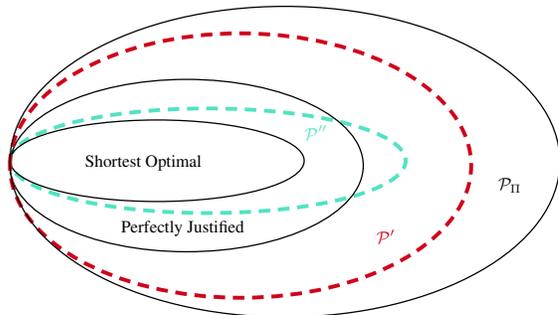


Figure 1: The space of all plans \mathcal{P}_Π for a given task. We introduce algorithms for fact-based task pruning, resulting in a subset of plans \mathcal{P}' that preserves perfectly justified plans. We also introduce algorithms that prune more aggressively (resulting in \mathcal{P}'') while preserving shortest optimal plans. None of the algorithms introduces any new plans outside \mathcal{P}_Π .

effort or automated analysis to remove extraneous information prior to searching for a plan.

Existing approaches that automatically prune irrelevant information from the agent’s world model try to strike a balance between pruning overhead and the search effort it saves (Nebel, Dimopoulos, and Koehler 1997; Boutilier, Brafman, and Geib 1998; Brafman 2001; Hoffmann and Nebel 2001b). Most modern planners base their analysis on the Fast Downward planning system (Helmert 2006), which performs causal graph analysis of state variables to safely and efficiently approximate which variables are goal-relevant. However, this process is overly conservative in the information it prunes.

In this paper, we build on these approaches and show that additional relevance analysis can significantly reduce search effort without adding substantial overhead. We argue that performing relevance analysis at the level of facts, as opposed to variables, offers several advantages. First, fewer actions will be marked as causally connected to the goal, since action preconditions and effects will have fewer facts in common than they will variables. Second, certain causal connections can be ignored if a precondition fact appears in the initial state and is unthreatened by any relevant action (Fig. 2, top-right). Third, we can form a reachability-

or

(iii) $eff(a_i)[v] = pre(a_j)[v]$ for some $j > i$ and $s_k[v] = eff(a_i)[v]$ for all $i < k \leq j$.

An action in a plan is redundant, or *unjustified*, if it can be removed from the plan without invalidating the plan. A plan where each action is justified is called a justified plan. A plan without a redundant subsequence is called perfectly justified plan (Fink and Yang 1992; Nebel, Dimopoulos, and Koehler 1997; Salerno, Fuentetaja, and Seipp 2023). We define justified plans and plan reduction as follows.

Definition 2 (Justified Plan) A plan π for the planning task Π is a justified plan if and only if all its actions are justified.

Definition 3 (Plan Reduction) Let π be a plan for a planning task Π . A subsequence ρ of π is a plan reduction of π if and only if ρ is also a plan for Π .

Definition 4 (Perfectly Justified Plan) A plan π for the planning task Π is a perfectly justified plan if and only if there is no plan reduction of π .

All shortest optimal plans are perfectly justified plans, and all perfectly justified plans are justified plans. The reverse direction does not necessarily hold for either relationship.

Preserving Action Applicability

Our approach starts by identifying a set of actions that the pruned task ought to retain, then it removes all *other* actions, along with any non-essential facts. This section introduces Algorithm 1 (PRUNE), which describes, for an arbitrary subset of actions, how to retain a sufficient set of facts for computing action applicability. Later sections will discuss how to compute the set of relevant actions.

Theorem 1 Given a task $\Pi = \langle \mathcal{V}, \mathcal{A}, s_0, G, cost \rangle$, a set of actions $\bar{\mathcal{A}} \subseteq \mathcal{A}$, and an action sequence $\pi = \langle a_0, \dots, a_{n-1} \rangle$ that only uses actions in $\bar{\mathcal{A}}$, let $\Pi' = \text{PRUNE}(\Pi, \bar{\mathcal{A}})$. Let π' be the corresponding action sequence in Π' with the same names and preconditions. Then π' is a plan for Π' if and only if π is a plan for Π .

Proof:

Let F be computed as in lines 1-4 of Algorithm 1. For each action $a \in \bar{\mathcal{A}}$, denote its corresponding action in \mathcal{A}' as a' , and observe that $eff(a') = F \cap eff(a)$ and that $pre(a') = pre(a)$.

The initial state is preserved fully, so $s_0 = s'_0$ and $pre(a_0) \subseteq s_0 \iff pre(a'_0) \subseteq s'_0$, hence Π' preserves the applicability of the first action.

Now suppose that actions 0 through $i \leq k$ are applicable in both tasks: $a_{0:i}$ and $a'_{0:i}$ are both applicable in s_0 . Let $s'_{i+1} = s_0[a'_{0:i}]$ and $s_{i+1} = s_0[a_{0:i}]$. Since $eff(a'_i) = F \cap eff(a_i)$, we have $s'_{i+1} = s_{i+1} \cap F$.

Case 1: If π is a plan, we have $pre(a_{i+1}) \subseteq s_{i+1} = s_i[a_i]$, which also holds if we project states and preconditions onto the facts in F : $pre(a_{i+1}) \cap F \subseteq s_{i+1} \cap F = s_i[a_i] \cap F$. And since $pre(a_{i+1}) \subseteq F$, it follows that $pre(a'_{i+1}) = pre(a_{i+1}) \cap F \subseteq s_i[a_i] \cap F = (s_i \cap F)[a'_i] = s'_i[a'_i] = s'_{i+1}$.

Algorithm 1: PRUNE

INPUT: $\Pi = \langle \mathcal{V}, \mathcal{A}, s_0, G, cost \rangle$ // Planning task
 $\bar{\mathcal{A}}$ // Set of actions to preserve

OUTPUT: Π' // Pruned planning task

1: $F = G \cup \bigcup_{a \in \bar{\mathcal{A}}} pre(a)$ // Precondition and goal facts

2: $\bar{\mathcal{V}} = \mathcal{V}(F)$ // Relevant variables

3: $F = F \cup \bigcup_{a \in \bar{\mathcal{A}}} eff(a)[\bar{\mathcal{V}}]$ // Add relevant effect facts

4: $F = F \cup s_0$ // Add initial state facts

5: $\mathcal{A}' = \emptyset$ // Initialize new action set

6: **for** $a_i \in \bar{\mathcal{A}}$ **do**

7: $a'_i = \langle pre(a_i), eff(a_i) \cap F \rangle$

8: $cost'(a'_i) = cost(a_i)$

9: $\mathcal{A}' = \mathcal{A}' \cup a'_i$ // Add to new action set

10: **end for**

11: **return** $\Pi' = \langle \mathcal{V}, \mathcal{A}', s_0, G, cost' \rangle$

Case 2: If π' is a plan, we have $pre(a'_{i+1}) \subseteq s'_{i+1} = s'_i[a'_i]$. Therefore, $pre(a_{i+1}) = pre(a_{i+1}) \cap F = pre(a'_{i+1}) \cap F \subseteq s'_{i+1} \cap F = s_{i+1} \cap F \subseteq s_{i+1}$.

In both cases, the next actions, a_{i+1} and a'_{i+1} , are applicable in states s_{i+1} and s'_{i+1} , as are the action sequences $a_{0:i+1}$ and $a'_{0:i+1}$ in s_0 , respectively. By induction on k , both plans are applicable and lead to $s_n = s_0[\pi]$ and $s'_n = s_0[\pi']$, respectively.

Case 1: If π is a plan, $G \subseteq s_0[\pi]$, and we have $G = G \cap F \subseteq s_0[\pi] \cap F = s_0[\pi']$, so π' is a plan for Π' .

Case 2: If π' is a plan, we have $G = G \cap F \subseteq s'_n \cap F = s_n \cap F \subseteq s_n$, so π is a plan for Π . \square

Corollary 1 Each plan in the pruned task Π' corresponds to a plan in the original task Π .

Proof:

This follows directly from Theorem 1, since $\bar{\mathcal{A}} \subseteq \mathcal{A}$. \square

After pruning, two optimizations are possible. First, we can remove variables that have only one value. Second, if we are only concerned with preserving optimal or perfectly justified plans, we can also prune actions with empty effects. We omit these optimizations to make the proofs simpler, but neither changes the underlying transition system.

Theorem 1 guarantees that (1) any plans we can express with the identified subset of actions will remain plans (and have the same cost) in the pruned task and (2) any plans we identify in the pruned task are also plans in the original task. The former ensures solution quality; the latter ensures solution validity. All that remains is to choose an appropriate subset of the actions.

From Variable-Level to Fact-Level Analysis

We start from a method that is equivalent to the removal of *irrelevant* variables (Helmert 2006).

In RELEVANCE-V (Alg. 2), we prune the task by performing (relaxed) backward reachability analysis from the

Algorithm 2: RELEVANCE-V (Variable Relevance)

INPUT: $\Pi = \langle \mathcal{V}, \mathcal{A}, s_0, G, cost \rangle$ // Planning task
OUTPUT: $\bar{\mathcal{A}}$ // Relevant actions

- 1: $\mathcal{V}^{(0)} = \mathcal{V}(G)$ // Initialize relevant variables with goal
- 2: $\mathcal{A}^{(0)} = \emptyset$ // Initialize relevant actions with empty set
- 3: **for** $i = 1, 2, \dots$ **do**
- 4: $\mathcal{A}^{(i)} = \mathcal{A}^{(i-1)} \cup \{a \in \mathcal{A} \mid \mathcal{V}(eff(a)) \cap \mathcal{V}^{(i-1)} \neq \emptyset\}$
- 5: $\mathcal{V}^{(i)} = \mathcal{V}^{(i-1)} \cup \bigcup_{a \in \mathcal{A}^{(i)}} \mathcal{V}(pre(a))$
- 6: **if** $\mathcal{V}^{(i)} = \mathcal{V}^{(i-1)}$ and $\mathcal{A}^{(i)} = \mathcal{A}^{(i-1)}$ **then break**
- 7: **end for**
- 8: **return** $\bar{\mathcal{A}} = \mathcal{A}^{(i)}$

goal (Fig. 2, top right). We start by marking the goal variables as relevant. We iteratively find actions that affect a relevant variable and mark their precondition variables as relevant. We stop when there is no change to the set of marked variables. The resulting set of actions can then be passed to PRUNE (Alg. 1) to simplify the task (see previous section).

Our first observation is that instead of reasoning on the level of variables, we can reason on the level of *facts*, i.e. variable values. This leads to RELEVANCE-F (Algorithm 3), wherein we start by marking the goal facts as relevant and iteratively add actions whose effects contain relevant facts.

Theorem 2 *Given a task Π , let $\bar{\mathcal{A}}_V = \text{RELEVANCE-V}(\Pi)$ and $\bar{\mathcal{A}}_F = \text{RELEVANCE-F}(\Pi)$. $\bar{\mathcal{A}}_V$ and $\bar{\mathcal{A}}_F$ each contain all actions from all justified plans for Π .*

Proof:

We first consider RELEVANCE-F (Algorithm 3).

Recall that a justified plan is one where each action establishes either a goal fact or a precondition of a subsequent action. Let π be any such plan, and let N be the length of π . Suppose at least one of the actions in π is not in $\bar{\mathcal{A}}_F$, and, without loss of generality, let a_j be the last such action.

All actions a_k , for $k > j$, are in $\bar{\mathcal{A}}_F$. So, in the final iteration i of the algorithm, $F^{(i-1)}$ contains their preconditions: $G \cup \bigcup_{k=j+1}^N pre(a_k) \subseteq F^{(i)} = F^{(i-1)}$, and $eff(a_j) \cap F^{(i-1)} = \emptyset$ (since $a_j \notin \bar{\mathcal{A}}_F$). But since a_j is justified for π , $eff(a_j) \cap (G \cup \bigcup_{k=j+1}^N pre(a_k)) \neq \emptyset$, which is a contradiction. Therefore, every action in π is in $\bar{\mathcal{A}}_F$, and this holds for any justified plan π .

Now note that for any set of facts F , we have that $(eff(a_i) \cap F \neq \emptyset) \implies \mathcal{V}(eff(a_i)) \cap \mathcal{V}(F) \neq \emptyset$. Whenever Algorithm 3 marks a fact relevant in $F^{(i)}$, Algorithm 2 marks *all* facts for the same variable relevant in $\mathcal{V}^{(i)}$. Thus, $\bar{\mathcal{A}}_F \subseteq \bar{\mathcal{A}}_V$, and RELEVANCE-V (Algorithm 2) preserves all actions on all justified plans. \square

Corollary 2 $\bar{\mathcal{A}}_V$ and $\bar{\mathcal{A}}_F$ each contain all actions from all perfectly justified plans for Π .

Corollary 3 $\bar{\mathcal{A}}_V$ and $\bar{\mathcal{A}}_F$ each contain all actions from all shortest optimal plans for Π .

Algorithm 3: RELEVANCE-F (Fact Relevance)

INPUT: $\Pi = \langle \mathcal{V}, \mathcal{A}, s_0, G, cost \rangle$ // Planning task
OUTPUT: $\bar{\mathcal{A}}$ // Relevant actions

- 1: $F^{(0)} = G$ // Initialize relevant facts with goal facts
- 2: $\mathcal{A}^{(0)} = \emptyset$ // Initialize relevant actions with empty set
- 3: **for** $i = 1, 2, \dots$ **do**
- 4: $\mathcal{A}^{(i)} = \mathcal{A}^{(i-1)} \cup \{a \in \mathcal{A} \mid eff(a) \cap F^{(i-1)} \neq \emptyset\}$
- 5: $F^{(i)} = F^{(i-1)} \cup \bigcup_{a \in \mathcal{A}^{(i)}} pre(a)$
- 6: **if** $F^{(i)} = F^{(i-1)}$ and $\mathcal{A}^{(i)} = \mathcal{A}^{(i-1)}$ **then break**
- 7: **end for**
- 8: **return** $\bar{\mathcal{A}} = \mathcal{A}^{(i)}$

The Power of Fact-Level Analysis

Our next observation is that the stopping criteria can be refined. Since we already track which actions are relevant, we can check whether any relevant action threatens a fact that appears in the initial state. An initial-state fact that is unthreatened by any relevant action is said to be *causally linked* to the initial state.¹ We proceed as before, but now we filter out these (currently) causally linked facts prior to computing relevant actions. Next, we add the preconditions of these relevant actions to our list of *unfiltered* relevant facts.² Then, we filter again (based on the new set of relevant actions), and repeat until we reach a fixed point. We can thus stop the iteration early whenever the only reason to mark further actions relevant would be due to them achieving causally linked facts (see Algorithm 4).

We show an example of this in Figure 2 (middle-right). Suppose the task has goal `has-shears` and initial state `has-stone-pickaxe`. Since the only precondition fact of `mine-iron` is `has-stone-pickaxe`, and neither `mine-iron` nor `craft-shears` threatens that fact, we need not mark as relevant `craft-stone-pickaxe`, nor any of its ancestors.

Theorem 3 *Given a task Π , $\bar{\mathcal{A}}_C = \text{RELEVANCE-FC}(\Pi)$ contains all actions from all perfectly justified plans for Π .*

Proof:

Suppose $\pi = \langle a_1, a_2, \dots, a_n \rangle$ is a perfectly justified plan for task Π . Suppose at least one of the actions in π is not in $\bar{\mathcal{A}}_C$, and, without loss of generality, let a_j be the last such action.

All actions a_k , for $k > j$, are in $\bar{\mathcal{A}}_C$. So, in the final iteration i of the algorithm, $F^{(i-1)}$ contains their preconditions: $G \cup \bigcup_{k=j+1}^N pre(a_k) \subseteq F^{(i)} = F^{(i-1)}$. Additionally, $eff(a_j) \cap F^{(i-1)} = \emptyset$ (since $a_j \notin \bar{\mathcal{A}}_C$).

There are two possible explanations for this: either $eff(a_j) \cap F^{(i-1)} = \emptyset$, or $eff(a_j) \cap F^{(i-1)} \neq \emptyset$ and

¹This is inspired by McAllester and Rosenblitt (1991), which describes causal links between actions.

²Storing the unfiltered relevant facts ensures that we still consider facts that were previously causally linked but that may have become threatened by a relevant action in a later iteration.

$(\text{eff}(a_j) \cap F^{(i-1)}) \subseteq C^{(i)}$. The former means no fact in a_j 's effect appears in the goal or any subsequent precondition. This is a contradiction as a_j is part of a perfectly justified plan.

The latter means any fact f in a_j 's effect that does appear in the goal or a subsequent precondition was already established by the initial state, and its value is unthreatened by any action in $\mathcal{A}^{(i)}$. This follows from the definition of $C^{(i)}$, but note that f may still be threatened by an action *not* in $\mathcal{A}^{(i)}$ (and a_j restores it). Therefore, if we delete a_j from the action sequence, the only way it could invalidate the plan is if a previous action $a'_j \notin \mathcal{A}^{(i)}$ in the plan threatens a causally linked fact.

Repeat the above analysis for the penultimate action not in $\mathcal{A}^{(i)}$, which now serves the same role as a_j . Again, it can be deleted. As before, the only way this could invalidate the plan is if a previous action $a''_j \notin \mathcal{A}^{(i)}$ in the plan threatens a causally linked fact.

Delete all the remaining actions not in $\mathcal{A}^{(i)}$, one by one, until none remain. At this point, the remaining sequence must be a plan, since no remaining actions threaten the causally linked facts, and none of the removed actions achieved any subsequent precondition or goal facts. (This is true even if the *initial* action in the plan was deleted, thus the sequence remains applicable in the initial state.) It is also a plan reduction, since it is shorter than π , which contradicts the fact that π was perfectly justified. \square

Corollary 4 $\bar{\mathcal{A}}_C$ contains all actions from all shortest optimal plans for Π .

Action Merging

Reasoning at the fact-level also allows us to perform the relevance analysis in a different labeled transition system.

Let $\Pi = \langle \mathcal{V}, \mathcal{A}, s_0, G, \text{cost} \rangle$ be a planning task, $\bar{\mathcal{A}} \subseteq \mathcal{A}$ be a set of relevant actions, and $\bar{\mathcal{V}} \subseteq \mathcal{V}$ be a set of relevant variables. Partition the actions in $\bar{\mathcal{A}}$ by cost and effects on variables in $\bar{\mathcal{V}}$. For each partition $\bar{\mathcal{A}}_k$, merge the actions to form a new action \bar{a}_k , as follows. For the precondition, compute the disjunction of the individual preconditions (DNF) and simplify the result: $\text{Pre}(\bar{a}_k) = \text{SIMPLIFYDNF}(\bigvee_{a \in \bar{\mathcal{A}}_k} \text{pre}(a))$. The aim of this process is to produce a DNF with fewer facts, though the resulting DNF may still include multiple disjunctions, which we then represent by separate actions. For the cost and effect, take the (common) cost and effect on variables in $\bar{\mathcal{V}}$ (guaranteed by the partition). Let \mathcal{A}' be the set of these merged actions, and define $\Pi' = \langle \mathcal{V}, \mathcal{A}', s_0, G, \text{cost} \rangle$. The labeled transition system of Π' can be obtained from the labeled transition system of Π via label reduction (Helmert, Haslum, and Hoffmann 2007; Sievers, Wehrle, and Helmert 2014).

Provided that only variables in $\bar{\mathcal{V}}$ appear in the merged preconditions, the labeled transition system only differs in its labels and its effects on variables outside of $\bar{\mathcal{V}}$. The backward and forward reachability of facts about variables in $\bar{\mathcal{V}}$ is the same in both. When $\bar{\mathcal{V}}$ is taken to be the set of (currently) relevant variables, this means we can alternatively perform

Algorithm 4: RELEVANCE-FC (Causal Links)

INPUT: $\Pi = \langle \mathcal{V}, \mathcal{A}, s_0, G, \text{cost} \rangle$ // Planning task

OUTPUT: $\bar{\mathcal{A}}$ // Relevant actions

```

1:  $F^{(0)} = G$  // Initialize relevant facts with goal facts
2:  $\mathcal{A}^{(0)} = \emptyset$  // Initialize relevant actions with empty set
3: for  $i = 1, 2, \dots$  do
4:    $T = (\bigcup_{a \in \mathcal{A}^{(i-1)}} \text{eff}(a)) \setminus s_0$  // Threatening facts
5:    $C^{(i)} = s_0[\mathcal{V} \setminus \mathcal{V}(T)]$  // Causally linked facts
6:    $\bar{F}^{(i)} = F^{(i-1)} \setminus C^{(i)}$  // Filter out causally linked facts
7:    $\mathcal{A}^{(i)} = \mathcal{A}^{(i-1)} \cup \{a \in \mathcal{A} \mid \text{eff}(a) \cap \bar{F}^{(i)} \neq \emptyset\}$ 
8:    $F^{(i)} = F^{(i-1)} \cup \bigcup_{a \in \mathcal{A}^{(i)}} \text{pre}(a)$ 
9:   if  $F^{(i)} = F^{(i-1)}$  and  $\mathcal{A}^{(i)} = \mathcal{A}^{(i-1)}$  then break
10: end for
11: return  $\bar{\mathcal{A}} = \mathcal{A}^{(i)}$ 

```

Algorithm 5: RELEVANCE-FCM (Casual Links & Merging)

INPUT: $\Pi = \langle \mathcal{V}, \mathcal{A}, s_0, G, \text{cost} \rangle$ // Planning task

OUTPUT: $\bar{\mathcal{A}}$ // Relevant actions

```

1:  $F^{(0)} = G$  // Facts from goal state
2:  $\mathcal{A}^{(0)} = \emptyset$  // Initialize relevant actions
3: for  $i = 1, 2, \dots$  do
4:    $T = (\bigcup_{a \in \mathcal{A}^{(i-1)}} \text{eff}(a)) \setminus s_0$  // Threatening facts
5:    $C^{(i)} = s_0[\mathcal{V} \setminus \mathcal{V}(T)]$  // Causally linked facts
6:    $\bar{F}^{(i)} = F^{(i-1)} \setminus C^{(i)}$  // Filter out causally linked facts
7:    $\mathcal{A}^{(i)} = \mathcal{A}^{(i-1)} \cup \{a \in \mathcal{A} \mid \text{eff}(a) \cap \bar{F}^{(i)} \neq \emptyset\}$ 
8:    $\mathcal{A}' = \text{REDUCE}(\Pi, \mathcal{A}^{(i)}, \mathcal{V}(F^{(i-1)}))$ 
9:    $F^{(i)} = F^{(i-1)} \cup \bigcup_{a \in \mathcal{A}'} \text{pre}(a)$ 
10:  if  $F^{(i)} = F^{(i-1)}$  and  $\mathcal{A}^{(i)} = \mathcal{A}^{(i-1)}$  then break
11: end for
12: return  $\bar{\mathcal{A}} = \mathcal{A}^{(i)}$ 

```

Algorithm 6: REDUCE

INPUT: $\Pi = \langle \mathcal{V}, \mathcal{A}, s_0, G, \text{cost} \rangle$ // Planning task

$\bar{\mathcal{A}} \subseteq \mathcal{A}$ // Subset of actions

$\bar{\mathcal{V}} \subseteq \mathcal{V}$ // Subset of variables

OUTPUT: $\bar{\mathcal{A}}'$ // Merged actions

```

1:  $\mathcal{P} \leftarrow$  Partition  $\bar{\mathcal{A}}$  by cost and effects on  $\bar{\mathcal{V}}$ .
2: Initialize  $\bar{\mathcal{A}}' = \emptyset$ .
3: for all  $\langle \bar{\mathcal{A}}_k, \text{eff}_k, \text{cost}_k \rangle \in \mathcal{P}$  do
4:    $\text{Pre}(\bar{a}'_k) = \text{SIMPLIFYDNF}(\bigvee_{a \in \bar{\mathcal{A}}_k} \text{pre}(a))$ 
5:    $\bar{\mathcal{A}}'_k = \{\langle \text{pre}(\bar{a}'_k), \text{eff}_k, \text{cost}_k \rangle \mid \text{pre}(\bar{a}'_k) \in \text{Pre}(\bar{a}'_k)\}$ 
6:    $\bar{\mathcal{A}}' = \bar{\mathcal{A}}' \cup \bar{\mathcal{A}}'_k$ 
7: end for
8: return  $\bar{\mathcal{A}}'$ 

```

relevance analysis in the label reduced transition system. In Algorithm 5, the relevance analysis considers the preconditions of \mathcal{A}' , rather than $\mathcal{A}^{(i)}$, when marking facts relevant. While the relevant variables and values in these labeled transition systems are the same, their sets of plans might differ and removing irrelevant facts may result in removing a perfectly justified plan.

For example, in the Minecraft domain (Fig. 2, bottom-right), the `hunt` action has precondition \neg hungry, whereas `gather` has no precondition and achieves the same (relevant) fact,³ `tribe-has-food`, for the same cost. Merging `hunt` and `gather` will simplify the preconditions to those of `gather` (i.e. empty), and thus `eat-cake` (which achieves \neg hungry) will not be marked relevant. If the task has goal `tribe-has-food` and initial state `has-cake \wedge hungry`, then `[eat-cake, hunt]` is a perfectly justified plan, and `[gather]` is a shortest optimal plan. RELEVANCE-FCM (Alg. 5) removes `eat-cake` and the perfectly justified plan, but not the optimal plan.

Theorem 4 *Given a task Π , $\bar{\mathcal{A}}_M = \text{RELEVANCE-FCM}(\Pi)$ contains all actions from all shortest optimal plans for Π .*

Proof:

Suppose $\pi = \langle a_1, a_2, \dots, a_n \rangle$ is a shortest optimal plan for task Π , with corresponding state sequence $\langle s_0, s_1, \dots, s_n \rangle$. Suppose at least one of the actions in π is not in $\bar{\mathcal{A}}_M$, and, without loss of generality, let a_j be the last such action.

All actions a_k , for $k > j$, are in $\bar{\mathcal{A}}_M$. So, in the final iteration i of the algorithm, $F^{(i-1)}$ contains their (possibly reduced and simplified) preconditions. If no partition contains more than one action, or if no partition's preconditions simplify, then $\bigcup_{a \in \mathcal{A}'} \text{pre}(a) = \bigcup_{a \in \mathcal{A}^{(i)}} \text{pre}(a)$ and Theorem 3 applies. Otherwise we must consider what happens during the label reduction and precondition simplification from $\mathcal{A}^{(i)}$ to \mathcal{A}' .

Suppose the partition computed by REDUCE contains an action set $\bar{\mathcal{A}}_p \subseteq \mathcal{A}^{(i)}$ with at least two actions whose preconditions simplify to form a merged action set $\bar{\mathcal{A}}'_p$. Specifically, if $P = \bigcup_{a \in \bar{\mathcal{A}}_p} \text{pre}(a)$ and $P' = \bigcup_{a \in \bar{\mathcal{A}}'_p} \text{pre}(a)$, then $P' \subset P$.

We are interested in the case where some fact $f \in \text{eff}(a_j)$ is not added to $F^{(i)}$ because of the merging process—that is, $f \in P \setminus P'$. If no such fact exists for any partition, then we can delete a_j from the plan using the same reasoning as in Theorem 3. Otherwise, let $v = \mathcal{V}(f)$, and notice that since $f \notin P'$, all values for v must appear in P , i.e. $\{\langle v, d \rangle \mid d \in \mathcal{D}(v)\} \subseteq P$. Thus, a_j achieves fact f in a precondition of some subsequent action a_f , but for each other value of v , there exists another action a'_f (at least one, but possibly more), whose cost and effect are the same and whose precondition is applicable.

We can delete a_j from the plan and replace a_f with the a'_f that is applicable in state s_{j-1} , i.e. $\text{pre}(a'_f)[v] =$

³`gather` also achieves `hungry`, but it is not task-relevant.

Algorithm 7: REACHABILITY

INPUT: $\Pi = \langle \mathcal{V}, \mathcal{A}, s_0, G, \text{cost} \rangle$ // Planning task
 OUTPUT: $\bar{\mathcal{A}}$ // (Relaxed) reachable actions

```

1:  $F^{(0)} = s_0$  // Initialize reachable facts with initial facts
2:  $\mathcal{A}^{(0)} = \emptyset$  // Initialize reachable actions with empty set
3: for  $i = 1, 2, \dots$  do
4:    $\mathcal{A}^{(i)} = \mathcal{A}^{(i-1)} \cup \{a \in \mathcal{A} \mid \text{pre}(a) \subseteq F^{(i-1)}\}$ 
5:    $F^{(i)} = F^{(i-1)} \cup \bigcup_{a \in \mathcal{A}^{(i)}} \text{eff}(a)$ 
6:   if  $F^{(i)} = F^{(i-1)}$  and  $\mathcal{A}^{(i)} = \mathcal{A}^{(i-1)}$  then break
7: end for
8: return  $\bar{\mathcal{A}} = \mathcal{A}^{(i)}$ 

```

Algorithm 8: RELEVANCEITERATION

INPUT: Π // Planning task
 OUTPUT: Π' // Pruned task

```

1: loop
2:    $\Pi' = \text{PRUNE}(\Pi, \text{RELEVANCE}(\Pi))$  // e.g. Algs. 2–5
3:    $\Pi' = \text{PRUNE}(\Pi', \text{REACHABILITY}(\Pi'))$  // e.g. Alg. 7
4:   if  $\Pi' = \Pi$  then break
5:    $\Pi = \Pi'$ 
6: end loop
7: return  $\Pi'$ 

```

$s_{j-1}[v]$. The consequence is that we can construct an even shorter plan than π of equal or lesser cost, which is a contradiction, since π is already shortest optimal. \square

Reachability and Relevance Iteration

The backwards goal-relevance analysis of Algorithm 5 may remove some facts or actions from the task that cause other facts or actions to become unreachable from the initial state in the pruned task. Thus, we can potentially achieve additional task pruning by running forward reachability analysis (Blum and Furst 1997; Hoffmann and Nebel 2001a; Alcázar and Torralba 2015) on the pruned task (see Algorithm 7).

Theorem 5 *Given a task Π , $\bar{\mathcal{A}}_R = \text{REACHABILITY}(\Pi)$ contains all actions from all plans for Π .*

Proof:

At each iteration of the algorithm, we mark as reachable all actions whose preconditions are applicable for the (currently) reachable facts. We then mark all reachable actions' effects as reachable facts. All actions on plans are applicable, so this will retain them all. \square

Finally, running Algorithm 7 after Algorithm 5 may once again remove actions or facts from the task and cause new facts or actions to become irrelevant, making further pruning possible. We therefore introduce one last algorithmic improvement, which is to repeatedly interleave relevance and reachability pruning (see Algorithm 8).

Theorem 6 *Given a task Π , RELEVANCEITERATION with RELEVANCE- $\{\mathcal{V}, \text{F}, \text{FC}\}$ preserves all perfectly justified*

IPC (all 65 domains)	None	FD	V	F	FC	FCM	FCMR	FCMRL
Actions - Sum	17001557	16699722	16699722	16141206	16128063	16128035	16009673	16009340
Facts - Sum	1138370	858628	858628	854767	852463	852463	849761	849733
Variables - Sum	363961	237096	237096	235923	235660	235660	234309	234295
Coverage - Sum	970	982	981	990	989	988	986	986
Pruning (s) - Geo. mean	None	None	0.07	0.07	0.07	0.09	0.11	0.11
Translation (s) - Geo. mean	0.18	0.18	0.25	0.25	0.25	0.28	0.29	0.30
Search (s) - Geo. mean	2.09	1.83	1.82	1.70	1.67	1.68	1.67	1.68
Total (s) - Geo. mean	2.20	1.92	1.91	1.78	1.76	1.76	1.73	1.74

IPC (20 improved domains)	None	FD	V	F	FC	FCM	FCMR	FCMRL
Actions - Sum	4898612	4679923	4679923	4121407	4108264	4108236	3989874	3989541
Facts - Sum	682666	426090	426090	422229	419925	419925	417223	417195
Variables - Sum	234224	118942	118942	117769	117506	117506	116155	116141
Task Size - Sum	40320929	36693077	36567962	34117422	34011938	34011802	32954966	32952845
Coverage - Sum	283	289	289	298	297	296	296	295
Pruning (s) - Sum	None	None	517.04	532.31	599.50	3892.72	4556.40	5068.87
Translation (s) - Sum	1039.34	1042.41	1460.95	1487.78	1521.78	3686.79	4306.96	6092.84
Search (s) - Sum	23286.24	18633.88	18924.82	13263.46	12987.00	13048.43	13390.98	13243.94
Total (s) - Sum	23292.94	18640.26	18931.01	13269.79	12993.06	13054.59	13396.53	13249.37

Table 1: Results for IPC benchmark tasks. Fact-based relevance analysis results in additional pruning. (Top) All domains. (Bottom) Subset of domains where F+ pruning improves over V in some task.

plans, with RELEVANCE-FCM preserves all shortest optimal plans, and never introduces any new plans.

Proof:

Follows directly from Theorems 1, 2, 3, 4, and 5. □

Experiments

We perform experiments on a wide range of benchmark planning tasks from the International Planning Competition (McDermott 2000; Long and Fox 2003; Coles et al. 2012; Vallati et al. 2015) and from our newly designed Minecraft planning domain (Figure 2). All experiments used identical hardware (AMD Ryzen 9 7900X3D CPU, 250 GB of RAM). The IPC experiments used a search budget of 30 minutes, and the Minecraft ones used 5 minutes. We include the Minecraft domain and problem files in the supplementary materials, and will release them upon publication, along with the implementations of the pruning algorithms.

We implement our algorithms in the Fast Downward planning system (Helmert 2006) as a new step after translation but prior to search. The default behavior of the translator module runs the causal-graph-based variable relevance analysis, converts to FDR, and then runs REACHABILITY.⁴ We include this behavior as a baseline (denoted FD in Table 1), and show that it is equivalent to disabling the causal graph analysis (`--keep-unimportant-variables`) and running RELEVANCE-V (denoted v) after FDR conversion and REACHABILITY. We then implement each successive upgrade: F considers facts, FC considers causal links to s_0 , FCM considers action merging, FCMR performs an extra

⁴Technically, it is $\text{PRUNE} \circ \text{REACHABILITY}$, but in this section we omit the PRUNE and function composition for brevity.

REACHABILITY step, and FCMRL runs RELEVANCEITERATION with all of these improvements enabled. The NONE baseline performs no pruning prior to search.

We measure the effectiveness of the pruning and the subsequent search. In Table 1, we record the numbers of total facts, actions, and variables after pruning, along with solution coverage and time. We include results for all 65 domains, as well as the subset of 20 domains with at least one task where fact-based analysis (F+) reduces task size relative to variables-based (v). Overall, we find that fact-based pruning reduces total task size and improves coverage, without increasing total time.

Table 2 shows pruning performance for the Minecraft domain as the number of agents increases for FD and FCM. Fact-based relevance analysis reduces actions by up to 99%, facts and variables by up to 90%, and solves 50 more problems than the baseline Fast Downward algorithm. We visualize the relative pruning performance (in number of actions) for both algorithms in Figure 3. The relative gains from fact-based pruning increase with task size.

Related Work

The idea of using goal information to guide the search process dates back to the earliest planners (Newell, Shaw, and Simon 1959; Fikes and Nilsson 1971), which used the goal mainly for means-ends analysis. Subsequent work introduced the notion of “goal regression”—the idea of applying actions in reverse, working backwards from the goal until reaching the initial state (Warren 1974; Waldinger 1975). These approaches are similar in spirit to ours, but did not remove irrelevant information prior to search.

In the context of pruning irrelevant information, Nebel, Dimopoulos, and Koehler (1997) pointed out that tasks con-

n_agents	Coverage		Actions		Facts		Variables		Pruning (s)	Total (s)	
	FD	FCM	FD	FCM	FD	FCM	FD	FCM	FCM	FD	FCM
1	10	10	22	6	31	8	15	4	0.01	0.01	0.01
2	10	10	67	12	55	10	28	5	0.01	54.48	0.01
5	0	10	587	30	127	16	64	8	0.01	-	0.01
10	0	10	4,177	60	251	26	126	13	0.02	-	0.02
12	0	10	7,125	72	300	30	150	15	0.05	-	0.10
15	0	10	13,764	90	369	36	184	18	0.08	-	0.97
20	0	10	32,349	120	485	46	243	23	0.21	-	44.14
25	0	0	62,938	150	609	56	304	28	0.55	-	-

Table 2: Results for Minecraft domain by number of agents. Fact-based relevance analysis reduces actions by up to 99%, facts and variables by up to 90%, and solves 50 more problems than the baseline Fast Downward algorithm.

tain two crucial sources of guidance: the initial state and the goal. The authors argued that if an action or fact is unreachable by forward-chaining from the initial state or backward-chaining from the goal, then it cannot be part of any solution. They further proved that fully removing all irrelevant information is as hard as planning itself (i.e. PSPACE-hard), and while their work included a polynomial-time approximation algorithm, it was not always solution preserving. Boutilier, Brafman, and Geib (1998) and Brafman (2001) respectively introduced the Reachable- k and Relevance- k algorithms, which were solution preserving; however, increasing amounts of pruning came with a combinatorial increase in computational complexity. Hoffmann and Nebel (2001b) instead considered pruning a relaxed version of the task, which allowed them to safely improve the planner’s heuristic computation, but they did not prune the task itself.

Modern planners try to strike a balance between pruning cost and planning cost. The Fast Downward planning system (Helmert 2006) performs causal graph analysis of state variables to approximate which variables are goal-relevant and then uses relaxed reachability analysis to ignore actions that are unreachable from the initial state. We prove that this process is guaranteed to preserve all actions from justified plans, but that it is far from optimal.

Conclusion

We introduce several algorithms for pruning irrelevant task information from planning tasks. Our algorithms consider relevance at the level of facts, rather than variables, which enables much more aggressive pruning. The algorithms iterate until they reach a fixed point, and fact-based relevance analysis allows us to terminate this iteration early—either because precondition facts are causally linked to the initial state, or because actions can be merged together with simplified preconditions. We formally characterize the set of plans preserved by each algorithm we introduce, and in particular show that they all preserve shortest optimal plans without introducing new ones. We test our algorithms on a wide range of planning tasks, and show that the approach can significantly reduce problem size and planning time when tasks contain irrelevant information.

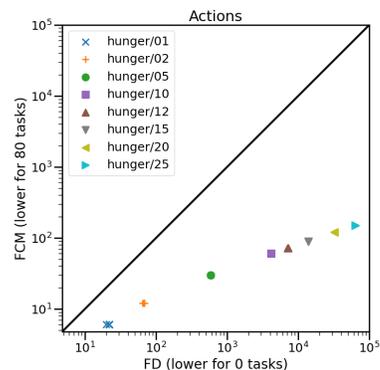


Figure 3: Action pruning comparison between FD and FCM.

References

- Alcázar, V.; and Torralba, Á. 2015. A Reminder about the Importance of Computing and Exploiting Invariants in Planning. In Brafman, R.; Domshlak, C.; Haslum, P.; and Zilberstein, S., eds., *Proceedings of the Twenty-Fifth International Conference on Automated Planning and Scheduling (ICAPS 2015)*, 2–6. AAAI Press.
- Bäckström, C.; and Nebel, B. 1995. Complexity Results for SAS⁺ Planning. *Computational Intelligence*, 11(4): 625–655.
- Blum, A.; and Furst, M. L. 1997. Fast Planning Through Planning Graph Analysis. *Artificial Intelligence*, 90(1–2): 281–300.
- Boutilier, C.; Brafman, R. I.; and Geib, C. 1998. Structured Reachability Analysis for Markov Decision Processes. In *Proceedings of the 14th Conference on Uncertainty in Artificial Intelligence (UAI 1998)*, 24–32.
- Brafman, R. 2001. On Reachability, Relevance, and Resolution in the Planning as Satisfiability Approach. *Journal of Artificial Intelligence Research*, 14: 1–28.
- Coles, A.; Coles, A.; García Olaya, A.; Jiménez, S.; Linares López, C.; Sanner, S.; and Yoon, S. 2012. A Survey of the Seventh International Planning Competition. *AI Magazine*, 33(1): 83–88.
- Fikes, R. E.; and Nilsson, N. J. 1971. STRIPS: A New Approach to the Application of Theorem Proving to Problem Solving. *Artificial Intelligence*, 2: 189–208.
- Fink, E.; and Yang, Q. 1992. Formalizing Plan Justifications. In *Proceedings of the 9th Canadian Conference on Artificial Intelligence (CAI 1992)*.
- Fink, E.; and Yang, Q. 1993. A Spectrum of Plan Justifications. In *Proceedings of the AAAI 1993 Spring Symposium*, 29–33. AAAI Press/MIT Press.
- Helmert, M. 2006. The Fast Downward Planning System. *Journal of Artificial Intelligence Research*, 26: 191–246.
- Helmert, M. 2009. Concise Finite-Domain Representations for PDDL Planning Tasks. *Artificial Intelligence*, 173: 503–535.

- Helmert, M.; Haslum, P.; and Hoffmann, J. 2007. Flexible Abstraction Heuristics for Optimal Sequential Planning. In Boddy, M.; Fox, M.; and Thiébaux, S., eds., *Proceedings of the Seventeenth International Conference on Automated Planning and Scheduling (ICAPS 2007)*, 176–183. AAAI Press.
- Hoffmann, J.; and Nebel, B. 2001a. The FF Planning System: Fast Plan Generation Through Heuristic Search. *Journal of Artificial Intelligence Research*, 14: 253–302.
- Hoffmann, J.; and Nebel, B. 2001b. RIFO Revisited: Detecting Relaxed Irrelevance. In Cesta, A.; and Borrajo, D., eds., *Proceedings of the Sixth European Conference on Planning (ECP 2001)*, 127–135. AAAI Press.
- Katz, M.; Röger, G.; and Helmert, M. 2022. On Producing Shortest Cost-Optimal Plans. In Chrupa, L.; and Saetti, A., eds., *Proceedings of the 15th Annual Symposium on Combinatorial Search (SoCS 2022)*, 100–108. AAAI Press.
- Long, D.; and Fox, M. 2003. The 3rd International Planning Competition: Results and Analysis. *Journal of Artificial Intelligence Research*, 20: 1–59.
- McAllester, D.; and Rosenblitt, D. 1991. Systematic Nonlinear Planning. In *Proceedings of the Ninth National Conference on Artificial Intelligence (AAAI 1991)*, 634–639. AAAI Press/MIT Press.
- McDermott, D. 2000. The 1998 AI Planning Systems Competition. *AI Magazine*, 21(2): 35–55.
- Nebel, B.; Dimopoulos, Y.; and Koehler, J. 1997. Ignoring Irrelevant Facts and Operators in Plan Generation. In Steel, S.; and Alami, R., eds., *Recent Advances in AI Planning. 4th European Conference on Planning (ECP 1997)*, volume 1348 of *Lecture Notes in Artificial Intelligence*, 338–350. Springer-Verlag.
- Newell, A.; Shaw, J. C.; and Simon, H. A. 1959. Report on General Problem-Solving Program. In *Proceedings of the 1st International Conference on Information Processing (IFIP 1959)*, 256–264.
- Salerno, M.; Fuentetaja, R.; and Seipp, J. 2023. Eliminating Redundant Actions from Plans using Classical Planning. In Marquis, P.; Son, T. C.; and Kern-Isberner, G., eds., *Proceedings of the Twentieth International Conference on Principles of Knowledge Representation and Reasoning (KR 2023)*, 774–778. IJCAI Organization.
- Sievers, S.; Wehrle, M.; and Helmert, M. 2014. Generalized Label Reduction for Merge-and-Shrink Heuristics. In Brodley, C. E.; and Stone, P., eds., *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence (AAAI 2014)*, 2358–2366. AAAI Press.
- Vallati, M.; Chrupa, L.; Grześ, M.; McCluskey, T. L.; Roberts, M.; and Sanner, S. 2015. The 2014 International Planning Competition: Progress and Trends. *AI Magazine*, 36(3): 90–98.
- Waldinger, R. 1975. Achieving several goals simultaneously. Technical Report 107, Stanford Research Institute.
- Warren, D. H. 1974. WARPLAN: A system for generating plans. Technical Report 76, University of Edinburgh.