

Approximate Novelty Search

¹ Anubhav Singh, ¹ Nir Lipovetzky, ² Miquel Ramirez, ³ Javier Segovia-Aguas

¹ School of Computing and Information Systems, University of Melbourne, Australia

² Electrical and Electronic Engineering, University of Melbourne, Australia

³ Dept. Information and Communication Technologies, Universitat Pompeu Fabra, Spain

anubhavs@student.unimelb.edu.au, {nir.lipovetzky, miquel.ramirez}@unimelb.edu.au, javier.segovia@upf.edu

Abstract

Width-based search algorithms seek plans by prioritizing states according to a suitably defined measure of novelty, that maps states into a set of novelty categories. Space and time complexity to evaluate state novelty is known to be exponential on the cardinality of the set. We present novel methods to obtain polynomial approximations of novelty and width-based search. First, we approximate novelty computation via random sampling and Bloom filters, reducing the runtime and memory footprint. Second, we approximate the best-first search using an adaptive policy that decides whether to forgo the expansion of nodes in the open list. These two techniques are integrated into existing width-based algorithms, resulting in new planners that perform significantly better than other state-of-the-art planners over benchmarks from the International Planning Competitions.

1 Introduction

Autonomous systems operating on the edge of computer networks or that only have occasional, sporadic access to vast, centralized computing resources require decision making algorithms that work under those conditions. Not only low response times are required to seek courses of action that are safe and effective, but also the memory available for such computation is limited. The need to adapt existing heuristic search algorithms such as A* to deal with time and space restrictions was recognized early on (Chakrabati et al. 1989) and followed-up recently (Vadlamudi, Aine, and Chakrabarti 2011; Dionne, Thayer, and Ruml 2012).

In this paper we look at *width-based search methods* (Lipovetzky and Geffner 2012), a family of algorithms that rely on heuristics that measure the *novelty* of a state, comparing its information content with that of states visited in the past. Originally developed in the context of classical planning (Geffner and Bonet 2013), when combined with other heuristics (Bonet and Geffner 2001; Hoffmann, Porteous, and Sebastia 2004; Lipovetzky and Geffner 2017a; Katz et al. 2017), width-based planners become state-of-the-art and competitive with portfolio solvers (IPC18). A major shortcoming of the latter width-based methods derived from best-first search (BFS) (Edelkamp and Schrödl 2012; Pearl 1983), such as those used by the planner DUAL-BFWS, is

that measuring novelty is exponential on the number of discrete levels or categories used to rank states. Lipovetzky and Geffner (2012) showed that an upper bound exists for any given classical planning instance, yet this result cannot be exploited, as this bound results in impractical runtimes to evaluate states, a crucial issue for the effectiveness of heuristic search methods.

We address this issue by proposing new methods to obtain polynomial approximations of novelty and to control the growth of the memory footprint of BFWS algorithms. The first contribution is an appraisal approximation of the novelty of state information by randomly sampling the space of possible valuations of state variables, and using Bloom filters for efficient, but imprecise, storage of state information. The second contribution is a novel form of best-first search which uses an adaptive policy that decides whether to delay the generation of successor states. This policy is derived from the analytical solution to an infinite-horizon Markov Decision Problem (MDP) (Bertsekas 2017), where its cost function controls the representation of different novelty categories in the open list.

The paper is structured as follows. We start discussing concisely background material covering classical planning, width-based search, and Bloom filters. Sections 3 and 4 expound the contributions of this paper, approximations of novelty measurements and search. We evaluate both over every benchmark in the IPC satisficing track. We finalize with a discussion of the importance and potential impact of our results.

2 Background

We follow Bonet & Geffner (2013) presentation of classical planning, an optimal control problem (Bertsekas 2017) where states are fully observable and made of a finite number of Boolean atoms, actions are finite and optimal solutions (plans) map the unique initial state into one of many goal states, in the minimum number of steps. The transition system for a classical planning problem P is defined as $S(P) = \langle S, s_0, S_G, A, T \rangle$ where S is the state space, $s_0 \in S$ is the initial state, $S_G \subseteq S$ is the set of goal states, A is the set of actions and $T : S \times A \rightarrow S$ is the transition function. We define the set of *applicable* actions $A(s)$ as the subset of $a \in A$, for which $T(s, a)$ is defined. A solution to $S(P)$ is a plan $\pi = \langle a_1, \dots, a_n \rangle$, with length n , that maps an initial

state $s_0 \in S$ to a goal state $s_n \in S_G$, such that each action is applicable in the corresponding state $a_i \in A(s_{i-1})$, along the induced sequence $s_0, \dots, s_i, \dots, s_n$ where $s_i = T(s_{i-1}, a_i)$.

We describe classical planning problems compactly using STRIPS (Nilsson and Fikes 1971). P is given as the tuple $P = \langle F, O, I, G \rangle$, where F is the collection of Boolean atoms or fluents, O is the set of operators, $I \subseteq F$ is a set of atoms that fully describe the initial state, and $G \subseteq F$ is a set of atoms in the goal state. The transition system $S(P)$ is obtained from P as $S = 2^{|F|}$, $s_0 = I$, $s_G = \{s \mid s \models G\}$, T and $A(s)$ follow from the effects and preconditions of $a \in O$. Optimal solutions are those sequences of actions that are the shortest too. In this paper we consider *satisficing* solutions, trading off between run time and quality.

The planners discussed in this paper are instances of the classic heuristic search algorithm named Best-First Search (Edelkamp and Schrödl 2012; Pearl 1983) or BFS for short. This algorithm searches for plans by extending incrementally all paths (nodes) in $S(P)$ starting from s_0 . The nodes are visited in the order specified by an *evaluation function* f defined over paths, and the algorithm terminates in the first path that ends on a state $s \in S_G$. BFS implicitly *enumerates* the states s in $S(P)$ by assigning to them a natural number, the *expansion order*, $e(s)$ (Dionne, Thayer, and Ruml 2012), which is 0 for s_0 and increases by one unit for every new state s extending an existing path. We denote the set of states generated before s as $\mathcal{P}(s) = \{s' \mid e(s') < e(s)\}$. As we will see in the next section, this ordering is crucial to define *novelty measures*.

2.1 Best-First Width Search (BFWS)

BFWS (Lipovetzky and Geffner 2017a) is a family of BFS algorithms where the evaluation function for a node n , $f(n)$ is defined as a tuple of functions

$$f(n) = (w, h_1, \dots, h_m)$$

where $w : S \rightarrow \mathcal{W}$ is the function measuring novelty, that maps states $s \in S$ into categories $\omega \in \mathcal{W}$, $\mathcal{W} \subset \mathbb{N}$, and $H = \{h_1, \dots, h_m\}$ is a set of suitably chosen functions. When inserting nodes n in the open list, BFWS algorithms sort in increasing order according to the first function in $f(n)$, breaking ties recursively with the provided h_i . These functions can also be used to partition the set of states generated before s as $\mathcal{P}(s, H) = \{s' \mid e(s') < e(s), h(s) = h(s') \forall h \in H\}$.

Definition 1. The novelty $w(s) = w_{(H)}(s)$ of a newly generated state s given a set of partition functions H over states $s \in S$ is k , iff (1) exists a tuple¹ $t \subseteq F$ of minimum size k , s.t. $s \models t$, (2) $\forall s' \in \mathcal{P}(s, H)$, $s' \not\models t$.

As noted in the introduction, BFWS algorithms are state-of-the-art over the IPC classical planning benchmarks. We illustrate the BFWS framework by discussing in detail BFWS(f_5), one of the best performing to-date in the Agile track (IPC 2018).

The evaluation function $f_5 = \langle w, \#g \rangle$ makes BFWS to expand first novel states, breaking ties with a simple goal counting heuristic (Nilsson and Fikes 1971) $\#g(s)$. The

novelty function uses two heuristic functions to partition the novelty space $w = w_{\langle \#g, \#r \rangle}$, one is $\#g(s)$, and the other, $\#r(s)$, counts the atoms p achieved along the path to s , such that $p \in R$, $R \subseteq F$, where R is selected by a *relevance analysis* procedure. R is meant to contain atoms which are *instrumental* to reach the goal efficiently, so for domain-independent planning, one can instance R as a set of landmarks (Hoffmann, Porteous, and Sebastia 2004), or the set of fluents which belong to positive effects of actions in the relaxed plan for s_0 (Hoffmann and Nebel 2001). Both of the above definitions of R were used in the planner DUAL-BFWS, and the later was used in BFWS(f_5).

Evaluating $w(s)$ requires to test states s to belong to the categories $\omega \in \mathcal{W}$. Lipovetzky and Geffner (2017a) define \mathcal{W} as the integer interval $[1, i + 1]$ where i is the size of the largest novel tuple generated by optimal plans for P . In this case, the test above requires to generate exhaustively all tuples t of size i present in state s and determine if they are present in $\mathcal{P}(s, H)$. An optimal procedure to implement Definition 1 follows. For each tuple of size $1 \leq l \leq i$, let $\beta_l(s) = \{t \mid t \subseteq F, s \models t, |t| = l\}$ ². Starting with $l = 1$, we enumerate tuples $t \in \beta_l(s)$, and then test if t is part of previously observed tuples in the set $\mathcal{N}(s) = \bigcup_{s' \in \mathcal{P}(s, H)} \beta_l(s')$.

If the test is negative for at least one tuple, then $w(s) := l$, otherwise, we need to test the elements of β_{l+1} , until $l = i$. If the test is positive for all $t \in \beta_i$, then the state is considered not to be novel, and $w(s) := i + 1$. This leads to the exponential time and space $O(|F|^i)$ requirements to evaluate $w(s)$ that we outlined in Section 1. While i is known to be generally way smaller than $|F|$, the bound i is usually high enough to render the novelty test up to i impractical. In the case of the IPC benchmarks, any value of $i > 2$ leads to very high runtimes to generate states. In practice, $w(s)$ is approximated by setting i to an arbitrary lower bound, which renders the evaluation of $w(s)$ to be tractable but may relegate states with valuable information to the back of the open list.

2.2 Bloom Filters

The *Bloom filters* (Bloom 1970; Louridas 2017) are a probabilistic data structure to represent sets efficiently, at the expense of allowing false positives when testing whether the set contains a given object. Typical implementations of Bloom filters consist of a bit-array v of size r , where all entries v_j are initially set to \perp , and K independent hash functions η that map objects into the range $[1, r]$. To add an object o as a member of the set represented by the Bloom filter, the K hash functions η_l are evaluated on o , so $v_{\eta_l(o)} := \top$ for $l = 1, \dots, K$. To test whether o is in the set, the hash functions η_l are evaluated, and if all $v_{\eta_l(o)} := \top$, then o is considered to be an element of the set. In comparison with a traditional hash table whose size grows with that of the range of possible objects, the Bloom filter has fixed-size r .

The choice of value for r determines the probability to obtain a *false positive*, that is, testing o for containment and getting a positive answer, when o has not been previously

¹Conjunction of atoms.

²We note that β_l can be iterated by lazily generating its elements.

added as a member. As noted in (Broder and Mitzenmacher 2004) the probability of a false positive is given by

$$P_f = \left(1 - e^{-\frac{Kq}{r}}\right)^K \quad (1)$$

where q is the *expected* number of different objects to be tested. The analytical solution to the problem of minimization of false positive rate with respect to K shows that P_f is minimized when $K = (r/q) \ln 2$ (Broder and Mitzenmacher 2004). Since the expected number of different objects, nodes in planning, tend to be larger than the memory, it follows that when $(r \ln 2) \leq q$, then $K = 1$ minimizes P_f .

3 Novelty Approximation

In this section, we describe an approximate measure of novelty for newly generated states, $\hat{w}(s)$, which is tractable and can be proved to be equal to $w(s)$ with positive probability. For that, Definition 1 is relaxed as follows

Definition 2. *The approximate novelty $\hat{w}(s) = \hat{w}_{(H)}(s)$ of a newly generated state s given a set of partition functions H over states $s \in S$ is k , iff (1) exists a tuple $t \subseteq F$ of minimum size k , s.t. $t \in Z_k(s)$, (2) $\forall s' \in \mathcal{P}(s, H)$, $\mathcal{O}(s', t) = \perp$.*

We have changed Definition 1 in two ways. First, we only test the tuples from a randomly sampled set $Z_l(s) \subseteq \beta_l(s)$, for $l = 1, \dots, i$. We require that the probability of every tuple in $\beta_l(s)$ being selected is uniformly distributed. For that we sample *without replacement* from the *discrete uniform distribution* over $\beta_l(s)$, with probability mass function $p_X(t) = z/|\beta_l(s)|$ for all $t \in \beta_l(s)$, $z = |Z_l(s)|$, representing the probability of occurrence of t in $Z_l(s)$. We note that tuples are sampled, independently, for each state. Second, we replace the condition $s' \not\models t$ for a random variable $\mathcal{O}(s', t)$ that models the runtime behavior of a Bloom filter. $\mathcal{O}(s', t)$ maps pairs s', t to \top with probability 1 when $s' \models t$ and $t \in Z_l(s')$, otherwise, it maps s', t to \perp with positive probability. The tractability of $\hat{w}(s)$ follows from requiring r , the number of entries in the Bloom filter, and $|Z_l(s)| \leq \bar{Z}$, where \bar{Z} is the maximum size of $Z_l(s)$, to be constants, e.g. $\bar{Z} = r = |F|$. It is trivial to note that the running time of any reasonable algorithm for computing $\hat{w}_{(H)}(s)$ as per Definition 2 is $O(i\bar{Z})$ and memory requirement is $O(ir)$, dropping the complexity of tuple membership checks from exponential to linear on i .

These two simple changes suffice to allow measures of novelty that are much finer than what can be obtained with highly optimized implementations of $w(s)$, but certainly, and as noted at the beginning of the section, there is a certain probability that $\hat{w}(s)$ and $w(s)$ will not be in agreement. The rest of this section is devoted to provide a probabilistic model of the rate at which approximate and actual novelty disagree.

Impact of Sampling. We proceed now to derive the probability of error induced by sampling from $\beta_l(s)$ following the discrete uniform distribution. By *error* we refer to the event of $w(s) \neq \hat{w}(s)$, that is $\hat{w}(s)$ is greater or less than $w(s)$ for a state s . We start by defining the probability γ_t , given s and

$\mathcal{P}(s, H)$, of a particular tuple $t \in \beta_l(s)$ observed as *new* in s , that is $t \notin \bigcup_{s' \in \mathcal{P}(s, H)} Z_l(s')$ and $t \in Z_l(s)$, as

$$\gamma_t = \left(\prod_{s' \in \mathcal{P}'_t(s)} \left(1 - \frac{z}{|\beta_l(s')|}\right) \right) \frac{z}{|\beta_l(s)|} \quad (2)$$

where $\mathcal{P}'_t(s)$ is defined as the set $\{s' \mid s' \in \mathcal{P}(s, H), s' \models t\}$, z is the sample size, $z/|\beta_l(s)|$ follows from the probability mass function p_X . The event of taking a sample at s is *independent* from that of sample at $s' \in \mathcal{P}(s, H)$ which allows us to use the *product rule*. Also, it follows that as $z \rightarrow |\beta_l(s)|$, the probability $\gamma_t \rightarrow 0$ when $\mathcal{P}'_t(s) \neq \emptyset$, and $\gamma_t \rightarrow 1$ when $\mathcal{P}'_t(s) = \emptyset$. That is, if all tuples $t \in \beta_l(s')$ are sampled in each $s' \in S$, as we do when computing novelty exactly, the probability of t being *new* in s is 0, if $t \in \bigcup_{s' \in \mathcal{P}(s, H)} Z_l(s')$, and 1 otherwise.

From Equation 2, we follow that the probability of tuple t not being *new* is $[1 - \gamma_t]$. We use this result to compute the probability that none of the tuples $t \in \beta_l(s)$ are *new*, assuming *independence* between different tuples to make the derivation tractable, as

$$p_l = \prod_{t \in \beta_l(s)} [1 - \gamma_t] \quad (3)$$

Using Equation 3, we can now define the probability of approximate novelty measure to be greater or smaller than actual, $P_H = P(\hat{w}(s) > w(s))$, $P_L = P(\hat{w}(s) < w(s))$ respectively, as

$$P_H = \prod_{i=1}^{w(s)} p_i, \quad P_L = \left(1 - \prod_{i=1}^{w(s)-1} p_i\right) \quad (4)$$

Finally, the probability of approximate and actual novelty measures to agree, $P_C = P(\hat{w}(s) = w(s))$, is :

$$P_C = \left(\prod_{i=1}^{w(s)-1} p_i \right) (1 - p_{w(s)}) \quad (5)$$

In Lemma 1 we prove the sum of probabilities in Eqs. 4 and 5 to be 1, and where $(P_L + P_H)$ is the total probability of error induced by sampling.

Lemma 1. *The sum of probabilities P_L , P_H and P_C is 1.*

Proof. Let $q = \prod_{i=1}^{w(s)-1} p_i$ be the probability of not finding new tuples with novelty below $w(s)$. Then, Equations are rewritten as $P_L = 1 - q$, $P_H = q p_{w(s)}$ and $P_C = q(1 - p_{w(s)})$ for a given novelty $w(s)$. Hence, $P_L + P_H + P_C = (1 - q) + q p_{w(s)} + q(1 - p_{w(s)}) = 1 - q + q p_{w(s)} + q - q p_{w(s)} = 1$, proving the correctness of equations. \square

Synergies between Sampling and Bloom Filters. The number of generated states $\mathcal{P}(s)$ before state s can be $O(b^d)$, that is, exponential on the branching factor b of the transition system $S(P)$ and the length d of the path to s from s_0 .

Therefore, replacing $\mathcal{N}(s) = \bigcup_{s' \in \mathcal{P}(s, H)} \beta_l(s')$ by a Bloom filter with r entries cannot come for free. While the Bloom filter will always give the correct answer to membership queries for tuples t that are in $\mathcal{N}(s)$, it can produce false positives for membership, as it incorrectly gives a positive answer for tuples $t' \notin \mathcal{N}(s)$.

We note that sampling from β_l enables the use of Bloom filters to “approximate” $\mathcal{N}(s)$. This is because it leads to a reduction of the probability of false positives P_f given in Equation 1, in comparison with what we would obtain from using β_l directly as in the algorithm for $w(s)$ given in Section 2.1. This observation follows from noting that q in Equation 1 is the expected number of distinct tuples t sampled during the search, and the rate of growth of this random variable is directly proportional to \bar{Z} . The smaller \bar{Z} is, the slower q will grow. From Equation 1, we can see that the probability of false positives P_f increases with the ratio q/r , which in turn depends only on q as r is a constant. Therefore, the rate of growth of P_f depends on \bar{Z} .

Finally, we note that P_f will be maximized when q is exponential on l , the maximum size of the tuples considered. Then, in principle, false positive probability increases too as l grows larger.

Total probability of error. To obtain the total probability of erroneously appraising the novelty of a newly generated state s , P_{error} , we incorporate Equation 1 into Equation 2

$$\gamma_t = \left[\left(\prod_{s' \in \mathcal{P}'(s)} \left(1 - \frac{z}{|\beta_l(s')|} \right) \right) \frac{z}{|\beta_l(s)|} (1 - P_f) \right] \quad (6)$$

used to evaluate Equation 4, from which then it follows that $P_{error} = P_L + P_H$.

Conjoining BFWS(f_5) and Novelty Approximation. Novelty approximation using sampling and Bloom Filter can be directly applied to BFWS. We replace the $w(s)$ in BFWS(f_5) with the approximation $\hat{w}(s)$ resulting in $\hat{f}_5 = \langle \hat{w}, \#g \rangle$.

Additionally, any reasonable implementation to compute $\hat{w}(s) = \hat{w}_{\langle \#g, \#r \rangle}(s)$ for BFWS(\hat{f}_5), as per Definition 2, needs to track the evaluations of partition functions $H = \{\#g, \#r\}$ for all observed tuples. This increases the space complexity by a factor of number of possible partitions, $|G| \cdot |F|$. We manage the increase in space complexity by employing a set of Bloom filters, a *bank V*, and then bounding the space available for novelty computation by a parameter D_{max} . In case D_{max} is sufficiently large to track all the tuples and evaluations of partition functions, we enable exact item membership tests. Otherwise, we use the *bank* of Bloom filters. Whenever a new partitioned space is observed, we assign it a Bloom filter from V . If the number of observed partitions exceeds $|V|$, we overlap them randomly, allowing different partitions to use the same Bloom filter. This results in a gradual decrease in the accuracy of novelty computation in exchange for space, the P_{error} increases as more partitions overlap.

The resulting planner BFWS(\hat{f}_5) has the following hyper-parameters, namely, the sample size \bar{Z} , the size of a Bloom filter r and the bound of space D_{max} . In Section 5, we present experimental evaluations with different choices of parameter values.

Increasing the novelty bound. As discussed in Section 2.1, any implementation of Definition 1 has a complexity of $O(|F|^i)$ rendering the computation impractical for many instances. Whereas, Definition 2 has linear complexity allowing us to compute $\hat{w}(s)$ for any value of $i \in [1, |F|]$. In the following section, we describe the impact of increasing $\mathcal{W} \in [1, i + 1]$ in the *polynomial* planners: BFWS(f_5) with novelty pruning (Lipovetzky and Geffner 2017b).

Theorem 1. *Let $P = \langle F, O, I, G \rangle$ be a STRIPS planning problem. The number of nodes generated for each novelty category $\omega \in \mathcal{W}$, when run with P as input, is less than or equal to $\binom{|F|}{\omega} \cdot |G| \cdot |F|$.*

Proof. The upper bound on the number of observed state partitions is given by $|G| \cdot |F|$. Also, the count of tuples of size ω , of atoms in F , is $\binom{|F|}{\omega}$. Hence, the number of nodes with $w(s) = \omega$ cannot exceed $\binom{|F|}{\omega} \cdot |G| \cdot |F|$. \square

Corollary 1. *Let $P = \langle F, O, I, G \rangle$ be a STRIPS planning problem, and BFWS(\hat{f}_5) the polynomial planner using Bloom filters introduced above. When run with P as its input, BFWS(\hat{f}_5) generates at most $|V| \cdot r$ nodes for each novelty category $\omega \in \mathcal{W}$.*

Proof. A Bloom filter represents $\binom{|F|}{\omega}$ tuples, but the number of true negatives is bound by the size of Bloom filter r . Also, we create a set of Bloom filters, the *bank V*, that represents the set of partitioned spaces of cardinality $|G| \cdot |F|$. Hence, the above bound holds. \square

From Theorem 1, we note that the bound on number of nodes with $w(s) = \omega + 1$ increases by $O(|F|)$ in comparison to those with $w(s) = \omega$, which makes nodes with large value of $\hat{w}(s)$ unlikely candidates for expansion. This leads us to another important issue afflicting BFWS algorithms, inherited from BFS, i.e. only a small fraction of the nodes that make it into the open list are ever considered for expansion. One possible method to address this challenge is to choose small size for r and V , and from Corollary 1 we can deduce that it will bound the nodes in each novelty category by $|V| \cdot r$, hence there are less nodes in each category. This makes it seem that nodes with high novelty $w(s)$ are now more likely to expand, however, in practice nodes receive a higher approximate novelty value $\hat{w}(s)$ because of increase in P_f . In the following section, we discuss a method that remedies this.

4 Best First Search with Open List Control

We propose an example of a novel methodology to design BFS algorithms that aim at controlling the rate of growth of nodes of each category $\omega \in \mathcal{W}$ in the open list. To do

this, we model the search as a discrete-time dynamical system subject to perturbation, and an optimal control problem (Bertsekas 2017) is formulated where optimal policies ensure that a rate of growth less than the branching factor b of $S(P)$ is sustained. With some simplifying assumptions, the optimal policy for this control problem can be derived analytically, as shown below, and integrated directly into the search algorithm. We model the evolution over time of the internal state (i.e. size of open lists) of a BFWS-like algorithm B , subject to function T_B , abstracting the instructions executed in one iteration of the expansion loop of B , as the dynamical system

$$x_{k+1} = T_B(x_k, u_k, c_k)$$

where, k is the index of current expansion, x_k is a suitably defined abstraction of the internal state of the search algorithm B at time k , $u_k \in [0, 1]^{|W|}$ is the *control action*, that prescribes the *pruning rate* for states with $w(s) = \omega$, and $c_k \in \mathbb{N}^{|W|}$ is the count of successor states s' at time k with $w(s') = \omega$. c_k is the *perturbation*, that is, an uncontrollable side-effect of node expansion that has been modeled as a uniformly distributed discrete random variable.

The information we track in x_k is given by the tuple $\langle n_e, n_v(\omega) \rangle$, where n_e is the number of expanded nodes so far, and $n_v(\omega)$ is the count of *novel states* visited for each novelty category $\omega \in W$. If $u = 0^{|W|}$, then T_B is deterministic and B behaves like a standard BFWS algorithm. Otherwise, the successors s' of state s pointed at by node n in the open list with $\min f(n)$ are generated with probability $1 - u^\omega$ when $w(s') = \omega$. If some s' is pruned, n is kept in a *holding queue* and re-expanded whenever the open list becomes empty. An important implementation detail for B is that it needs to maintain $|W|$ open lists in parallel, as keeping smaller open lists is often more performant (Burns et al. 2012), also greatly facilitating implementation and computation of states x_k .

In order to formulate an optimal control problem, we need first to specify a cost function. Considering a very large number of stages or expansions, we can reasonably assume that the horizon is infinite, and define the average cost per stage function (Bertsekas 2017), with a policy $\pi = \{\mu_0, \mu_1, \dots\}$, of the form

$$J_\pi = \lim_{N \rightarrow \infty} \frac{1}{N} \mathbf{E}_{c_k} \left[\sum_{k=0}^{N-1} g(x_k, \mu_k(x_k), c_k) \right] \quad (7)$$

The choice of the cost per stage $g(x_k, \mu_k(x_k), c_k)$ is dictated by the need to seek a trade-off between the number of states in the open list with $\hat{w}(s) = \omega$ growing too large and missing out useful novel states. We define g_k as follows

$$g_k = \sum_{\omega \in W} \left[c_k^\omega (1 - \mu_k^\omega(x_k)) + \frac{1}{(1 - \mu_k^\omega(x_k))} \right] \quad (8)$$

where the expected count of successor nodes, with novelty $\hat{w}(s) = \omega$, added to the open list at time k is given by $c_k^\omega (1 - \mu_k^\omega(x_k))$, and the second term is the inverse rate of node generation, as we want every possible value of novelty to be

represented in the open list with positive probability. Using Equation 8 we can rewrite J_π as

$$J_\pi = \lim_{N \rightarrow \infty} \frac{1}{N} \mathbf{E}_{c_k} \left[\sum_{k=0}^{N-1} g_k \right] \quad (9)$$

We make an assumption to facilitate obtaining the optimal policy, namely, μ_k will converge to some stationary μ as $k \rightarrow \infty$, so it can be used to estimate the cost of future stages accurately. Also, the expected value of uniformly distributed random variable c_k^ω is calculated from $n_v(\omega)$ and n_e , as $E[c_k^\omega] = n_v(\omega)/n_e$. It follows then from Equation 9 that

$$J_\pi = \sum_{\omega \in W} \left(\frac{n_v(\omega)(1 - \mu^\omega(x_k))}{n_e} + \frac{1}{(1 - \mu^\omega(x_k))} \right) \quad (10)$$

We note that J_π is strictly convex and differentiable over $\mu^\omega \in [0, 1)$, and optimal values for the control inputs correspond with optimal solutions of the optimization problem

$$\min_{\mu \in [0, 1]^{|W|}} J_\pi$$

Such a solution is directly obtained from Equation 10 from the solution of the differential equation $\partial J_\pi / \partial \mu^\omega = 0$,

$$\mu^\omega(x_k) = \begin{cases} 1 - \left(\frac{n_e}{n_v(\omega)} \right)^{\frac{1}{2}}, & \text{if } n_e/n_v(\omega) < 1 \\ 0, & \text{otherwise} \end{cases} \quad (11)$$

Note that the *holding queue* follows the well-established practice of segmenting the search frontier into multiple queues (Richter and Westphal 2010). We use the optimal policy that we derived above to control the different queues based on their $w(s)$ values, ensuring that each queue represents every category $\omega \in W$ with positive probability. The queues are accessed sequentially, expanding all the nodes in the current queue before switching to the next. Also, the way we implement it, the queues are generated lazily, following a partial expansion of nodes whose successor falls in the subsequent queues.

5 Experimental Evaluation

In order to evaluate the impact *novelty approximation* and *open list control* has on width-based planners, we implemented different instantiations of BFWS(f_5): *complete* as described in the Section 2, or *incomplete* if nodes with novelty greater than a given bound are pruned (Lipovetzky and Geffner 2017b). We used the *Downward Lab's* experiment module (Seipp et al. 2017) on a server with Intel Xeon Processors (2 GHz) with a 1800 *sec* and 8 *GB time* and *memory* limit, respectively. All BFWS planners are implemented in C++ using the planning modules from *LAPKT* (Ramirez, Lipovetzky, and Muise 2015) and *grunder* from *Tarski* (Frances and Ramirez 2019). We use every benchmark in the IPC *satisficing* track to evaluate the correctness of the novelty approximation \hat{w} , and performance of new planners that use \hat{w} . In case a domain has appeared over

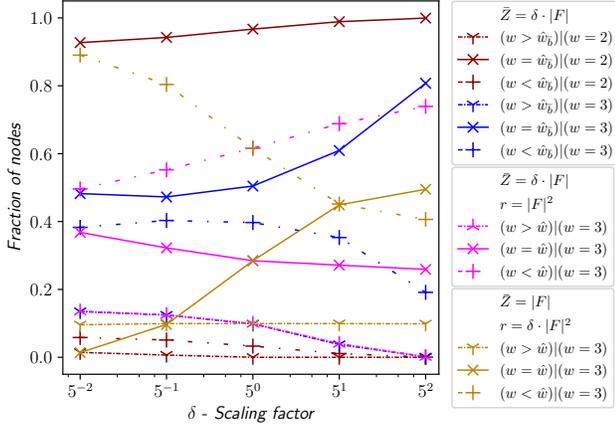


Figure 1: Variation in the rate of *accurate* ($w = \hat{w}$), *lower* ($w > \hat{w}$) and *higher* ($w < \hat{w}$) approximation of novelty w over different sizes of sample (\bar{Z}) and Bloom Filter (r).

multiple IPCs, we used the problem set from the most recent IPC. We compare our new planners against notable *polynomial* planners: BFWS(f_5) with novelty pruning and $\langle 1, 2\text{-C-M} \rangle$, a sequential polynomial planner (Lipovetzky and Geffner 2017b), as well as two state-of-the-art planners DUAL-BFWS (Lipovetzky and Geffner 2017a) and LAMA-first (Richter and Westphal 2010). We show that the introduction of these methods has a significant impact on the performance of the BFWS algorithms.

Correctness of novelty approximation. We evaluate the reliability of the novelty approximation by observing the effect on rate of correct and incorrect (*lower* or *higher*) approximation of novelty over varying sizes of sample \bar{Z} and Bloom Filter r , scaled by a multiplicative factor δ . The novelty approximation \hat{w} is *correct* or *accurate* if $\hat{w}(s) = w(s)$. We limit the maximum size of tuple evaluated to 3, as higher order computations for exact novelty w were infeasible within the practical constraints of time and memory. Thus, $w : S \rightarrow \mathcal{W}$, where $\mathcal{W} = [1, 4]$, and $w = 4$ represents all nodes with $w > 3$. To distinguish the impact of sampling from that of Bloom Filter, we capture the results of novelty approximation *with* and *without* Bloom Filter, hereafter, represented as \hat{w} and $\hat{w}_{\bar{b}}$, respectively. We capture the statistics from 1200 solved instances in IPC satisfying benchmarks.

From Fig. 1, we note that the rate of *correct* approximate novelty ($w = \hat{w}_{\bar{b}}$) increases with sample size \bar{Z} , when Bloom Filters are not used. This backs up our analysis in Section 3 that the accuracy of novelty approximation is likely to increase with sample size. We also observe that the rate does not decrease below 1/2 even for $w = 3$, where the sample is order of $1/|F|^2$ smaller than the exhaustive set. This is a significant improvement over a trivial method of using a coin toss to determine whether or not a tuple t of size l is *new* in state s , which has probability $1/8$ of selecting *correct* novelty, given $w = 3$.

While we note that the novelty approximation without Bloom Filters performs satisfactorily in terms of correct-

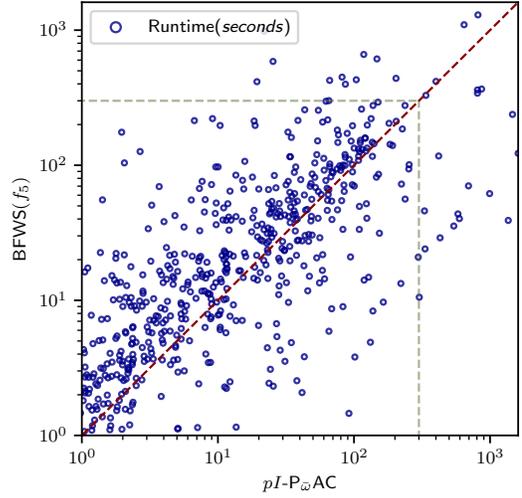


Figure 2: Pairwise comparison of runtime, over All IPC satisfying benchmarks, between BFWS(f_5) and $pI\text{-P}_{\bar{w}}\text{AC}$.

ness, it is still infeasible to store exhaustive set of tuples $\beta_l(s)$ of size $O(|F|^l)$, when $l \geq 3$, for many IPC problems. As discussed in Section 3, we address this by using Bloom Filters for evaluation of $\hat{w} \geq 3$. With this addition, we observe a slight decrease in the rate of *correct* novelty approximation, which is the consequence of false positives, discussed in Section 2.2. Also, we observe that the trend along sample size is reversed, i.e., the rate of *correct* novelty approximation now decreases with increase in sample size. The trend is in line with the theoretical analysis in the Section *Synergies between Sampling and Bloom Filters*. On the other hand, increasing the size of Bloom Filter r improves the results, as the false positive rate decreases.

Performance over benchmarks. Hereafter, we represent a particular configuration of BFWS planner as ' $pI\text{-}(P|\mathcal{L})_{\bar{w}}\text{AC}$ '. The *prefix* ' p ' refers to the use of novelty based pruning for nodes with $\hat{w}(s) > \bar{w}$. ' I ' refers to BFWS called sequentially until the problem is solved over $\bar{w} \in [1, |F|]$. $P_{\bar{w}}$ refers to BFWS(f_5) planner with the set of possible novelty categories $\mathcal{W} = [1, \bar{w} + 1]$. $\mathcal{L}_{\bar{w}}$ refers to BFWS(f_5) with the goal counting heuristics replaced by landmark counts (Richter, Helmert, and Westphal 2008), that is $f_{\mathcal{L}} = \langle w, h_L \rangle$. ' A ' denotes that $\hat{w}(s)$ is used instead of $w(s)$, and ' C ' denotes that BFWS is modified to control open list growth as described in Section 4. All ' AC ' planners were run 4 times with different seeds, so we report the *mean and standard deviation* of statistics of interest.

We set the sample size $\bar{Z} = |F|$ so as to maintain a linear time complexity. We found that D_{max} values between 100 MB and 1GB had similarly good results for ' $pI\text{-P}_{\bar{w}}\text{AC}$ ', we show the results for $D_{max} = 500 \text{ MB}$. For the Bloom Filters size r , we didn't observe much variation between $100 \text{ KB} / 8 \cdot 10^5 \text{ bits}$ and $10 \text{ MB} / 8 \cdot 10^7 \text{ bits}$, with $D_{max} = 1 \text{ GB}$. In our final implementation, we set an initial value of $r = |F|^2$, subject to increase when $\bar{w}r \cdot |V| < D_{max} \wedge |V| = |G| \cdot |F|$, and decrease when $\bar{w}r > D_{max} \wedge |V| = 1$.

	$\bar{w} = 1$	$\bar{w} = 2$	$\bar{w} = 3$	$\bar{w} \geq 4$
# Instances	100.00 %	18.92%	3.68%	1.05%

Table 1: % of instances across all IPC satisficing benchmarks where a node of novelty \bar{w} was recorded in found plans.

A total of 103 instances out of 1691 used the *bank* of Bloom filters V , described in Section 2, ensuring that novelty computation does not exceed D_{max} . Lastly, we use the solution to the problem of minimizing P_f in Equation 1 to chose the number of hash functions as $K = \ln 2 (r/q)$, where $q = \binom{|F|}{\omega}$.

Looking back at the motivation, a key driver for introducing the novelty approximation was to enable novelty computation for values greater than 2, which was infeasible for many IPC domains with the exact novelty definition. The results for p -P3A in Table 2, show that our hypothesis was indeed correct as computing higher novelties with approximation improves coverage. This is substantiated in Table 1 which shows that $\approx 5\%$ of the solved instances had one or more nodes with $\bar{w}(s) \geq 3$ in the solution plan. Moreover, the coverage of approximate planners with $\bar{\omega} = 2$, P_2A and p - P_2A , improves in comparison to P_2 and p - P_2 , respectively, which indicates that there is no apparent demerit of using novelty approximation. The improvement can be attributed to *polynomial* time and space complexity of $\hat{w}(s)$ allowing for additional search capacity.

Though $BFWS(\hat{f}_5)$ performs satisfactorily, it has a key shortcoming which impacts the search within the limited time environment, i.e. for large instances of domains with width $i > 2$, the $BFWS$ search driven by the evaluation function $f_5 = \langle w, \#g \rangle$ exhausts all the available time in expanding nodes with $w(s) \leq 2$. Moreover, the issue gets compounded for domains with high branching factor as the open list doesn't fit within the memory bounds. We address both the issues by applying the open list control discussed in Section 4. In our implementation, the *control* is not applied to child nodes with novelty $w(s)=1$, as the maximum count of such nodes is small, $O(|F|)$, and have minimal impact on space. Note that this method will not cause the search to become incomplete. However, if we choose not to maintain the *holding queue*, we get a search that is *incomplete* and terminates early. Introducing the *open list control* in $BFWS(\hat{f}_5)$ leads to noticeable improvement in coverage of P_2AC and P_3AC which can be observed in Table 2. We do not report tables on plan length due to space limits as plan length remains similar for all configurations.

At this point, we discuss a new planner, where we iteratively run the *polynomial* $BFWS(\hat{f}_5)$ with novelty based pruning, sequentially increasing the number of novelty categories \mathcal{W} at each iteration, $\mathcal{W} = [1, \bar{\omega} + 1]$, over $\bar{\omega} \in [1, |F|]$. We denote the planner as ' pI - $P_{\bar{\omega}}AC$ ' where I stands for *iterative*. Informally, its major advantage is that it taps into the low polynomial space and time complexity of p - $P_{\bar{\omega}}AC$ with small $\bar{\omega}$ values as well as the greater coverage with larger $\bar{\omega}$. This can be observed in Table 2, which shows a signifi-

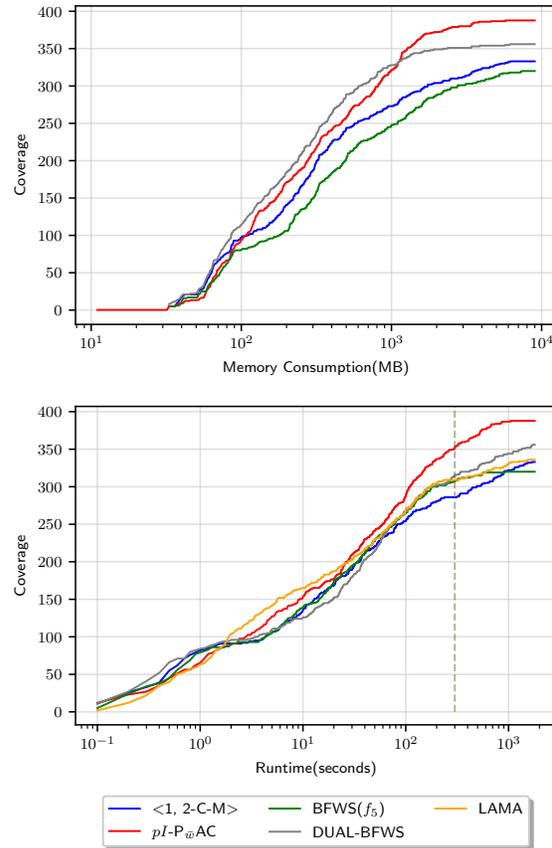


Figure 3: Coverage over memory(MB) and time(seconds) on IPC 2014 and 2018 satisficing benchmarks.

cant jump in coverage compared to $BFWS(f_5)$ with novelty based pruning (p - P_2) and $\langle 1, 2-C-M \rangle$ (B3).

The coverage is also higher than the state-of-the-art $LAMA$ -first (B1) and $DUAL$ - $BFWS$ (B2). Moreover, from Fig. 2, we note that ' pI - $P_{\bar{\omega}}AC$ ' planner has better runtime performance than $BFWS(f_5)$, the winner of Agile track (IPC 2018). It solved 59 more instances than $BFWS(f_5)$ across every IPC satisficing benchmarks with a 300 sec time and 8 GB memory limit. At the same time. Fig. 3 confirms that the space and time consumption is much less than the baseline $BFWS$ planners. It is worth pointing that ' pI - $P_{\bar{\omega}}AC$ ' is probabilistically *incomplete*. Also, we did not observe any difference in coverage of ' pI - $P_{\bar{\omega}}AC$ ' with or without the *holding queue*, as the nodes pruned at one iteration get selected in subsequent iterations with positive probability.

Discussion We show that *approximate novelty search* greatly improves the performance over baseline $BFWS$ planners. The ability to compute $\hat{w} > 2$ using novelty approximation, within practical constraints of time and memory, allows us to use ' pI - $P_{\bar{\omega}}AC$ ' configuration that beats the state-of-the-art. This is impressive for a sequential polynomial planner which uses simple goal counting heuristics $\#g(s)$ and relaxed plan counter $\#r(s)$ along with \hat{w} to direct the search. Also, we can observe that certain domains were af-

domain	B1	B2	P ₂	P ₂ A	P ₂ AC	P ₃ A	P ₃ AC	p -P ₂	B3	p -P ₂ A	p -P ₃ A	pI -P _{$\bar{\omega}$} AC	pI -L _{$\bar{\omega}$} AC
agricola (20)	12	8	11	11±1.0	12±1.7	15±0.8	16±1.3	10	15	10±0.6	15±0.8	16±1.3	12±0.5
airport (50)	34	47	46	46±0.6	46±0.6	44±0.0	44±0.6	46	47	46±0.6	45±1.0	46±0.6	46±0.5
assembly (30)	30	30	30	30±0.6	29±1.0	30±0.6	30±0.6	30	30	30±0.5	30±0.0	30±0.0	30±0.0
caldera (20)	16	20	15	16±1.0	20±0.5	16±1.0	18±0.5	19	20	20±0.5	18±0.5	20±0.5	20±0.0
cavediving (20)	7	7	7	7±0.0	8±0.6	8±0.0	8±0.5	1	8	2±2.9	8±0.5	9±1.0	8±0.5
childsnack (20)	6	10	0	4±1.3	5±1.7	3±1.9	6±0.6	0	2	5±0.5	6±0.6	8±1.3	8±1.3
citycar (20)	5	20	5	5±0.0	20±0.0	5±0.0	20±0.6	20	20	20±0.5	5±0.0	20±0.0	20±0.0
data-network (20)	13	11	9	12±1.7	19±1.0	11±2.5	18±0.5	16	14	17±1.0	16±0.6	18±0.6	18±0.6
depot (22)	20	22	22	22±0.0	22±0.0	22±0.0	22±0.0	22	22	22±0.0	22±0.0	22±0.0	22±0.0
flashfill (20)	14	16	12	14±2.4	14±0.6	14±2.2	14±1.0	15	9	14±2.4	14±1.9	14±1.0	14±1.0
floortile (20)	2	2	1	2±0.5	2±0.0	2±0.0	2±0.0	0	1	1±0.0	2±0.5	2±0.0	2±0.0
hiking (20)	20	12	12	14±2.1	8±0.8	18±1.0	20±0.5	9	13	12±1.8	20±0.0	20±0.0	19±0.5
maintenance (20)	11	17	17	16±0.5	16±0.6	16±0.5	17±0.5	17	17	16±0.5	16±0.5	17±0.5	17±0.5
mprime (35)	35	35	32	30±0.6	35±0.0	31±0.5	34±0.8	35	35	35±0.0	32±0.8	35±0.0	35±0.0
mystery (30)	19	19	19	19±0.0	19±0.0	19±0.5	19±0.5	19	18	19±0.0	19±0.5	19±0.5	19±0.5
nomystery (20)	11	19	13	14±1.0	12±1.0	13±0.5	14±1.0	13	13	12±1.7	14±0.6	15±1.0	18±1.4
nurikabe (20)	9	14	16	14±0.6	15±1.3	14±0.6	15±2.1	16	16	14±0.5	14±0.6	15±1.0	14±0.6
org-synth-split (20)	12	11	5	6±0.5	3±0.8	7±1.0	5±1.4	4	3	4±0.5	6±1.0	7±0.0	6±0.5
parcprinter (20)	20	16	9	5±1.0	5±1.9	5±0.8	6±1.0	9	16	6±1.0	5±1.0	8±0.0	6±0.5
pathways-neg (30)	24	30	23	30±0.6	29±1.5	30±0.6	29±0.5	24	27	30±0.5	30±0.6	30±0.0	30±0.0
pegsol (20)	20	20	20	20±0.0	20±0.5	20±0.0	20±0.0	5	20	12±1.5	18±0.5	20±0.0	20±0.0
pipeworld-nt (50)	43	50	50	50±0.0	50±0.0	50±0.0	50±0.0	50	50	50±0.0	50±0.0	50±0.0	50±0.0
pipeworld-t (50)	43	38	43	42±0.5	42±0.6	42±0.5	43±0.8	41	39	41±1.5	42±0.5	42±1.2	43±1.5
psr-small (50)	50	50	48	49±0.5	49±0.5	50±0.0	50±0.0	31	46	34±1.3	43±0.8	49±0.5	48±0.6
rovers (40)	40	37	39	40±0.0	40±0.0	40±0.0	40±0.0	39	38	40±0.0	40±0.0	40±0.0	40±0.0
satellite (36)	36	31	27	30±0.8	32±0.6	30±0.5	30±0.0	27	31	32±0.5	30±0.5	34±0.6	34±0.6
schedule (150)	150	149	149	149±1.0	149±0.8	149±1.0	150±0.6	149	149	149±1.0	149±1.0	149±1.0	149±1.0
settlers (20)	18	8	7	6±1.0	12±0.6	6±1.3	10±0.0	10	11	9±1.0	6±1.0	12±0.6	17±1.3
snake (20)	5	12	19	16±0.5	15±0.8	17±0.5	17±0.5	18	3	16±0.5	17±0.5	20±0.5	20±0.5
sokoban (20)	19	17	14	15±0.5	10±1.0	16±0.5	14±0.8	6	13	4±0.6	11±0.5	16±0.6	16±0.6
spider (20)	16	14	13	15±1.0	14±1.0	15±1.0	14±1.3	13	11	15±1.0	15±1.0	14±1.0	15±0.5
storage (30)	20	28	29	30±0.6	30±0.6	30±0.6	30±0.5	30	29	30±0.6	30±0.6	30±0.5	30±0.6
termes (20)	16	9	9	10±0.0	8±0.6	9±1.0	8±1.3	1	6	2±0.5	6±1.9	7±1.4	10±1.3
tetris (20)	16	16	20	20±0.0	20±0.0	20±0.0	20±0.0	20	18	20±0.0	20±0.0	20±0.0	20±0.0
thoughtful (20)	15	20	20	20±0.0	20±0.0	20±0.0	20±0.0	20	20	20±0.0	20±0.0	20±0.0	20±0.0
tidybot (20)	17	18	19	20±0.0	20±0.5	20±0.0	20±0.0	20	20	20±0.0	20±0.0	20±0.0	19±0.5
tpp (30)	30	29	29	30±0.6	30±0.0	29±0.0	30±0.0	30	30	30±0.0	30±0.0	30±0.0	30±0.0
transport (20)	16	20	20	20±0.0	20±0.0	20±0.0	20±0.0	20	20	20±0.0	20±0.0	20±0.0	20±0.0
trucks-strips (30)	18	16	9	9±0.8	9±1.3	9±0.8	10±1.4	11	8	12±1.8	11±1.3	12±1.3	12±0.8
Total (1691)	1456	1496	1436	1455±8.7	1476±4.2	1463±8.9	1502±4.9	1414	1456	1438±5.9	1462±8.0	1524±2.5	1516±5.0

Table 2: Coverage over all satisficing benchmarks from IPCs: *complete* - B1: LAMA-*first*, B2: DUAL-BFWS and 'P...', and *polynomial incomplete* - B3: $\langle 1, 2\text{-C-M} \rangle$ and ' p -P...'. P _{$\bar{\omega}$} refers to BFWS(f_5) planner with $\mathcal{W} = [1, \bar{\omega} + 1]$, L _{$\bar{\omega}$} is BFWS($f_{\mathcal{L}}$), which uses *Landmarks*, ' I ' stands for *Iterative*, ' A ' for *approximate*, and ' C ' for *control over open list*. The *mean coverage* is shown along with the *standard deviation* for the planners that use sampling. Domains which are fully solved by all planners are omitted but included in supplementary material (Singh et al. 2021). The best results are highlighted in bold.

ected more than others. Specifically, the domains *citycar*, *data-network*, *hiking* and *satellite* benefited significantly.

We found that the *open list control* significantly benefited the domains *citycar* and *data-network* which have a high branching factor but solvable with $\hat{\omega} \leq 2$. *Citycar* in particular was fully solvable with $\hat{\omega} = 1$ and discarding nodes with $w(s) > 1$ didn't impact the order of expansion. *Hiking* and *satellite* on the other hand required expansion of nodes of $w(s) > 2$, and the increased coverage highlights the importance of policy based control of different novelty categories in the open list. *Childsnack* and *Floortile* however showed no improvement, which is a combined effect of high width and the fact that our goal count heuristic $\#g(s)$ is not informed enough.

6 Conclusion

The proposed methods of novelty approximation and open list control in *BFWS* not only have positive impact on coverage but also on the overall time and space complexity of the

search, resulting in new state-of-the-art planners over satisficing benchmarks from every IPC since 1998 and more significantly the last 2 IPCs (2014 and 2018). These results strongly suggest that probabilistically complete search algorithms are a promising research direction in classical planning. This is specially crucial in limited time and memory environments where the search must work within hard constraints on time and memory. However, we must note that approximate novelty search is by no means a silver bullet, and certain domains including *Childsnack* and *Floortile* still remain unsolvable. We hope this work brings about the insights to develop the next generation of classical planners, that scale up better as the intractability of the benchmarks ramps up and tackle the inherent limitations of BFS.

Acknowledgements

Anubhav Singh is supported by Melbourne Research Scholarship established by the University of Melbourne.

Javier Segovia-Aguas is supported by TAILOR, a project

funded by EU H2020 research and innovation programme no. 952215, an ERC Advanced Grant no. 885107, and grant TIN-2015-67959-P from MINECO, Spain.

This research was supported by use of the Nectar Research Cloud, a collaborative Australian research platform supported by the National Collaborative Research Infrastructure Strategy (NCRIS).

References

- Bertsekas, D. P. 2017. *Dynamic Programming and Optimal Control*. Athena Scientific, 4th edition.
- Bloom, B. H. 1970. Space/Time Trade-Offs in Hash Coding with Allowable Errors. *Communications of the ACM* 13(7): 422–426.
- Bonet, B.; and Geffner, H. 2001. Planning as Heuristic Search. *Artificial Intelligence* 129(1–2): 5–33.
- Broder, A.; and Mitzenmacher, M. 2004. Network Applications of Bloom Filters: A Survey. *Internet Mathematics* 1(4): 485–509.
- Burns, E. A.; Hatem, M.; Leighton, M. J.; and Ruml, W. 2012. Implementing Fast Heuristic Search Code. In *Proc. of the Annual Symposium on Combinatorial Search*.
- Chakrabati, P. P.; Ghose, S.; Acharya, A.; and de Sarkar, S. C. 1989. Heuristic Search in Restricted Memory. *Artificial Intelligence* 41: 197–221.
- Dionne, A. J.; Thayer, J. T.; and Ruml, W. 2012. Deadline-Aware Search Using On-Line Measures of Behavior. In *Proc. of the Annual Symposium on Combinatorial Search*.
- Edelkamp, S.; and Schrödl, S. 2012. *Heuristic Search – Theory and Applications*. Morgan Kaufmann.
- Frances, G.; and Ramirez, M. 2019. Tarski - An AI Planning Modeling Framework. <https://github.com/aig-upf/tarski>. Accessed: 2019-11-11.
- Geffner, H.; and Bonet, B. 2013. *A Concise Introduction to Models and Methods for Automated Planning*. Morgan & Claypool Publishers.
- Hoffmann, J.; and Nebel, B. 2001. The FF Planning System: Fast Plan Generation Through Heuristic Search. *Journal of Artificial Intelligence Research* 14: 253–302.
- Hoffmann, J.; Porteous, J.; and Sebastia, L. 2004. Ordered Landmarks in Planning. *Journal of Artificial Intelligence Research* 22: 215–278.
- IPC18. 2018. International Planning Competition 2018 – Classical Tracks. <https://ipc2018-classical.bitbucket.io/>. Accessed: 19-01-2020.
- Katz, M.; Lipovetzky, N.; Moshkovich, D.; and Tuisov, A. 2017. Adapting Novelty to Classical Planning as Heuristic Search. In *International Conference on Automated Planning and Scheduling*, 172–180.
- Lipovetzky, N.; and Geffner, H. 2012. Width and Serialization of Classical Planning Problems. In *Proc. of European Conference in Artificial Intelligence (ECAI)*.
- Lipovetzky, N.; and Geffner, H. 2017a. Best-First Width Search: Exploration and Exploitation in Classical Planning. In *Proc. of the AAAI Conference on Artificial Intelligence*.
- Lipovetzky, N.; and Geffner, H. 2017b. A polynomial planning algorithm that beats LAMA and FF. In *International Conference on Automated Planning and Scheduling*.
- Louridas, P. 2017. *Real-World Algorithms: A Beginner's Guide*. MIT Press.
- Nilsson, N.; and Fikes, R. 1971. STRIPS: A new approach to the application of theorem proving to problem solving. *Artificial Intelligence* 1: 27–120.
- Pearl, J. 1983. *Heuristics*. Addison-Wesley.
- Ramirez, M.; Lipovetzky, N.; and Muise, C. 2015. Lightweight Automated Planning ToolKit. <http://lapkt.org/>. Accessed: 2020-01-20.
- Richter, S.; Helmert, M.; and Westphal, M. 2008. Landmarks Revisited. In *Proc. of the AAAI Conference on Artificial Intelligence*.
- Richter, S.; and Westphal, M. 2010. The LAMA planner: Guiding cost-based anytime planning with landmarks. *Journal of Artificial Intelligence Research* 39: 127–177.
- Seipp, J.; Pommerening, F.; Sievers, S.; and Helmert, M. 2017. Downward Lab. URL <https://doi.org/10.5281/zenodo.790461>.
- Singh, A.; Lipovetzky, N.; Ramirez, M.; and Segovia-Aguas, J. 2021. Technical Appendix. <https://doi.org/10.5281/zenodo.4606931>. Accessed: 2021-3-15.
- Vadlamudi, S. G.; Aine, S.; and Chakrabarti, P. P. 2011. MAWA*: A Memory-Bounded Anytime Heuristic Search Algorithm. *IEEE Transactions on Systems, Man, and Cybernetics* 41(3): 725–735.