Novelty vs. Potential Heuristics: A Comparison of Hardness Measures for Satisficing Planning

Simon Dold, Malte Helmert

University of Basel, Switzerland

{simon.dold, malte.helmert}@unibas.ch

Abstract

Classical planning considers a given task and searches for a plan to solve it. Some tasks are harder to solve than others. We can measure the "hardness" of a task with the novelty width and the correlation complexity. In this work, we compare these measures.

Additionally, we introduce the river measure, a new measure that is based on potential heuristics and therefore similar to the correlation complexity but also comparable to the novelty width. We show that the river measure is upper bounded by the correlation complexity and by the novelty width $+1$.

Furthermore, we show that we can convert a planning task with a polynomial blowup of the task size to ensure that a heuristic of dimension 2 exists that gives rise to backtrack-free search.

Introduction

In satisficing planning we search for a solution to a planning task without considering the length or cost of the plan. As the search space of planning tasks is often extremely large, uninformed searches like blind search are infeasible. Therefore, this problem is tackled with informed search such as heuristic search. A heuristic guides the search by providing an estimate of which parts of the search space should be explored earlier. Potential heuristics are a class of heuristics introduced by Pommerening et al. (2015). Their expressiveness depends on their dimension.

Potential heuristics are mostly used in optimal planning. Pommerening et al. (2015), and Pommerening, Helmert, and Bonet (2017) showed that low dimensional potential heuristics with interesting properties for optimal planning can be synthesized in polynomial time. Corrêa and Pommerening (2019) looked at the dimension necessary to express the perfect heuristic on IPC domains. Seipp et al. (2016) introduced correlation complexity, a measure that checks what dimension is necessary for a potential heuristic to give rise to backtrack-free search for satisficing planning. It expresses the degree of inter-relatedness between state variables of a planning task. In other words, how many conjunct facts have to be taken into account for a planning task.

The Novelty Width[1] (Lipovetzky and Geffner 2012) is an alternative measure that looks at the degree of interrelatedness between state variables of a planning task from a different angle. It is based on the *novelty* of a state. The novelty of a state is k if at least k facts must be considered at once to detect that a state was not yet generated in the search. Lipovetzky and Geffner (2012) implemented a novelty based search algorithm with run time exponential only in k. Later they refined the approach (Lipovetzky and Geffner 2014), combined it with heuristics, and showed competitiveness (Lipovetzky and Geffner 2017b). Yet another usage of novelty is to have it as a companion for heuristic search (Lipovetzky and Geffner 2017a; Francès et al. 2018; Corrêa and Seipp 2022).

The two mentioned measures have a core difference. One talks about the search from the initial state to a goal state, the other about walks to a goal state from any state that is reachable and solvable.

To build a bridge for comparison we introduce a new measure, the river measure. It essentially is equal to the correlation complexity but is only interested in the descending walks starting in the initial state. This is less restrictive and more natural. We demonstrate the relationship of the river measure to both the correlation complexity and variants of the novelty width. Furthermore, we introduce a task conversion that reduces the river measure to 2, which highlights an oddity of this hardness measure.

Background

We consider planning tasks in the SAS^+ formalism (Bäckström and Nebel 1995). A **planning task** is a tuple $\Pi = \langle V, I, O, \gamma \rangle$. V is the finite set of **state variables**, each $v \in V$ corresponds to a finite domain $dom(v)$. A **fact** is a tuple of a state variable and an element of its domain $(v \mapsto d)$ with $v \in V, d \in dom(v)$. A **partial assignment** p is a set of facts where each contained fact belongs to a different variable. We denote the set of variables used in p as $vars(p)$. We refer to an individual fact corresponding to a variable v in p with $p[v] = d$ if $(v \mapsto d) \in p$. A **state** is a partial assignment with a fact for each state variable. We say a partial assign-

[1] In the original work it is simply called "width". We call it "novelty width" to emphasize the distinction between the different, independent concept of width by Chen and Giménez (2007).

ment p **agrees** with a state s if $p \subseteq s$. The state I represents the **initial state**.

The set O contains the **operators**. Each operator $o \in O$ consists of two partial assignments: a **precondition** $pre(o)$ and an **effect** $eff(o)$. An operator o is **applicable** in state s if $pre(o) \subseteq s$. Its application results in the **successor state** $s[\![o]\!]$ where $s[\![a]\!][v] = eff(a)[v]$ if $v \in vars(eff(a))$ and $s[\![a]\!][v] = s[v]$ otherwise. The set of successors of a state $succ(s)$ is the set of individual successor states of $s[\![o]\!]$ with any applicable operator $o \in O$.

The partial assignment γ represents the goal of the task. Each state that agrees with γ is a **goal state**. We call a list of states $\pi = (\pi_0, \dots, \pi_n)$ a **walk** of **length** n if π_{i+1} is a successor of π_i for each $i \in \{0, \dots, n-1\}$ and a **plan** if π_n is a goal state. A plan is **optimal** if its length is minimal (optimal planning considers a cost for each operator. In this work we ignore the cost completely. This is equivalent to assuming a cost of 1 for each operator). The state s is **reachable** if a walk from I to s exists and **solvable** if a plan starting in s exists.

A **heuristic** is a function that maps each state of Π to a value in $\mathbb{R} \cup \{\infty\}$. The idea of a heuristic is to be a guidance through the state space. We look at some definitions that describe different escalations of guidance.

Definition 1 (Enforcing)**.** *We say a heuristic h **enforces a state** s' on s if $s' \in succ(s)$ and $h(s) \leq h(\hat{s})$ for each $\hat{s} \in succ(s) \setminus \{s'\}$ and $h(s) > h(s')$.*

*We say a heuristic h **enforces a walk** $\pi = (s_0, \dots, s_n)$ if h enforces s_i on s_{i-1} for each $0 < i \leq n$.*

In other words, h enforces s' on s if s' is the only improving successor of s under h.

Definition 2 (Hinting)**.** *We say a heuristic h **hints at a state** s' on s if $s' \in succ(s)$ and $h(s) > h(s')$.*

*We say a heuristic h **hints at a walk** $\pi = (s_0, \dots, s_n)$ if h hints at s_i on s_{i-1} for each $0 < i \leq n$.*

In other words, h hints at s' on s if s' is an improving successor of s under h.

It is easy to see that each heuristic that enforces a state or walk also hints at it.

A **potential heuristic** (Pommerening et al. 2015) is a heuristic that is computed with a weighted count of the partial assignments that agree with the given state.

$$h^{pot}(s) = \sum_{p \in \mathcal{P}} (w(p) \cdot [p \subseteq s])$$

where $\mathcal{P}$ is the set of all possible partial assignments for the task, $[p \subseteq s]$ is in the Iverson bracket notation, and $w(p)$ is the **weight** for the partial assignment p. In practice, most of the weights are 0. The **dimension** of a potential heuristic is $max_{p \in \mathcal{P}, w(p) \neq 0} |p|$.

Different Measures

We want to compare these measures:

- (Optimistic/Effective) Novelty Width ((o/e)NW)
- Correlation Complexity (CC)
- River Measure (RM)

Algorithm 1: `Novelty Width Search`

Data: planning task $\Pi = \langle V, I, O, \gamma \rangle$, width $k \in \mathbb{N}$
Result: plan π

1 **if** $\gamma \subseteq I$ **then**
2   **return** empty plan
3 **end**
4 *open* := $[I]$
5 *closed* := $\emptyset$
6 insert each $p \subseteq I$ with $|p| = k$ into *closed*
7 **while** *open is not empty* **do**
8   *current* := pop first element of *open*
9   **foreach** *candidate* $\in$ *succ*(*current*) **do**
10   **if** $\gamma \subseteq$ *candidate* **then**
11   **return** extracted plan to *candidate*
12   **else if** $\exists p \subseteq$ *candidate with* $|p| = k, p \notin$ *closed* **then**
13   insert each $p \subseteq$ *candidate* with $|p| = k$ into *closed*
14   append *candidate* to *open*
15   **end**
16   **end**
17 **end**
18 **return fail**

Novelty Width

Novelty Width Search (NWS) (Lipovetzky and Geffner 2012)[2] is a variation of Breadth First Search (BFS). The pseudo-code is shown in Algorithm 1. The difference to BFS is the stricter pruning. NWS prunes not only the states that are duplicates (i.e. the states in the closed set) but all states with a novelty larger than a given parameter k. We call them **approximate duplicates**. This check is done in Line 12. The other difference to BFS is the initialization in Line 6. Instead of inserting the initial state I into the closed set all partial assignments of size k that agree with I are inserted. Besides these two differences, it is the same as BFS. The novelty of a state s is the size of the smallest partial assignment p that agrees with s but no other state that was generated previously in the search. The amount of partial assignments with size k is exponential only in k. Note that the novelty of a state is not a property of the state itself but is dependent on the previous behavior of the search.

The novelty width is defined with the graph $\mathcal{G}^i$. Its nodes are partial assignments of size i.

Definition 3. *For $\Pi = \langle V, I, O, \gamma \rangle$ the graph $\mathcal{G}^i$ is defined inductively by:*

- *the partial assignment p of size $|p| = i$ is a root in $\mathcal{G}^i$ iff p agrees with I.*
- *$\langle p, p' \rangle$ is a directed edge in $\mathcal{G}^i$ iff p is in $\mathcal{G}^i$ and for every optimal plan $\pi = (\pi_0, \dots, \pi_n)$ for $\langle V, I, O, p \rangle$ there is a*

[2]In the original work it is called i-width search, *IW(i)*. They defined it on the STRIPS(Fikes and Nilsson 1971) formalism. This is no issue as the modification to BFS also works for the SAS^+ version.

successor π_{n+1} of π_n such that π followed by π_{n+1} is an optimal plan for $\langle V, I, O, p' \rangle$ and $|p'| = i$.

The partial assignment p' could be only reachable by an operator o with $|pre(o)|$ of arbitrary size if every optimal plan to p leads to a state that agrees with $pre(o)$. We can view this as the combination of (1) state s agrees with p and (2) s is an optimal goal[3] state for $\langle V, I, O, p \rangle$, gives rise to the implicit information that s agrees with $pre(o)$. This implicit information could be too large to be expressed explicitly as a partial assignment in the size, we restrict ourselves to.

Definition 4 (Novelty Width)**.** *Given a planning task $\Pi = \langle V, I, O, \gamma \rangle$. A partial assignment p is of **novelty width** 0 if it is true in I. Otherwise, its novelty width is the smallest w such that $\mathcal{G}^w$ contains p.*

*The **novelty width of a planning task** is the novelty width of γ. Denoted as* $\text{NW}(\Pi)$.

Note that the goal γ is not contained in $\mathcal{G}^k$ for any $k \in \mathbb{N}$ for unsolvable tasks. This implies $\text{NW}(\Pi)$ is unbounded if Π is unsolvable.

If k is set to the Novelty Width of the task NWS is guaranteed to find an optimal plan. However, NWS can find a sub-optimal plan with k less than the novelty width. This can happen if the states of the optimal plan are categorized as an approximate duplicate due to large novelty while all states in a sub-optimal plan provide a lower novelty. Lipovetzky and Geffner (2012) defined the effective novelty width as:

Definition 5 (Effective Novelty Width (eNW))**.** *We call the smallest k for which NWS(k) finds a satisficing plan for Π the **effective novelty width** of* Π, $\text{eNW}(\Pi)$.

It is possible for two different states which are in the same depth of the search tree to be seen as approximate duplicates of each other, this makes the effective novelty width dependent on the tie-breaking strategy.

Definition 6 (Tie-Breaking Strategy)**.** *We call a function a **tie-breaking strategy** for planning task $\Pi = \langle V, I, O, \gamma \rangle$ if it maps any state in Π to a total order of the operators O.*

The tie-breaking strategy is responsible for the order of iteration in Line 9. This influences not only the order of states in the open list (Line 14) but also which states will be pruned (Line 12). With a tie-breaking strategy, NWS is deterministic.

Definition 7 (Optimistic Novelty Width (oNW)**.** *We call the smallest k for which a tie-breaking strategy exists such that NWS(k) finds a satisficing plan for Π the **optimistic novelty width** of* Π, $\text{oNW}(\Pi)$.

By definition $\text{oNW}(\Pi) \leq \text{eNW}(\Pi)$. In the introducing paper, Lipovetzky and Geffner (2012) pointed out that the novelty width is an upper bound of the effective novelty width $\text{eNW}(\Pi) \leq \text{NW}(\Pi)$. Similar to NW, $\text{eNW}(\Pi)$ and $\text{oNW}(\Pi)$ are unbounded if Π is unsolvable.

It is easy to see that each state in a plan found by NWS must contain a novel partial assignment. There are at most $\binom{|V|}{k} \cdot D^k$-many partial assignments of size k, where D is the maximal domain size. This implies that the maximum length of a plan found by NWS with input width k is uppper bounded by $(|V| \cdot D)^k$. Additionally, the time complexity of NWS is $\mathcal{O}((|V| \cdot D)^{2k})$ because each iteration of the inner loop adds at least one partial assignment to the closed set and the checks and insertions from lines 12 and 13 are of time complexity $\mathcal{O}(\binom{|V|}{k}) \subset \mathcal{O}(|V|^k)$ (assuming a constant time for each individual check/insertion).

[3] Meaning this state is the end of at least one optimal plan for the given task.

Correlation Complexity

For local search, one would prefer a state space topology where the only local minima are at goal states, i.e., every non-goal state has a successor with a lower heuristic value (Hoffmann 2001). This implies that greedy search algorithms like greedy best-first search or simple hill-climbing directly reach the goal without back-tracking.

For the correlation complexity, this requirement is relaxed to a subset of the state space namely the alive states.

A state is **alive** if it is reachable and solvable. **Descending** means each (non-goal) alive state has a successor with a lower heuristic value. **Dead-end avoiding** means no alive state has a successor with a lower heuristic value that is unsolvable. The correlation complexity is based on descending and dead-end avoiding (DDA) heuristics.

Definition 8. *Let h be a heuristic for a planning task Π with alive states S_A and goal states S_G. If the following two conditions hold, then h is **DDA** for Π.*

$$\forall s \in (S_\text{A} \setminus S_\text{G})\, \exists t \in \mathit{succ}(s) : h(t) < h(s)$$

$$\forall s \in (S_\text{A} \setminus S_\text{G}) \forall t \in \mathit{succ}(s) : h(t) < h(s) \rightarrow t \in S_\text{A}$$

We see such a heuristic hints, on every alive state at a walk that ends in a goal state. Additionally, all walks hinted at by h on an alive state end in a goal.

Definition 9. *The **correlation complexity** of a task* $\text{CC}(\Pi)$ *is the smallest dimension for a potential heuristic on Π that is DDA.*

Note that every heuristic is trivially DDA on unsolvable tasks because no alive state exists. This implies $\text{CC}(\Pi) = 0$ if Π is unsolvable.

(o/e)NW vs. CC

In the introducing paper of the CC Seipp et al. (2016) compared it to the NW. They explained two scenarios where one measure is arbitrary, and the other is constant. Their arguments also work for the comparison to oNW and eNW. We reiterate that.

The measures (o/e)NW and CC express a notion of how hard a planning task is. One significant difference is that the (o/e)NW only considers the search from the initial state to a goal state, while the CC considers the search from all alive states to a goal state.

This encodes some form of pessimism into the CC. To make this more apparent consider the following modification of a solvable task: Modify the task such that there is a single artificial goal state and an artificial initial state with two successors. One to the original initial state the other is a shortcut transition from the new initial state to the new goal. The artificial goal state is reachable by all original goal states. This modification reduces the (o/e)NW to 1, which matches the intuition that this task is not as hard as the original. Yet, CC stays the same because the set of alive states did not change (except for being projected into a space where the additional variables exist). In contrast, the oNW contains a form of optimism with its optimization over tie-breaking strategies.

Looking at a task that encodes a binary counter gives us the opposite effect. This task contains state variables $v_0, \dots, v_n$ with a domain of $\{0, 1\}$ to represent each bit. In the initial state, each bit is set to 0 and in the goal each is set to 1. Each of the 2^n operators has a representation of one integer as precondition and a representation of the following integer as effect. There the correlation complexity is 1 (simply choose the weight such that the variable representing the ith bit set to 1 is -2^i) while the (o/e)NW is the amount of bits.

River Measure

We want to get rid of this pessimism and concentrate on the search from the initial state to a goal state while staying close to the idea of the CC. Thereby we also remove the dependency on tie-breaking strategies. For that, we introduce a new measure.

For intuition, we want to look at planning tasks with heuristics, analogous[4] to a landscape. We interpret the state as latitude-longitude position and the heuristic as its altitude, a successor state corresponds to a neighboring position. If someone pours water onto one position it would flow to the neighboring positions with lower altitude and moves further to a neighbor with even lower altitude.

For a planning task with a DDA heuristic, we could pour the water onto any alive state and it would flow to a goal position, eventually. However, in satisficing planning we do not need guidance from all alive states to goal states but only one from the initial state. If we are guaranteed that the water flows from the initial position to a goal position, like a river flows from the source to the sea, we are happy. We see the goal states as drains that remove the water from the system.

For this reason, we shift our focus from the alive states to the *wet* states.

Definition 10. *A state s in Π is **wet** under a heuristic h if h hints at a walk π from the initial state and π contains s not after a goal state (the initial state is trivially wet).*

In other words, a state is wet if it is reachable by a descending walk that does not pass the goal.

Definition 11. *Let h be a heuristic for a planning task Π with the wet states S_{W} and goal states S_{G}. If the following two conditions hold, then h is **WDDA** for Π.*

$$\forall s \in (S_{\mathrm{W}} \setminus S_{\mathrm{G}})\ \exists t \in succ(s) : h(t) < h(s)$$

$$\forall s \in (S_{\mathrm{W}} \setminus S_{\mathrm{G}}) \forall t \in succ(s) : h(t) < h(s) \rightarrow t \in S_{\mathrm{W}}$$

The first condition guarantees that each non-goal state in S_{W} has successor with a lower heuristic value and the second guarantees that such a state is part of S_{W}. We see for a WDDA heuristic all walks it hints at on the initial state lead to a goal (not necessarily end in one).

WDDA retains the key characteristics of DDA: for any task, greedy search algorithms using a heuristic satisfying the WDDA property are led towards a goal state without encountering local minima and hence without back-tracking.

With that, we introduce the **river measure** (RM).

Definition 12. *The RM of a task* $\mathrm{RM}(\Pi)$ *is the smallest dimension of a potential heuristic on* Π *that is* wet descending and dead-end avoiding *(WDDA).*

This means each wet state (that is not a goal state) has a successor with a lower heuristic value. Successors of a wet state (that is not a goal state) with a lower heuristic value are wet, too.

Note that no heuristic exists on unsolvable tasks that is WDDA. This implies $\mathrm{RM}(\Pi)$ is unbounded if Π is unsolvable.

To talk about individual states of a task we define the RM of a state.

Definition 13. *The **RM of a state** s in $\Pi = \langle V, I, O, \gamma \rangle$ is* $\mathrm{RM}(s) = \mathrm{RM}(\langle V, s, O, \gamma \rangle)$.

CC vs. RM

The most apparent similarity between the two is the disregard for the difficulty to actually synthesize a potential heuristic with the required property. Additionally, it is easy to see that each DDA heuristic is also WDDA. This implies that $\mathrm{CC}(\Pi) \geq \mathrm{RM}(\Pi)$. Helmert et al. (Helmert et al. 2022) studied multiple modifications of the DDA property. Each of them implies the WDDA property. From this implication and their results, we conclude that the synthesis of a WDDA heuristic is in Σ_2^p, which is the second level of the polynomial hierarchy.

We saw in the comparison to (o/e)NW that even in easy tasks "bad" alive states could cause a large CC. If a single alive state exists that requires d dimension in the potential heuristic to construct a descending walk to the goal from it, then the CC is at least d. Even if that state is far away from the initial state and the goal state. The state could also be reachable only by traversing through goal states. However, focusing on individual states does not reveal the entire picture.

The maximal RM over all alive states is not guaranteed to be the CC of the task, we consider the following planning task as a counterexample. A binary counter that counts from 11 down to 00, but it is first decided if a little endian or big endian counting is used. The state space is depicted in Figure 1.

For all states s with $(v \mapsto \text{little endian})$ the weights $w(b_0 \mapsto 1) = 1$ and $w(b_1 \mapsto 1) = 2$ provide a potential heuristic of dimension 1 that is WDDA from s.

For all other states s' the weights $w(v \mapsto \text{undecided}) = 1$, $w(v \mapsto \text{little endian}) = 2$, $w(b_0 \mapsto 1) = 2$ and $w(b_1 \mapsto 1) = 1$ provide a potential heuristic of dimension 1 that is WDDA from s'.

[4] This mental model is the same as the one used for high watermarks (Wilt and Ruml 2014). However, in this work, we avoid local minima between the initial state and the goal completely instead of filling them up.

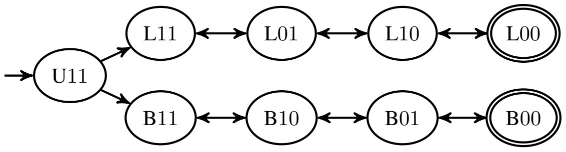

Figure 1: Reachable state space of the little/big endian binary counter. The first symbol inside the node indicates the assignment of v with U for *undecided*, L for *little endian*, and B for *big endian*. The second symbol shows the assignment of bit b_0 and the third of bit b_1.

However, the CC of the task is at least 2. This can be shown with the criterion[5] from Seipp et al. (2016) since the operators $\langle\{v \mapsto \text{big endian}, b_0 \mapsto 1, b_1 \mapsto 0\}, \{b_0 \mapsto 0, b_1 \mapsto 1\}\rangle$ and $\langle\{v \mapsto \text{littel endian}, b_0 \mapsto 0, b_1 \mapsto 1\}, \{b_0 \mapsto 1, b_1 \mapsto 0\}\rangle$ are both critical but inverse of each other.

The CC is equal to the RM if only a single walk π from the initial state to the goal exists. A WDDA heuristic h hints only at π. Since this is the only walk to a goal the states in π are exactly the alive states. We see that the wet states are the alive states and conclude that h is also DDA.

oNW vs. RM

A plan found by NWS with input width k can be enforced by a potential heuristic of dimension $k+1$. For the constructive proof, we use discriminating facts.

Definition 14 (Discriminating Facts). *Let s, s' be states. If $s \neq s'$ then there exists at least one fact $(v \mapsto d') \in s'$ with $(v \mapsto d') \notin s$. We call the set of such facts the **discriminating facts** from s to s'. We denote the set of discriminating facts as $\delta(s, s') := s' \setminus s$.*

Let $\delta^\times(s, s')$ be one arbitrarily chosen element from the discriminating facts $\delta(s, s')$.

Theorem 1. *If NWS with input width k finds a plan π for $\langle V, I, O, \gamma\rangle$ (with any tie-breaking strategy) then a potential heuristic of dimension $k+1$ exists that enforces π on I.*

Proof. We assume $k < |V|$ because a potential heuristic of dimension $|V|$ holds no restriction. We can assign each state a heuristic value individually. It is trivial that an enforcing potential heuristic of dimension $|V|$ exists.

We look at the search tree of an NWS with input width k that found a plan π for a planning task $\Pi = \langle V, I, O, \gamma\rangle$. That plan π traverses through the states $\pi_0, \pi_1, \ldots, \pi_L$ with $\pi_0 = I$. We will construct a potential heuristic h that enforces π on I.

Let p_i^π be one arbitrarily chosen partial assignment that was novel in π_i with $|p_i^\pi| = k$. The idea is to give p_i^π a strong negative weight to ensure $h(\pi_{i-1}) > h(\pi_i)$. However, we have to watch out for successors other than π_i and give them strong positive weights to ensure that h enforces π_i on π_{i-1}.

We can divide $\mathit{succ}(\pi_{i-1})$ into the disjoint sets $\{\pi_i\}, \{\pi_{i-1}\} \cap \mathit{succ}(\pi_{i-1}), S_i, C_i, Z_i$ with:

- $S_i := \{s_i \in \mathit{succ}(\pi_{i-1}) \mid p_i^\pi \subseteq s_i, s_i \neq \pi_i\}$
- $C_i := \{c_i \in \mathit{succ}(\pi_{i-1}) \mid p_i^\pi \nsubseteq c_i, p_{i-1}^\pi \subseteq c_i, c_i \neq \pi_{i-1}\}$
- $Z_i := \{z_i \in \mathit{succ}(\pi_{i-1}) \mid p_i^\pi \nsubseteq z_i, p_{i-1}^\pi \nsubseteq z_i\}$

The set S_i contains the successors that also agree with the novel assignment p_i^π. The set C_i contains the successors that also agree with the assignment p_{i-1}^π that was novel in π_{i-1} and do not agree with the novel assignment p_i^π. The set Z_i contains the successors that do not agree with p_i^π nor with p_{i-1}^π.

Consider the sets of facts:

- $S_i^\delta := \{\delta^\times(\pi_i, s_i) \mid s_i \in S_i\}$
- $C_i^\delta := \{\delta^\times(\pi_{i-1}, c_i) \mid c_i \in C_i\}$

and the set of partial assignments:

- $S_i^p := \{p_i^\pi \dot\cup \{\delta_i^s\} \mid \delta_i^s \in S_i^\delta\}$
- $C_i^p := \{p_{i-1}^\pi \dot\cup \{\delta_i^c\} \mid \delta_i^c \in C_i^\delta\}$

A successor x of π_{i-1} agrees with a fact in S_i^δ, implies $x \neq \pi_i$ and a successor x of π_{i-1} agrees with a fact in C_i^δ, implies $x \neq \pi_{i-1}$.

We now assign weights to the partial assignments in $\{p_i^\pi\}$, S_i^p, and C_i^p.

It could be the case that one $c_i \in C_i$ agrees with more p_j^π with $j < i$ than π_{i-1} does. This would lead to more negative summands in the evaluation of the potential heuristic, which could lead to $h(c_i) < h(\pi_{i-1})$ meaning h would not enforce π_i on π_{i-1}. To compensate for that we give all $p_{i-1}^\pi \cup \{\delta_i^c\}$ in C_i^p a strong positive weight.

Let $n := |V|$ be the number of state variables in the considered task Π and the auxiliary constant $\Omega := \binom{n}{k+1} + 2$. We chose the weights for the potential heuristic with:

- $w(p_i^\pi) := -\Omega^{2\cdot i}$
- $w(s_i^p) := +\Omega^{2\cdot i+1}$ for all $s_i^p \in S_i^p$
- $w(c_i^p) := +\Omega^{2\cdot i-1}$ for all $c_i^p \in C_i^p$

for each $i \in \{0, \ldots, L\}$. For all other partial assignments the weight is 0. We see that the largest partial assignment with a non-zero weight is of size $k+1$ and therefore h is of dimension $k+1$.

For any state x the heuristic h is evaluated with

$$h(x) = \sum_{j=0}^{L} -\Omega^{2j} \cdot [p_j^\pi \subseteq x] + \sum_{j=0}^{L} \sum_{s^p \in S_j^p} \Omega^{2j+1} \cdot [s^p \subseteq x] + \sum_{j=0}^{L} \sum_{c^p \in C_j^p} \Omega^{2j-1} \cdot [c^p \subseteq x]$$

Since p_j^π does not agree with any $x_i \in \mathit{succ}(\pi_{i-1})$ where $j > i$, we can reduce $h(x_i)$ to

$$h(x_i) = \sum_{j=0}^{i} -\Omega^{2j} \cdot [p_j^\pi \subseteq x_i] + \sum_{j=0}^{i} \sum_{s^p \in S_j^p} \Omega^{2j+1} \cdot [s^p \subseteq x_i] + \sum_{j=0}^{i} \sum_{c^p \in C_j^p} \Omega^{2j-1} \cdot [c^p \subseteq x_i]$$

[5]Theorem 5 (Seipp et al. 2016).

For the lower bound of $h(\pi_i)$, we consider the case that π_i agrees with all partial assignments with a negative weight and none of the partial assignments with a positive weight. We conclude the lower bound:

$$\sum_{j=0}^{i} -\Omega^{2j} + 0 + 0 \leq h(\pi_i)$$

For the upper bound, we consider that p_i^π agrees with π_i and there are $\binom{n}{k+1}$ partial assignments of size $k+1$ that agree with a given state. We conclude the upper bound:

$$\begin{aligned} h(\pi_i) &\leq -\Omega^{2i} + \binom{n}{k+1}\Omega^{2\cdot(i-1)+1} + 0 \\ &\leq -(\Omega \cdot \Omega^{2i-1}) + \binom{n}{k+1}\Omega^{2\cdot(i-1)+1} \\ &< -\Omega^{2i-1} \text{ for } \Omega > \binom{n}{k+1} + 1 \\ &\leq \sum_{j=0}^{i-1} -\Omega^{2j} \\ &\leq h(\pi_{i-1}) \end{aligned}$$

and $h(\pi_i) \leq 0$.

For any $s_i \in S_i$ there is at least one $s^p \in S_i^p$ that agrees with s_i. This implies the lower bound:

$$\begin{aligned} h(s_i) &\geq \sum_{j=0}^{i} -\Omega^{2j} + 1 \cdot \Omega^{2i+1} + 0 \\ &\geq -2 \cdot \Omega^{2i} + \Omega^{2i+1} \text{ for } \Omega > 2 \\ &\geq 0 \text{ for } \Omega > 2 \\ &\geq h(\pi_{i-1}) \end{aligned}$$

For any $c_i \in C_i$ there is at least one $c^p \in C_i^p$ that agrees with c_i. The partial assignment p_i^π does not agree with c_i. This implies the lower bound:

$$\begin{aligned} h(c_i) &\geq \sum_{j=0}^{i-1} -\Omega^{2j} + 0 + 1 \cdot \Omega^{2i-1} \\ &\geq -2 \cdot \Omega^{2i-2} + \Omega^{2i-1} \text{ for } \Omega > 2 \\ &\geq 0 \text{ for } \Omega > 2 \\ &\geq h(\pi_{i-1}) \end{aligned}$$

The partial assignments p_i^π and p_{i-1}^π do not agree with any $z_i \in Z_i$. This implies the lower bound:

$$\begin{aligned} h(z_i) &\geq \sum_{j=0}^{i-2} -\Omega^{2j} + 0 + 0 \\ &\geq h(\pi_{i-1}) \end{aligned}$$

We conclude that each successor $x \in \mathit{succ}(\pi_{i-1}) \setminus \{\pi_i\}$ holds $h(\pi_{i-1}) \leq h(x)$ and $h(\pi_{i-1}) > h(\pi_i)$. This implies that h enforces π_i on π_{i-1} for each $i \in \{0, \ldots, L\}$ and therefore h enforces π.

□

Note that NWS can find a plan with k smaller than the novelty width. We remember, NWS finds a satisficing plan with the fitting tie-breaking strategy and $k = \text{oNW}(\Pi)$. In the proof above it is not used that k is the novelty with but only that NWS with input width k found a (satisficing) plan.

This shows the upper bound for the river measure $\text{RM}(\Pi) \leq \text{oNW}(\Pi) + 1 \leq \text{NW}(\Pi) + 1$.

The representation of each weight in above proof requires a number of bits polynomial in $|V|$ and exponential in oNW. To be more precise $\mathcal{O}(\log((|V|^2)^{|V|^{\text{oNW}}})) = \mathcal{O}(|V|^{\text{oNW}} \cdot 2 \cdot \log(|V|))$.

Conversion

We take a look at how we can change the hardness of a planning task according to the measures in a way that contradicts the intuition of hard and easy tasks.

There is a way to convert a planning task to an equivalent task in the sense that the original is solvable iff the converted task is solvable and a plan for one is easily translatable into a plan for the other.

The conversion is linear in the optimal plan length of the original task. The number of operators grows linearly by the optimal plan length. The conversion has an oNW of 1. This implies that the RM of such a converted task is always 2.

We add a step counter to the planning task. The step is counted by each operator. The plan length does not change but the amount of operators increases at least by a factor of the shortest plan length.

Definition 15 (E-Conversion)**.** *Let $\Pi = \langle V, I, O, \gamma \rangle$ be a planning task with optimal plan length L. The **E-conversion** of Π is the planning task $\Pi^E = \langle V^E, I^E, O^E, \gamma \rangle$ with*

$$\begin{aligned} V^E =& V \cup \{t\} \\ O^E =& \{\langle \mathit{pre}(o) \cup \{(t \mapsto i-1)\}, \\ & \mathit{eff}(o) \cup \{(t \mapsto i)\}\rangle | i \in \{1, \ldots, L\}, o \in O\} \\ I^E =& I \cup \{(t \mapsto 0)\} \end{aligned}$$

and a domain of $\{0, \ldots, L\}$ for the new state variable t.

It is easy to see that the E-conversion Π^E is identical to the original task Π except for the counter that updates with every step. A projection to the original state variables returns the original task. We can easily extract a plan for the original task from a plan for its E-conversion.

Theorem 2. *Let $\Pi = \langle V, I, O, \gamma \rangle$ be a planning task.*

$$\text{NW}(\Pi^E) \leq \text{NW}(\Pi).$$

Proof. Looking at the graphs $\mathcal{G}^i(\Pi)$ and $\mathcal{G}^i(\Pi^E)$ we see that $\mathcal{G}^i(\Pi)$ is a subgraph of $\mathcal{G}^i(\Pi^E)$, meaning if γ is a node in $\mathcal{G}^i(\Pi)$, it is one in $\mathcal{G}^i(\Pi^E)$, too but not necessarily vice versa. □

Theorem 3. *Let $\Pi = \langle V, I, O, \gamma \rangle$ be a planning task.*

$$\text{oNW}(\Pi^E) = 1.$$

Proof. With the right tie-breaking strategy and input width 1, NWS iterates through the foreach loop in a way that in each depth it first adds the state π_i of a plan into the open list. This is possible as $\{(t \mapsto i)\} \subseteq \pi_i$ is a partial assignment of size 1 that is not in the closed set, yet. □

Theorem 4. *A plan enforcing, potential heuristic of dimension* 2 *exists for the E-conversion* Π^E *of any planning task* Π.

Proof. Since $\{(v'_i \mapsto 1)\}$ is a novel partial assignment in each state π_i^E the optimistic novelty width is 1. We conclude with Theorem 1 that a plan enforcing potential heuristic of dimension 2 exists. □

The $\mathrm{eNW}(\Pi^E)$ is in between $\mathrm{oNW}(\Pi^E) = 1$ and $\mathrm{NW}(\Pi^E) \leq \mathrm{NW}(\Pi)$.

The time complexity of running NWS with input width k on task Π^E is $\mathcal{O}((|V| \cdot D + L)^k)$ where L is the length of a plan for Π and $L \leq (|V| \cdot D)^{\mathrm{oNW}(\Pi)}$. Resulting in a time complexity of $\mathcal{O}((|V| \cdot D)^{k \cdot \mathrm{oNW}(\Pi)})$. That means if the difference between $\mathrm{NW}(\Pi)$ and $\mathrm{oNW}(\Pi)$ is large, then running NWS on Π could in theory be beneficial. The question is if such tasks do exist.

This shows that the measures eNW, oNW, and RM are very dependent on the representation of the task.

Discussion

We see that adding state variables can reduce the hardness of a task. This might contradict one's intuition that tasks of larger size tend to be harder. After all the state space increases.

However, this conversion is not meant to help the planner but the pair of heuristic and algorithm. The assignment of the additional state variable in a given state implies together with the behavior of the algorithm and heuristic further, compact information about the walk to said state.

We can view this as the combination of s agrees with $p \supseteq \{(t \mapsto x)\}$ ($x \in \mathit{dom}(t)$) and s is reached by a descending walk under h gives rise to further implicit information about s. Similarly to the edges in the $\mathcal{G}^i$ graph, such implicit information could be too large to be expressed explicit as a partial assignment in the size, we restrict ourselves to.

Without these, it would have to rely on features of greater size to extract the necessary information to provide the guidance we ask for.

For the NWS which works without a heuristic, this argument can not be made. Even though the oNW is reduced by the conversion the measure becomes utterly useless. The idea was to have a modified duplicate check that prunes more states than BFS. With the conversion, this aspect is lost. States in the converted task are never seen as approximate duplicates of each other unless they would be true duplicates in the original task and reached by a walk of equal length. This happens because the added state variables provide a novel fact after each step interacting with the counter.

Adding unnecessary state variables to reduce the hardness can be viewed the other way around. Removing unnecessary state variables might increase the RM. We look at the Visit-All domain from the international planning competition (IPC 2011). The task we consider has an agent that starts in the middle of a 3 by 3 grid. The goal is to visit the corner locations 0-0, 0-2, 2-0, and 2-2. A potential heuristic of dimension 1 with the weight $w(\{\texttt{location-x-y-is-visited} \mapsto \mathit{true}\}) = -1$ for each location would provide a WDDA heuristic. However, planners performing a goal-based relevance analysis remove the `location-x-y-is-visited` variables if they are not goal variables. This includes all planners that use the preprocessing step of FF (Hoffmann and Nebel 2001) or the translation step of Fast Downward (Helmert 2009). Removing these seemingly irrelevant variables causes an RM of 2, while the original task has an RM of 1.

In any plan the agent has to move onto one of the edge locations (0-1, 1-0, 1-2, or 2-1) in the first step. For a WDDA heuristic, the weight for this location must be lower than for the middle one because no fact changes besides these two. The agent has to get to at least a second edge location to solve the task. This second edge location has to have a better heuristic value than the first one for a descending heuristic. However, all edge locations are reachable in the initial state. The agent could get to the edge location with the lowest weight first and end up being stuck in a local minimum.

The oNW is not affected in this example. Each state of the optimal plan provides a novel fact about the agent's position. Therefore the oNW is 1. The CC is also not affected because the state with the agent on the middle location and all edge locations visited is an alive state. Therefore the CC is 2.

Conclusion

We introduced the river measure as a new measure for the "hardness" of planning tasks. We compared it to the existing measures and showed that it is upper bounded by the correlation complexity and by the effective novelty width +1. $\mathrm{RM}(\Pi) \leq \mathrm{CC}(\Pi)$ and $\mathrm{RM}(\Pi) \leq \mathrm{oNW}(\Pi) + 1 \leq \mathrm{eNW}(\Pi) + 1 \leq \mathrm{NW}(\Pi) + 1$. Additionally, we explained a conversion for arbitrary planning tasks with a growth of task description size linear to the size of the original plan that reduces the river measure to 2.

This indicates the expressive power of potential heuristics with bounded dimension and the value seemingly unnecessary state variables can provide.

Acknowledgements

We have received funding for this work from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement no. 817639). Moreover, this research was partially supported by TAILOR, a project funded by the EU Horizon 2020 research and innovation programme under grant agreement no. 952215.

References

Bäckström, C.; and Nebel, B. 1995. Complexity Results for SAS$^+$ Planning. *Computational Intelligence*, 11(4): 625–655.

Chen, H.; and Giménez, O. 2007. Act Local, Think Global: Width Notions for Tractable Planning. In *Proc. ICAPS 2007*, 73–80.

Corrêa, A. B.; and Pommerening, F. 2019. An Empirical Study of Perfect Potential Heuristics. In *Proc. ICAPS 2019*, 114–118.

Corrêa, A. B.; and Seipp, J. 2022. Best-First Width Search for Lifted Classical Planning. In *Proc. ICAPS 2022*, 11–15.

Fikes, R. E.; and Nilsson, N. J. 1971. STRIPS: A New Approach to the Application of Theorem Proving to Problem Solving. *AIJ*, 2: 189–208.

Francès, G.; Geffner, H.; Lipovetzky, N.; and Ramiréz, M. 2018. Best-First Width Search in the IPC 2018: Complete, Simulated, and Polynomial Variants. In *IPC-9 Planner Abstracts*, 23–27.

Helmert, M. 2009. Concise Finite-Domain Representations for PDDL Planning Tasks. *AIJ*, 173: 503–535.

Helmert, M.; Sievers, S.; Rovner, A.; and Corrêa, A. B. 2022. On the Complexity of Heuristic Synthesis for Satisficing Classical Planning: Potential Heuristics and Beyond. In *Proc. ICAPS 2022*, 124–133.

Hoffmann, J. 2001. Local Search Topology in Planning Benchmarks: An Empirical Analysis. In *Proc. IJCAI 2001*, 453–458.

Hoffmann, J.; and Nebel, B. 2001. The FF Planning System: Fast Plan Generation Through Heuristic Search. *JAIR*, 14: 253–302.

Lipovetzky, N.; and Geffner, H. 2012. Width and Serialization of Classical Planning Problems. In *Proc. ECAI 2012*, 540–545.

Lipovetzky, N.; and Geffner, H. 2014. Width-based Algorithms for Classical Planning: New Results. In *Proc. ECAI 2014*, 1059–1060.

Lipovetzky, N.; and Geffner, H. 2017a. Best-First Width Search: Exploration and Exploitation in Classical Planning. In *Proc. AAAI 2017*, 3590–3596.

Lipovetzky, N.; and Geffner, H. 2017b. A polynomial planning algorithm that beats LAMA and FF. In *Proc. ICAPS 2017*, 195–199.

Pommerening, F.; Helmert, M.; and Bonet, B. 2017. Higher-Dimensional Potential Heuristics for Optimal Classical Planning. In *Proc. AAAI 2017*, 3636–3643.

Pommerening, F.; Helmert, M.; Röger, G.; and Seipp, J. 2015. From Non-Negative to General Operator Cost Partitioning. In *Proc. AAAI 2015*, 3335–3341.

Seipp, J.; Pommerening, F.; Röger, G.; and Helmert, M. 2016. Correlation Complexity of Classical Planning Domains. In *Proc. IJCAI 2016*, 3242–3250.

Wilt, C.; and Ruml, W. 2014. Speedy versus Greedy Search. In *Proc. SoCS 2014*, 184–192.