

Novel Is Not Always Better: On the Relation between Novelty and Dominance Pruning

Joschka Groß

Saarland University
Saarland Informatics Campus
Saarbrücken, Germany
s8jagros@stud.uni-saarland.de

Álvaro Torralba

CISPA Heimholz Center
Saarland Informatics Campus
Saarbrücken, Germany
torralba@cs.uni-saarland.de

Maximilian Fickert

Saarland University
Saarland Informatics Campus
Saarbrücken, Germany
fickert@cs.uni-saarland.de

Abstract

Novelty pruning is a planning technique that focuses on exploring states that are novel, i.e., those containing facts that have not been seen before. This seemingly simple idea has had a huge impact on the state of the art in planning though its effectiveness is not entirely understood yet.

We relate novelty to dominance pruning, which compares states to previously seen states to eliminate those that are provably worse in terms of goal distance. Novelty can be interpreted as an unsafe approximation of dominance, where states containing novel facts are relevant because they enable new paths to the goal and, therefore, they are less likely to be dominated by others. This provides a framework to understand the success of novelty, resulting in new variants that combine both techniques.

Introduction

Novelty pruning is a technique that fosters the exploration of *novel* states (Lipovetzky and Geffner 2012). Novelty is measured as the size of the smallest tuple of facts that was not contained in any previously seen state, so the most novel states are those in which a fact is true for the first time. This seemingly simple idea has had a huge impact on the landscape of satisficing planning algorithms, leading to various combinations of novelty and heuristic search (Lipovetzky and Geffner 2017a; 2017b; Katz et al. 2017; Fickert 2018). Also, since novelty does not directly depend on the actions, it can be used even if a full action model is not available (Lipovetzky, Ramírez, and Geffner 2015; Shleyfman, Tuisov, and Domshlak 2016; Francès et al. 2017).

However, the effectiveness of novelty is not entirely understood. The success of novelty is commonly explained as a matter of balancing exploration and exploitation, where following the heuristic corresponds to exploitation and expanding novel states corresponds to exploration. But there are also common situations where exploring novel facts is not necessarily a good strategy. Consider a task where a truck must deliver several packages spending at most 100 units of fuel. In that case, driving back and forth between any two locations will lead to novel states because they have less fuel

available, even though this is clearly not progressing towards the goal. Intuitively, it would be better to focus on states that change the position of the truck and/or packages, but this requires analyzing the action model to distinguish which facts are most relevant to achieve the goal.

To shed light on the effectiveness of novelty-based methods, we consider their relation to dominance pruning methods. Dominance pruning techniques eliminate states that are provably “worse” than other previously seen states (Torralba and Hoffmann 2015; Torralba 2017; 2018). In the example above, states where the truck and packages are in the same position with lower fuel are dominated and can be pruned.

We define a family of unsafe dominance pruning methods that generalizes previous methods for both dominance and novelty pruning. We explore different members of this family, including variants that consider different sets of projections, as well as label-dominance (LD) novelty. LD novelty combines both approaches, only considering a fact novel if no *better* fact, in terms of the outgoing plans that it enables, has been seen before. The results of the evaluation of our new variants on IPC benchmarks confirm that using different projections leads to an improvement in many domains. LD novelty can also lead to more effective pruning without being less safe than standard novelty pruning, though novelty heuristic methods do not benefit from a direct integration.

Characterizing novelty as an unsafe dominance pruning method leads to a new interpretation about why exploring novel states is preferable. What really matters is whether the state has a path to the goal that was not possible from any state seen before. We argue that this view is consistent with previous approaches such as novelty heuristics, but putting the emphasis on directly comparing the set of plans, rather than facts, can lead to a better understanding of when novelty works and how to improve it. We suggest a new method to prioritize the expansion of non-novel states that aims to estimate the probability of not being dominated, consistently improving current state-of-the-art novelty techniques.

Background

We consider classical planning in the finite-domain representation (FDR) (Bäckström and Nebel 1995). A *classical planning task* is a 4-tuple $\langle \mathcal{V}, \mathcal{A}, s_{\mathcal{I}}, \mathcal{G} \rangle$, where \mathcal{V} is a set of

state variables, each with a finite domain \mathcal{D} , \mathcal{A} are the actions, each defined as a set of *preconditions pre* and *effects eff* (partial variable assignments), $s_{\mathcal{T}}$ is the *initial state*, a complete variable assignment, and \mathcal{G} is the goal, a partial variable assignment. A variable/value pair $var = val$, with $var \in \mathcal{V}$ and $val \in \mathcal{D}_{var}$, is called a *fact*. We denote the set of all facts by \mathcal{F} , and the set of all states (complete variable assignments) by \mathcal{S} . An action a is *applicable* in a state s if all its preconditions are satisfied in s . The result of applying a in s is the state $s[[a]]$, which is the same state as s , except for all variables v on which eff_a is defined, $s[[a]](v) := eff_a(v)$. A solution to a planning task is a sequence of actions π (a *plan*) leading from I to a state that satisfies the goal, and it is called *optimal* if it has minimal length ($|\pi|$) among all plans. We denote $h^*(s)$ to the length of an optimal plan from s .

Consider a task with variables $V = \{t, p, f\}$ that describe the position of a truck with $\mathcal{D}_t = \{T_A, T_B\}$, the position of a package with $\mathcal{D}_p = \{P_A, P_B, P_T\}$, and the amount of available fuel with $\mathcal{D}_f = \{F_{10}, \dots, F_0\}$ respectively. $I = \langle T_A, P_A, F_{10} \rangle$ and $G = \{P_B\}$. There are actions that allow to load/unload the package at the location of the truck, or drive between A and B consuming an unit of fuel.

A *transition system* (TS) is a tuple $\Theta = \langle \mathcal{S}, L, T, s^I, \mathcal{S}_G \rangle$ where \mathcal{S} is a set of *states*, L is a finite set of *labels*, $T \subseteq \mathcal{S} \times L \times \mathcal{S}$ is a set of *transitions*, $s^I \in \mathcal{S}$ is the *start state*, and $\mathcal{S}_G \subseteq \mathcal{S}$ is the set of *goal states*. A planning task induces a *state space*, which is a TS where: \mathcal{S} is the set of all states; $s^I = s_{\mathcal{T}}$; $s \in \mathcal{S}_G$ iff $\mathcal{G} \subseteq s$; $L = \mathcal{A}$, and $s \xrightarrow{a} s[[a]] \in T$ if a is applicable in s . We will use $s \in \Theta$ to refer to states in Θ and $s \xrightarrow{a} t$ to refer to their transitions.

Novelty Pruning

Given a set of states seen so far \mathcal{T} , the *novelty* of a state s is the size of the smallest tuple of facts in s that is not contained in any $s' \in \mathcal{T}$. This concept has been exploited in different forms, the simplest one being *novelty pruning*, where any state that does not qualify as novel with respect to previously generated states is pruned. We denote the novelty pruning method that prunes all states with novelty greater than k by \mathcal{N}_k (k -novelty pruning). This kind of pruning is used in *iterated width search* (IW) (Lipovetzky and Geffner 2012), where a single iteration $IW(k)$ is a breadth-first search with \mathcal{N}_k , thus expanding at most $|\mathcal{F}|^k$ states. This relates to the theoretical notion of *width*, which guarantees $IW(w)$ will find a solution on tasks with width w or less.

Recently, new novelty measures have been introduced that combine novelty with heuristic functions (Katz et al. 2017; Lipovetzky and Geffner 2017a), considering only states with greater or equal heuristic value for the computation of novelty. We briefly summarize some of the definitions by Katz et al. here, as we will introduce new variants of them. In the following definitions, we assume that there is a given set of states seen so far \mathcal{T} and a heuristic $h : \mathcal{S} \mapsto \mathbb{N}_0$. The *novelty score* of a fact f is defined as $N(f) = \inf_{s \in \mathcal{T}, f \in s} h(s)$ and the novelty score of a fact f in a state s is $N(f, s) = N(f) - h(s)$ if $f \in s$. The *binary novelty heuristic* h_{BN}

separates novel states from non-novel ones:

$$h_{BN}(s) = \begin{cases} 0 & \text{if } \exists f \in s, N(f, s) > 0 \\ 1 & \text{otherwise} \end{cases}$$

A similar novelty heuristic, which we denote by $h_{BN}^{\bar{}}$, has been introduced by Lipovetzky and Geffner (2017a). Under $h_{BN}^{\bar{}}$, a state is novel if it contains a fact that that was not contained in any previously seen state with the same heuristic value, whereas h_{BN} adds the restriction that the new heuristic value must be smaller than the ones seen before.

These binary novelty heuristics only separate novel states from non-novel ones, but it can be useful to have a more fine-grained separation. The *quantified novelty heuristic* h_{QN} also differentiates novel states from each other: $h_{QN}(s) = |\mathcal{V}| - \sum_{f \in s} N^+(f, s)$ where $N^+(f, s)$ is 1 if $N(f, s) > 0$, and 0 otherwise. The *quantified both heuristic* h_{QB} additionally differentiates non-novel states from each other:

$$h_{QB}(s) = \begin{cases} h_{QN}(s) & \text{if } h_{QN}(s) < |\mathcal{V}| \\ |\mathcal{V}| + \sum_{f \in s} N^-(f, s) & \text{otherwise} \end{cases}$$

where $N^-(f, s)$ is 1 if $N(f, s) < 0$, and 0 otherwise.

Dominance Pruning

Dominance methods also compare a state s to a set \mathcal{T} of already seen states to prune s if a *better* state t has been seen. Different notions of dominance differ on how they compare states, which is formalized in terms of a relation. A *dominance relation* is a relation $\preceq \subseteq \mathcal{S} \times \mathcal{S}$ if $s \preceq t$ implies $h^*(t) \leq h^*(s)$. Any such relation can be used to safely prune a state s if $\exists t \in \mathcal{T}$ s.t. $g(t) \leq g(s)$ and $s \preceq t$.

A relation \preceq is *goal-respecting* if whenever $s \preceq t$, $t \in \mathcal{S}_G \vee s \notin \mathcal{S}_G$. A relation \preceq is a *simulation* relation (Milner 1971) if, whenever $s \preceq t$, for all $s \xrightarrow{l} s'$, there exists $t \xrightarrow{l} t'$ s.t. $s' \preceq t'$. A *cost-simulation* allows the transition from t to use a different label, i.e., whenever $s \preceq t$, for every $s \xrightarrow{l} s'$, there exists $t \xrightarrow{l'} t'$ s.t. $s' \preceq t'$. Any goal-respecting cost-simulation is a dominance relation that leads to safe pruning.

Computing a relation directly on Θ is not feasible because there are exponentially many states. Instead, a compositional approach is followed where the task is described as a set of TSs $\{\Theta_1, \dots, \Theta_n\}$ such that $\Theta = \Theta_1 \otimes \dots \otimes \Theta_n$. The simplest choice, that we follow throughout this paper, is to use the atomic projections onto the variables of the planning task, so that we have a Θ_i for each variable in V (Helmert et al. 2014). We denote the projection of s on Θ_i by s_i , which corresponds to the fact $s(v_i)$. A set of relations $\{\preceq_1, \dots, \preceq_n\}$ is then computed on $\{\Theta_1, \dots, \Theta_n\}$. A relation for the complete state space Θ can be obtained as \preceq s.t. $s \preceq t$ iff $s_i \preceq_i t_i$ for each Θ_i . In other words, the individual relations \preceq_i compare facts associated with the same variable, and a state dominates another if it is at least as good according to all projections/variables. We say that a relation on facts \preceq_i is *compositionally safe* if, whenever $s_i \preceq_i t_i$, replacing the fact s_i by t_i in any state will not increase goal distance. In our running example, replacing lower amounts of fuel by larger amounts of fuel will never increase goal

distance. Similarly, replacing the location of a package by setting it to its goal location is never detrimental.

Label dominance (LD) simulation uses a fixpoint polytime algorithm to compute a set of coarser relations $\preceq^{LD} = \{\preceq_1^{LD}, \dots, \preceq_n^{LD}\}$ such that their combination is guaranteed to be a cost-simulation for the entire state space Θ , so they are compositionally safe (Torralba and Hoffmann 2015).

Combining Novelty and Dominance

Dominance and novelty pruning techniques are related as both compare newly generated states to the set of already seen states \mathcal{T} . Let $P = \{\Theta_1, \dots, \Theta_n\}$ be the set of atomic projections and $\{\preceq_1, \dots, \preceq_n\}$ a compositionally safe dominance relation on P . Then, a new state s can be pruned if:

- Safe dominance: $\exists t \in \mathcal{T} : \forall \Theta_i \in P : s_i \preceq_i t_i$
- Novelty \mathcal{N}_1 : $\forall \Theta_i \in P : \exists t \in \mathcal{T} : s_i = t_i$

The fundamental difference between novelty and safe dominance pruning is that, by swapping the quantifiers, the former allows using different states t^1, \dots, t^k for each variable/projection. To understand this difference, we analyze the relation between facts and the plans enabled by them. Each state s in a transition system has a corresponding set of plans $\vec{A}(s)$. These are all possible paths starting from s and ending at some goal state, possibly containing cycles and self-loop transitions. Since we have defined a TS Θ_i for each variable v_i , we can also associate each fact of the planning task $s_i \in \Theta_i$ with a set of (abstract) plans $\vec{A}(s_i)$. The state space of the planning task Θ is equal to the product $\bigotimes_{i \in [1, n]} \Theta_i$, so the set of valid plans for a state $s \in \Theta$ is the intersection of the set of valid plans of each of its facts, $\vec{A}(s) = \bigcap_{i \in [1, n]} \vec{A}(s_i)$.

If each relation \preceq_i is a simulation relation for Θ_i and $s_i \preceq_i t_i$, then all paths for s_i are paths for t_i , $\vec{A}(s) \subseteq \vec{A}(t)$. It is clear then why this is compositionally safe: by the properties of intersection, replacing s_i by t_i can only grow the set of plans for a state. However, simulation relations are too strict. An LD-simulation obtains coarser relations that lead to more pruning by allowing t_i to use a different plan (e.g., a shorter one) such that if the plan from s_i is valid in all other projections, the plan from t_i must be valid too.

Note that requiring that a single state t dominates s is unnecessarily strict, because each plan from s could be proven to be dominated by paths starting at different $t \in \mathcal{T}$. We say that a set of states \mathcal{T} dominates s if for each $\vec{a}_s \in \vec{A}(s)$, there exist $t \in \mathcal{T}$ and $\vec{a}_t \in \vec{A}(t)$ such that $|\vec{a}_t| \leq |\vec{a}_s|$. However, in the compositional approach, the same state t must be used to prove that s is dominated for all projections. When novelty pruning uses different states t^i to dominate s for each projection Θ_i , pruning is not safe anymore because s may have a plan that is present on t_i^i but is spurious — t^i may not have such a path on a different projection. We model this through the parameter \mathcal{Q} that represents the sets of variables that must be dominated by the same state.

The other difference between novelty and dominance pruning is that the latter is defined over an arbitrary set of relations \mathcal{R} and projections P .

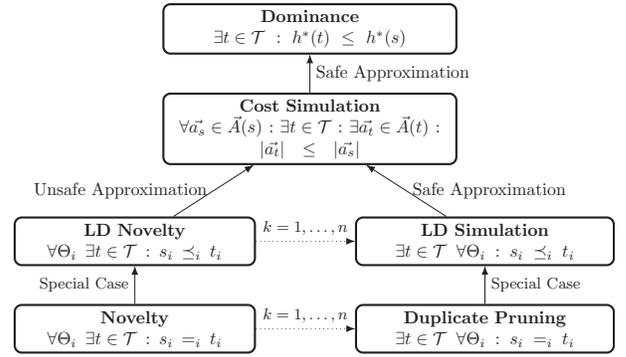


Figure 1: The connection between different notions of dominance and novelty summed up in an illustration. The formulas below each concept describe when a new state s is pruned, given a set of already seen states \mathcal{T} .

Definition 1 (Unsafe Dominance) Let $P = \{\Theta_1, \dots, \Theta_k\}$ be a set of projections. Let $\mathcal{R} = \{\preceq_1, \dots, \preceq_k\}$ be a set of relations on P and let $\mathcal{Q} \subseteq 2^{2^P}$ be a set of subsets of P . A newly generated state s is pruned under unsafe-compositional dominance if:

$$\forall \mathcal{Q} \in \mathcal{Q} : \exists t \in \mathcal{T} : \forall \Theta_i \in \mathcal{Q} : s_i \preceq_i t_i$$

In this paper we fix P to be the set of atomic projections. This definition generalizes previous definitions of dominance and novelty pruning, which correspond to appropriate choices of the remaining two parameters: \mathcal{Q} , and \mathcal{R} .

Figure 1 shows a high-level diagram of how different instantiations relate to each other. As subset selector we consider \mathcal{Q}_K which returns all subsets of size K . All of them approximate the notion of a cost-simulation by computing a compositional relation. Then, safe dominance pruning corresponds to using $\mathcal{R} = \preceq^{LD}$ and $\mathcal{Q} = \mathcal{Q}_k$, whereas the novelty pruning method \mathcal{N}_i corresponds to $\mathcal{R} = \preceq^=$ and $\mathcal{Q} = \mathcal{Q}_i$. Combining these ideas results in a new method, which we call LD-Novelty (\mathcal{N}_i^{LD}), where $\mathcal{R} = \preceq^{LD}$ and $\mathcal{Q} = \mathcal{Q}_i$.

A New View On Novelty

The unsafe dominance pruning framework provides a new interpretation on why novelty pruning works. Novelty approximates the ideal notion of dominance where a state is pruned if there is another state that is closer to the goal. To do so, it follows an approach similar to a cost simulation where the set of plans possible from s is compared to the set of plans possible from all other states. By comparing the sets of plans under different projections of the planning task, novelty identifies states having paths that no other previous state had. According to this view, novel states are preferred because they are less likely to be dominated by others, and the success of novelty approaches is then linked to the correlation between novel and dominated states.

We identify three sources of error (E1)–(E3) that may cause a state to be incorrectly classified as novel or non-novel. On the one hand, a state s may be misclassified as novel due to some fact s_i (or fact tuple) that has not been

seen yet, but it is dominated because there exists $t \in \mathcal{T}$ such that $h^*(t) \leq h^*(s)$. This may be because (E1) the set of paths that are possible from s_i is dominated by the set of paths that are possible from t_i , i.e., for all $\vec{a}_s \in s_i$ there exists $\vec{a}_t \in t_i$ such that \vec{a}_t is as good as \vec{a}_s . Another reason (E2) is that s_i has a plan better than any of t_i , but the path is spurious on s .

On the other hand, s may be misclassified as non-novel despite being the closest state to the goal. In that case, in all projections, there exists some t_i for which the optimal plan for s is valid. However, (E3) the path from t may be spurious so s is not dominated and wrongly classified as non-novel.

The two parameters \mathcal{R} and \mathcal{Q} offer a trade-off between these sources of error. Coarser relations (\mathcal{R}) directly aim to ameliorate (E1), and could be useful whenever too many states are being considered novel. This leads to more pruning, possibly increasing the error (E3) as one would expect from arbitrary relations. Compositionally safe relations (e.g. \preceq^{LD}) are particularly suitable for this purpose and should not be greatly affected by this since, everything else being equal, if the path from t is spurious, the path from s is also spurious. As pruning is unsafe the “everything else being equal” does not hold, but in that case one could expect s being novel due to a different variable.

Considering more or larger projections (\mathcal{Q}) can help to ameliorate (E3) by increasing the number of states that are considered novel. Indeed, by choosing $\mathcal{Q} = \mathcal{Q}_n$ pruning becomes safe and no state is incorrectly classified as non-novel. However, this is very likely to increase (E1) significantly reducing the number of non-novel states.

Next, we consider how this view fits with different novelty-based algorithms and how our new variants of novelty pruning can be applied in that context.

Iterated Width

As a means to explain the success of novelty, previous work introduced the notion of *width* of a planning task. The width of a planning task w is the minimum k such that there exists a list of k -tuples of facts t^1, \dots, t^m such that $t^1 \subseteq s_{\mathcal{I}}, \mathcal{G} \subseteq t^m$, and for any $i \in [1, m-1]$, every optimal plan for t^i can be extended by one action into an optimal plan for t^{i+1} .

The $IW(k)$ algorithm that combines breadth-first search with k -novelty pruning is guaranteed to find a solution on any planning task with width k . In other words, the *width* of a planning task is a sufficient condition to guarantee that novelty pruning \mathcal{N}_k is safe for breadth-first search, i.e., preserves at least one solution, though not necessarily optimal.

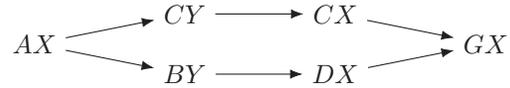
We define a new variant of iterated width, $IW^{\preceq}(k)$ where \mathcal{N}_k is replaced by \mathcal{N}_k^{LD} . This variant has more aggressive pruning, where some facts are not considered novel even if they have never been seen, because other (better) facts have been seen before. In our example, all facts involving the remaining amount of fuel will be ignored by $IW^{\preceq}(1)$ because the initial state has a fact with the maximum amount of fuel. While this is intuitively helpful, it can potentially eliminate useful states too. For example, if $\mathcal{T} = \{\langle T_A, F_6, P_T \rangle, \langle T_B, F_2, P_A \rangle\}$ have already been seen, then the potentially useful state $s = \langle T_B, F_5, P_T \rangle$ will be pruned.

However, standard novelty would also prune $\langle T_B, F_6, P_T \rangle$, which is strictly better than s .

Using compositionally safe relations guarantees the pruning to be safe whenever $k = n$. However, the following counterexample shows that \mathcal{N}_w^{LD} is not safe anymore for tasks of width w .

Theorem 1 *There exist planning tasks with width 1 where $IW^{\preceq}(1)$ does not find a solution.*

Proof. Consider the following planning task with two variables v_1 and v_2 with domains $\mathcal{D}_{v_1} = \{A, B, C, D, G\}$ and $\mathcal{D}_{v_2} = \{X, Y, Z\}$. The initial state is $\langle v_1 = A, v_2 = X \rangle$ and the goal is $v_1 = G$. All actions have two preconditions, and they result in the following state space:



The task has width 1 because of the path $A \rightarrow B \rightarrow D \rightarrow G$, so the goal can be reached by a path where there is a novel fact at each step. Moreover, a relation such that $B \preceq C$ is compositionally safe because replacing B by C in any state does not increase goal distance. However, $IW^{\preceq}(1)$ does not find a solution because after expanding AX , CY and BY are generated but BY is pruned because of CY , and CX is pruned because it is not novel. \square

The reason is that, whenever a tuple of facts t^i dominates another s^i , then by the properties of compositional dominance, if s^i has a path to the goal, so does t^i . However, by the definition of width, the path from s^i is guaranteed to not be pruned by novelty pruning, whereas the same does not necessarily hold for the path from t^i . However, this does not imply that LD-novelty pruning \mathcal{N}_i^{LD} is less safe than \mathcal{N}_i . Indeed, opposite examples that are only solved when using a dominance relation also exist. This may happen if the dominance relation prunes some states that do not lead to the goal. As those states are pruned, others in a path to the goal may become novel. Our hypothesis, supported by the empirical data of the experimental results section, is that if \mathcal{N}_i^{LD} uses compositionally safe relations, this pruning will not be less safe than \mathcal{N}_i in practice.

Binary Novelty Heuristics

To obtain state of the art performance, novelty is combined with heuristics to guide the search towards the goal. Here, we analyze the novelty heuristics introduced in the background under our interpretation of novelty as an approximate method to determine whether the state has a path to the goal that was not present in any previous state.

The binary novelty heuristics h_{BN} and $h_{\text{BN}}^{\overline{\cdot}}$ consider a state novel if it has a novel fact. The main difference with respect to \mathcal{N}_1 is that they take a heuristic function into account for the novelty. This perfectly fits our view on novelty. A fact should be considered novel if it enables a new path that was not possible on any state in \mathcal{T} . The heuristic refines our belief on which paths are possible from s to the goal. Since it estimates the minimum length of a plan from s , it restricts the set of valid plans to the subset of plans that have length

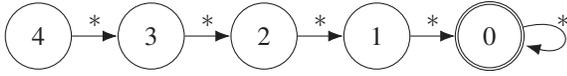


Figure 2: Θ_h for a heuristic with maximum value of 4.

greater than or equal to $h(s)$, excluding all shorter plans. If the heuristic value is lower than any other state with the same fact seen so far, the set of plans that are possible according to that projection is not dominated by any previous state.

To include the information of the heuristic in our framework, one can extend the partition $P = \{\Theta_1, \dots, \Theta_n\}$ with an additional transition system Θ_h that represents the information about the length of a plan, like the one in Figure 2. Then, by choosing $\mathcal{Q} = \{\{\Theta_i, \Theta_h\} \text{ for } i \in [1, n]\}$, we obtain both methods depending on our choice for \mathcal{R} : h_{BN}^- corresponds to using the identity relation, and h_{BN} corresponds to $\mathcal{R} = \{\preceq_1^-, \dots, \preceq_n^-, \preceq_h^{LD}\}$.

Using LD novelty in this context $\mathcal{R} = \{\preceq_1^{LD}, \dots, \preceq_n^{LD}, \preceq_h^{LD}\}$ means that anytime that a fact is seen in a state with a given heuristic value h , all dominated facts (e.g. having less fuel) are considered to be seen with h too.

Quantitative Novelty Heuristics

A limitation of the previous approaches is that they only divide states into novel and non-novel without further separation. The quantified novelty heuristics h_{QN} and h_{QB} provide a numeric value that prioritizes some novel and non-novel states over others. Ideally, this priority corresponds to the confidence of the novelty assessment, estimating the probability that the state has a better plan than other seen states.

The quantification is different for novel and non-novel states. For novel states, they simply count the number of novel facts. The intuition is that if more facts are novel, there are more novel paths so it is less likely that all of them are dominated by other paths or spurious.

For a non-novel state s , h_{QB} counts the number of facts where $N(f, s) < 0$, i.e., how many facts have been seen before in some other state $t \in \mathcal{T}$ with $h(t) < h(s)$. However, this is not entirely consistent with our interpretation since it does not directly aim to measure the probability that the assessment in s being non-novel is accurate. A state s is incorrectly classified as non-novel whenever $h^*(s) < h^*(t)$ for all $t \in \mathcal{T}$. This means that there is a plan \vec{s} for s , such that no other $t \in \mathcal{T}$ has a better path. However, since the state was classified as non-novel this means that for all projections $Q \in \mathcal{Q}$ there is some $t \in \mathcal{T}$ whose projection onto Q has a plan \vec{a}_t at least as good as any plan \vec{s}_Q for s in Q . The classification is incorrect due to (E3) whenever all possible such paths from all such t are spurious. In principle, this is not directly related to whether $h(s) = h(t)$ or $h(s) > h(t)$.

We propose an alternate measure for quantifying the non-novel states. In general, the likelihood of all plans being spurious is inversely proportional to the number of states t that have at least one such path. Therefore, our new priority function for non-novel states replaces N^- by:

$$N_{\rightarrow}^-(f, s) = 1 - \frac{1}{|\{t \in \mathcal{T} \mid s[i] \preceq_i t[i] \wedge h(t) \leq h(s)\}|}$$

This can be efficiently computed by keeping a counter for each fact-heuristic value pair (s_i, h) on the number of states seen with that heuristic value and the same or a better fact.

Experiments

We implemented the described techniques in Fast Downward (Helmert 2006). We run experiments on a cluster of machines with Intel Xeon E5-2660 CPUs with a clock rate of 2.2GHz using the Lab framework (Seipp et al. 2017). The time and memory limits are set to 30 minutes and 4GB. We consider all STRIPS domains from the satisficing tracks of all International Planning Competitions (IPC) up to 2018.

Effective Width Analysis

We first evaluate our hypothesis that using compositionally safe relations can increase the amount of pruning, while remaining as safe. To do so, we measure the *effective width*, i.e., how many instances are solved in practice by $\text{IW}(k)$ compared to $\text{IW}^{\preceq}(k)$ for different relations \preceq .

| Relation | $w_e = 1$ | $w_e = 2$ | $w_e > 2$ |
|-------------------|-----------|-----------|-----------|
| $\preceq^=$ | 40.42% | 46.25% | 13.32% |
| \preceq^{LD} | 40.35% | 45.82% | 13.83% |
| \preceq^{LD2} | 38.80% | 38.99% | 22.21% |
| \preceq^{LD1} | 36.81% | 21.32% | 41.87% |
| \preceq^{LD0} | 23.82% | 0.00% | 76.18% |
| \preceq^{LDinv} | 21.80% | 27.86% | 50.34% |

Table 1: Percentage of tasks with effective width of 1, 2, or greater per domain on average on tasks with a single goal. Safe relations ($\preceq^=$ and \preceq^{LD}) are compared against unsafe approximations of \preceq^{LD} .

Table 1 shows the effective width of instances with a single goal, as reported in the evaluation by Lipovetzky and Geffner (2012). For each IPC instance, we generate at most 100 instances by selecting single goals at random if the task has more than 100 goals. The identity relation ($\preceq^=$) is the baseline and corresponds to standard novelty pruning. LD-simulation (\preceq^{LD}) is a compositionally safe relation. The results are consistent with our hypothesis, as there is no significant change in the effective width of the instances between these two configurations. Fig. 3 compares the number of expansions of $\text{IW}(2)$ and $\text{IW}^{\preceq}(2)$ to show that indeed, pruning is more aggressive when using \preceq^{LD} . This is especially the case for instances with a single goal fact, since the dominance relations can be much coarser in that case as many variables become irrelevant. It is remarkable that the number of instances solved is practically equal, even though the number of expansions is reduced by up to three orders of magnitude. The picture is similar for $\text{IW}(1)$.

To analyze whether the same result would be obtained by relations that are not compositionally safe, we introduce relations without this property. For this, we use the same fixpoint algorithm used to compute LD-simulation relations (Torralba and Hoffmann 2015) but stopping it ear-

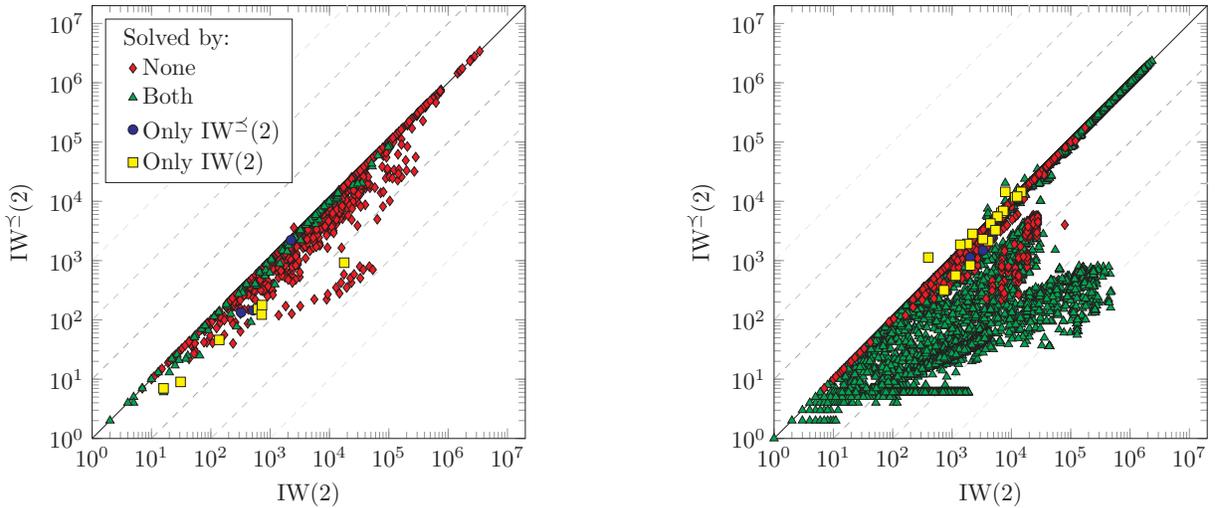


Figure 3: Expansions made by IW(2) and $IW^=(2)$ on IPC instances (left) and single goal instances (right).

lier, before the fixpoint is reached. This algorithm initializes the set of relations with a coarser relation \preceq^{LD0} where a fact t_i dominates another s_i if the goal-distance of t_i in the projection Θ_i is not greater than that of s_i . Therefore, all facts belonging to a variable with no goal defined are considered equivalent according to this relation. Afterwards, this relation is iteratively refined, by removing all pairs (s_i, t_i) that do not fulfill the properties of LD-simulation under an overapproximated label dominance, resulting in relations $\preceq^{LD1}, \preceq^{LD2}, \dots$ until it converges on \preceq^{LD} . All these variants result in more aggressive pruning than \preceq^{LD} , and they are not guaranteed to be compositionally safe. We also compare against the inverse relation, \preceq^{LDinv} , defined as $s \preceq^{LDinv} t \Leftrightarrow t \preceq^{LD} s$. \preceq^{LDinv} considers non-novel any fact if some fact that is worse according to \preceq^{LD} has been seen before, resulting in extremely unsafe pruning. Indeed, the results of Table 1 show how IW(1) and IW(2) with non-compositionally safe relations solve much fewer tasks, significantly increasing the effective width. Note that, with unsafe relations the iterated-width search is no longer complete so the effective width could be infinity in many tasks.

Novelty Heuristics

We focus our evaluation on the quantified both heuristic with underlying delete relaxation (h^{FF}) heuristics (Hoffmann and Nebel 2001). Following the evaluation in Katz et al.’s work (2017), the heuristic is used in greedy best-first search with lazy evaluation, without other search enhancements like preferred operators. Ties between states with the same h_{QB} values are broken by the underlying heuristics. We analyze, one by one, the impact of changing the relations (\mathcal{R}), sets of projections (\mathcal{Q}), and priority of non-novel states (N^-).

Relations (\mathcal{R}) Table 2 shows a comparison of the coverage obtained by different variants of the quantified both heuristics when using $\preceq^=$ and \preceq^{LD} as relations (Q_1 and Q^{pre} are defined below). The comparison shows that the

| | Atomic Q_1 | | Projections Q^{pre} | |
|----------------|--------------|-----------|-----------------------|-----------|
| | N^- | N^- | N^- | N^- |
| $\preceq^=$ | 1564 | 6 16 1598 | 1618 | 11 8 1593 |
| \preceq^{LD} | 1499 | 6 18 1526 | 1528 | 16 8 1493 |

Table 2: Analysis of $\preceq^=$ vs. \preceq^{LD} , and N^- vs. N^- for two choices of \mathcal{Q} . The large numbers show the total coverage of each configuration. The small numbers on each edge show the number of domains in which one configuration has higher coverage than the other, e.g. the configuration with atomic projections, N^- , and $\preceq^=$ beats the same configuration but using \preceq^{LD} in 14 domains and loses in 6.

baseline $\preceq^=$ consistently obtains better coverage, though using dominance relations does help in a few domains. This does not accurately represent the impact of \preceq^{LD} on search effort, since results are biased by the cost of computing the dominance relation, which may be prohibitive on instances with a very large number of actions and/or variables. Fig. 4a directly compares the expansions made by both configurations. The impact of using a dominance relation can be very positive whenever this helps to explore a better set of novel states, avoiding the expansion of any novel state. But there are also cases where reducing the number of novel states hurts performance. In some cases the results are affected by the effect of tie-breaking, but there are domains where using \preceq^{LD} consistently reduces the number of expansions (Driver-Log, Gripper, Maintenance, Rovers, Satellite, TidyBot, and Woodworking), and others where it is consistently worse (Parking, TPP, and NoMystery). A reason for this is that the dominance-based novelty is more greedy about achieving sub-goals because they dominate other facts, which is a

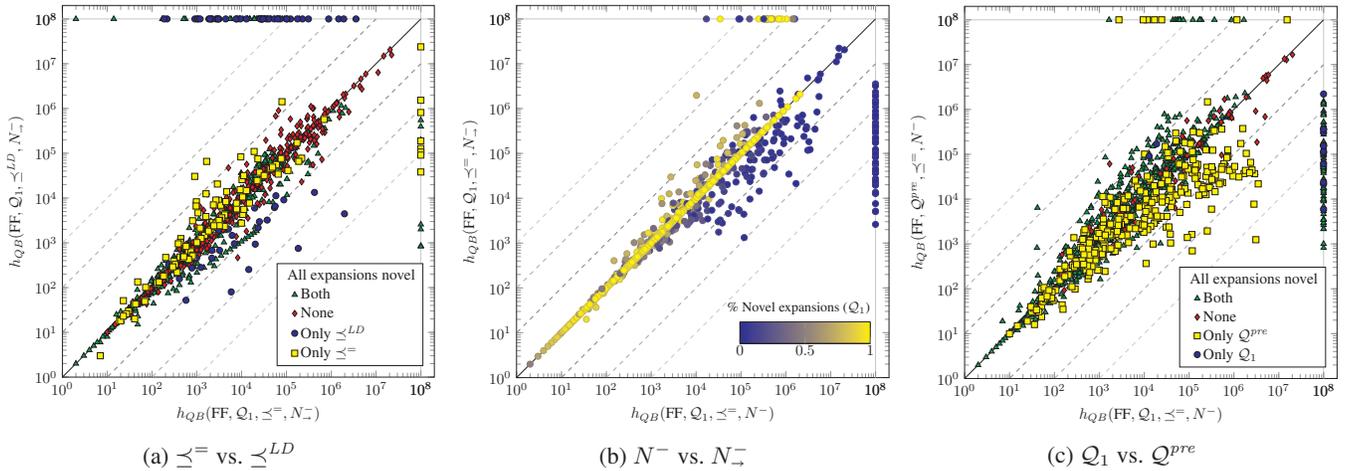


Figure 4: Expansions of lazy greedy best-first search with h_{QB} (h^{FF}) when changing (a) \mathcal{Q} , (b) \mathcal{R} , and (c) the priority for non-novel states. Colors provide information about the number of novel/non-novel states explored on solved tasks: Figs (a) and (c) show whether all expansions were novel and Fig (b) shows the percentage of novel states explored by the baseline.

good strategy for some domains but not for others.

Quantification of Non-Novel States The new priority function for non-novel states (N_{\downarrow}) outperforms the baseline if atomic projections are used (see Table 2). Fig. 4b shows a detailed picture of the number of expansions in that case. Clearly, when the baseline expands a big percentage of non-novel states, there is more margin of improvement. In those cases, N_{\downarrow} can reduce the number of expansions significantly, in many cases over one order of magnitude. If larger projections are used (Q^{pre} , see below), however, N_{\downarrow} is not as beneficial. The reason is that using larger projections greatly increases the number of novel states, so that the priority for non-novel states has a lower impact and the overhead of computing N_{\downarrow} over N_{\downarrow} reduces the performance.

Projections (\mathcal{Q}) We consider several methods to generate projections for the planning task. Q_1 and Q_2 consider all projections of size 1 and 2 respectively. Other variants try to select variables that are related to each other, either in the causal graph (*cg*) (Helmert 2004), or because there is an action whose precondition depends on both variables (*pre*). $Q_{1,2}^{cg}$ and $Q_{1,2}^{pre}$ use all atomic projections plus all projections that contain exactly two related variables. Q^{pre} and Q^{cg} use a single projection per variable, but including as many directly related variables as possible up to a limit of 1000 abstract states. Table 3 shows a comparison of all these configurations (using $\leq^=$) and Fig. 4c directly compares the best configuration, Q^{pre} , against the baseline Q_1 , which corresponds to the standard h_{QB} introduced in previous work. Overall, considering projections is beneficial, mainly because by greatly increasing the number of novel states, it often avoids exploring any non-novel state. The best performance is obtained by configurations that do not choose too many projections and do not impose a strict limit of 2 on the number of variables in a single projection. The

| | Q_1 | Q_2 | $Q_{1,2}^{cg}$ | $Q_{1,2}^{pre}$ | Q^{cg} | Q^{pre} | Total |
|-----------------|-------|-------|----------------|-----------------|----------|-----------|-------|
| Q_1 | – | 14 | 8 | 9 | 8 | 9 | 1564 |
| Q_2 | 17 | – | 6 | 6 | 8 | 6 | 1551 |
| $Q_{1,2}^{cg}$ | 20 | 15 | – | 7 | 10 | 10 | 1609 |
| $Q_{1,2}^{pre}$ | 17 | 16 | 8 | – | 9 | 7 | 1618 |
| Q^{cg} | 20 | 20 | 15 | 13 | – | 6 | 1630 |
| Q^{pre} | 17 | 17 | 13 | 15 | 8 | – | 1634 |

Table 3: Pairwise comparison of different projections. The value in row r and column c indicates the number of domains in which the configuration in row r has higher coverage than the one in column c . The last column shows total coverage.

| | BFWS | MERW. | Q_1 | N_{\downarrow} | Q^{pre} | Total |
|------------------|------|-------|-------|------------------|-----------|-------|
| BFWS | – | 10 | 17 | 13 | 16 | 1628 |
| MERW. | 14 | – | 17 | 12 | 12 | 1637 |
| Q_1 | 13 | 13 | – | 6 | 8 | 1608 |
| N_{\downarrow} | 14 | 14 | 14 | – | 9 | 1626 |
| Q^{pre} | 15 | 15 | 15 | 12 | – | 1658 |

Table 4: Pairwise comparison to state-of-the-art planners.

impact of different mechanisms to choose related variables (*cg* or *pre*) is not too large, but in general it seems that *pre* is slightly better. Choosing variables that appear together in the precondition of actions may be helpful since it influences whether the paths in the projection are spurious or not.

Comparison to State of the Art We compare our best configurations against the state-of-the-art novelty-based planners Dual-BFWS (Lipovetzky and Geffner 2017a; Francès et al. 2018), which uses h_{BN}^{\downarrow} with multiple base heuristics based on delete relaxation and landmarks (Porteous, Sebastia, and Hoffmann 2001), and MERWIN (Katz et al. 2018), which combines h_{QB} with the red-black partial delete relaxation heuristic (Domshlak, Hoffmann, and

Katz 2015). In this comparison, we combine both h^{FF} and the landmark-count heuristic (Porteous, Sebastia, and Hoffmann 2001; Richter, Helmert, and Westphal 2008) in our configurations (using the novelty heuristic based on both as the main evaluator, then breaking ties by h^{FF} and h^{LM} in that order), which yielded the best results in our experiments. Table 4 shows a domain-wise comparison with the aforementioned planners. Our configurations are (1) Q_1 , our baseline using atomic projections and identity relations, (2) N_- , the same as (1) but separating non-novel states by N_- , and (3) Q^{pre} , the same as (1) but using Q^{pre} projections. Our configurations are very competitive, with Q^{pre} beating both Dual-BFWS and MERWIN in terms of overall coverage. Overall, all shown configurations are very complementary (in particular Q^{pre} and Dual-BFWS), even though they are all based on the same technique (i.e. novelty pruning).

Conclusion

Novelty and dominance pruning are closely related methods that compare newly generated states against all previously known states. While novelty methods are traditionally defined in terms of novel facts, we re-define them as an approximate dominance pruning method that compares the outgoing action paths in a set of projections. This enables several extensions, such as combining notions of dominance and novelty pruning to avoid considering facts novel that do not enable better paths to the goal.

Acknowledgments This work has been partially supported by BMBF through funding for CISPA (grant no. 16KIS0656), and DFG through grant 389792660 as part of TRR 248 (see <https://perspicuous-computing.science>).

References

Bäckström, C., and Nebel, B. 1995. Complexity results for SAS⁺ planning. *Computational Intelligence* 11(4):625–655.

Domshlak, C.; Hoffmann, J.; and Katz, M. 2015. Red-black planning: A new systematic approach to partial delete relaxation. *Artificial Intelligence* 221:73–114.

Fickert, M. 2018. Making hill-climbing great again through on-line relaxation refinement and novelty pruning. In Bulitko, V., and Storandt, S., eds., *Proceedings of the 11th Annual Symposium on Combinatorial Search (SOCS'18)*. AAAI Press.

Francès, G.; Ramírez, M.; Lipovetzky, N.; and Geffner, H. 2017. Purely declarative action descriptions are overrated: Classical planning with simulators. In Sierra, C., ed., *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI'17)*, 4294–4301. AAAI Press/IJCAI.

Francès, G.; Lipovetzky, N.; Geffner, H.; and Ramírez, M. 2018. Best-first width search in the IPC 2018: Complete, simulated, and polynomial variants. In *IPC 2018 planner abstracts*.

Helmert, M.; Haslum, P.; Hoffmann, J.; and Nissim, R. 2014. Merge & shrink abstraction: A method for generating lower bounds in factored state spaces. *Journal of the Association for Computing Machinery* 61(3):16:1–16:63.

Helmert, M. 2004. A planning heuristic based on causal graph analysis. In Koenig, S.; Zilberstein, S.; and Koehler, J., eds., *Proceedings of the 14th International Conference on Automated Plan-*

ning and Scheduling (ICAPS'04), 161–170. Whistler, Canada: Morgan Kaufmann.

Helmert, M. 2006. The Fast Downward planning system. *Journal of Artificial Intelligence Research* 26:191–246.

Hoffmann, J., and Nebel, B. 2001. The FF planning system: Fast plan generation through heuristic search. *Journal of Artificial Intelligence Research* 14:253–302.

Katz, M.; Lipovetzky, N.; Moshkovich, D.; and Tuisov, A. 2017. Adapting novelty to classical planning as heuristic search. In *Proceedings of the 27th International Conference on Automated Planning and Scheduling (ICAPS'17)*, 172–180. AAAI Press.

Katz, M.; Lipovetzky, N.; Moshkovich, D.; and Tuisov, A. 2018. MERWIN planner: Mercury enhanced with novelty heuristic. In *IPC 2018 planner abstracts*.

Lipovetzky, N., and Geffner, H. 2012. Width and serialization of classical planning problems. In Raedt, L. D., ed., *Proceedings of the 20th European Conference on Artificial Intelligence (ECAI'12)*, 540–545. Montpellier, France: IOS Press.

Lipovetzky, N., and Geffner, H. 2017a. Best-first width search: Exploration and exploitation in classical planning. In Singh, S., and Markovitch, S., eds., *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI'17)*, 3590–3596. AAAI Press.

Lipovetzky, N., and Geffner, H. 2017b. A polynomial planning algorithm that beats LAMA and FF. In *Proceedings of the 27th International Conference on Automated Planning and Scheduling (ICAPS'17)*, 195–199. AAAI Press.

Lipovetzky, N.; Ramírez, M.; and Geffner, H. 2015. Classical planning with simulators: Results on the atari video games. In Yang, Q., ed., *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI'15)*, 1610–1616. AAAI Press/IJCAI.

Milner, R. 1971. An algebraic definition of simulation between programs. In *Proceedings of the 2nd International Joint Conference on Artificial Intelligence (IJCAI'71)*, 481–489. London, UK: William Kaufmann.

Porteous, J.; Sebastia, L.; and Hoffmann, J. 2001. On the extraction, ordering, and usage of landmarks in planning. In Cesta, A., and Borrajo, D., eds., *Proceedings of the 6th European Conference on Planning (ECP'01)*, 37–48. Springer-Verlag.

Richter, S.; Helmert, M.; and Westphal, M. 2008. Landmarks revisited. In Fox, D., and Gomes, C., eds., *Proceedings of the 23rd National Conference of the American Association for Artificial Intelligence (AAAI'08)*, 975–982. Chicago, Illinois, USA: AAAI Press.

Seipp, J.; Pommerening, F.; Sievers, S.; and Helmert, M. 2017. Downward Lab. <https://doi.org/10.5281/zenodo.790461>.

Shleyfman, A.; Tuisov, A.; and Domshlak, C. 2016. Blind search for atari-like online planning revisited. In Kambhampati, S., ed., *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI'16)*, 3251–3257. AAAI Press/IJCAI.

Torralba, Á., and Hoffmann, J. 2015. Simulation-based admissible dominance pruning. In Yang, Q., ed., *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI'15)*, 1689–1695. AAAI Press/IJCAI.

Torralba, Á. 2017. From qualitative to quantitative dominance pruning for optimal planning. In Sierra, C., ed., *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI'17)*, 4426–4432. AAAI Press/IJCAI.

Torralba, Á. 2018. Completeness-preserving dominance techniques for satisficing planning. In Lang, J., ed., *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI'18)*, 4844–4851.