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 is an effective approach to classical planning which has produced many successful algorithms over the years. A key feature which distinguishes width-based search from classic heuristic search algorithms is the use of *specific structural properties* of the *explored state space* to guide the *exploration* and goal-directed heuristic measures for *exploitation*. The *structural properties* are captured as an n-ary relation over the *fluents* which is processed to compute the *state novelty*. The size of the relation and the time complexity of computing novelty measure is exponential on the arity n. Approximate novelty search introduces novel polynomial approximations of *state novelty* and width-based search. It uses *Bloom filter* to *efficiently* represent the interpretation of the relational predicate and *random sampling* in the computation of *state novelty*. It also uses an adaptive policy which decides to delay the generation of successor states. In this paper, we explain the integration of these two techniques into the *polynomial-time* variant of Best-First Width Search (BFWS), one of the most successful width-based algorithm in *satisficing* planning.

.

Introduction

Width-based search algorithms rely on the notion of *state novelty* which is an *orthogonal* measure to *goal-directed* heuristics. While the heuristics provide an approximation of the distance to the goal, the novelty measures instead captures how *novel* the state is with respect to the *explored* state space. Several width-based search algorithms have been proposed (Lipovetzky and Geffner 2014; Lipovetzky et al. 2014; Lipovetzky and Geffner 2017a,b) out of which best-first width search (BFWS) has been the most acclaimed. A major shortcoming of the width-based methods is that complexity of computing novelty measure is exponential on the number of discrete level or categories used to rank the states. While there exists an upper bound on the number of novelty categories required to solve a given classical planning instance (Lipovetzky and Geffner 2012), a large bound results in impractical space and time requirements for novelty computation. Approximate Novelty Search (Singh et al. 2021) proposes a probabilistic approximation of novelty measure which trades off accuracy of novelty computation for commitments on space and time complexity. This allows width-based search algorithms to tap into the search space associated with higher novelty categories. Next, we present a brief account of best-first width search and novelty approximation, along with the description of the planner configurations that we have submitted in *agile* and *satisficing* track of the IPC.

Approximate Novelty Search

BFWS (Lipovetzky and Geffner 2017a) is a best-first search algorithm which uses a tuple of functions $f(n) = (w, h_1, \ldots, h_m)$ to guide the search, where $w : S \mapsto \mathcal{W}$ measures the novelty of a state, $\mathcal{W} \in \mathbb{N}$ is the set of novelty categories and $H = \{h_1, \ldots, h_m\}$. BFWS algorithm sorts the nodes in order of importance using the first function in $f(n)$, recursively breaking ties using the next function provided in $f(n)$. The *approximation* of BFWS (Singh et al. 2021) uses the same notion to guide the search with two differences (1) $f(n) = (\hat{w}, h_1, \ldots, h_m)$, where $\hat{w} : S \mapsto \mathcal{W}$ is a function measuring the *approximate* novelty, (2) it uses an adaptive policy, derived from the analytical solution to an infinite-horizon Markov Decision Problem (MDP), that decides whether to forgo the expansion of nodes in the open list. These improvements result in a state-of-the-art BFWS planner over IPC *satisficing* benchmarks by simply pairing novelty measure with goal-counting heuristic $\#g$, i.e. $f(n) = (\hat{w}, \#g)$.

Sequential *polynomial approximate* BFWS(f_5)

In this planner, we make sequential calls to the *polynomial approximate* BFWS(f_5) (Singh et al. 2021) with novelty based pruning until we run out of time. We denote the sequential configuration as 'pI-BFWS(f_5)AC', were p stands for novelty based *pruning*, I for iterative, A for novelty *approximation* and C for adaptive *control* of the open-list. The *polynomial approximate* BFWS(f_5) is denoted as 'p-BFWS$(f_5)_{\bar{\omega}}$AC', where the set of novelty categories $\mathcal{W} = [1, \bar{\omega} + 1]$ and the nodes with $\hat{w}(n) > \bar{\omega}$ are pruned.

We start by calling 'p-BFWS$(f_5)_{\bar{\omega}}$AC' with $\bar{\omega} = 1$, i.e. nodes with $\hat{w}(n) > 1$ are pruned. At each subsequent call, we increase the novelty bound $\bar{\omega}$ by 1. At small values of $\bar{\omega}$ the planner taps into the low polynomial search

---

space of BFWS(f_5) with a small probability of error in the *novelty* computation. As the value of $\bar{\omega}$ grows large it becomes harder to compute novelty *exactly*. Indeed, the original BFWS(f_5) would exceed the space and time limits for $\bar{\omega} > 2$ on many IPC benchmark domains. 'p-BFWS$(f_5)_{\bar{\omega}}$AC' allows us to tap into that space by trading off the quality of novelty computation for time and space guarantees. We have entered this planner into *agile* and *satisficing* track, with one difference in the *satisficing* submission - once 'pI-BFWS(f_5)AC' finds a solution we call the implementation of weighted A* used in LAMA (Richter and Westphal 2010) to improve the plan quality until timeout.

References

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. 2014. Width-based algorithms for classical planning: New results. In *Proc. of European Conference in Artificial Intelligence (ECAI)*, 1059–1060. IOS Press.

Lipovetzky, N.; and Geffner, H. 2017a. Best-First Width Search: Exploration and Exploitation in Classical Planning. In *Proc. of the AAAI Conference (AAAI).*

Lipovetzky, N.; and Geffner, H. 2017b. A polynomial planning algorithm that beats LAMA and FF. In *Int'l Conference on Automated Planning and Scheduling (ICAPS).*

Lipovetzky, N.; Ramirez, M.; Muise, C.; and Geffner, H. 2014. Width and inference based planners: Siw, bfs (f), and probe. *Proc. of the Int'l Planning Competition (IPC)* 43.

Richter, S.; and Westphal, M. 2010. The LAMA planner: Guiding cost-based anytime planning with landmarks. *Journal of Artificial Intelligence Research (JAIR)* 39: 127–177.

Singh, A.; Lipovetzky, N.; Ramirez, M.; and Segovia-Aguas, J. 2021. Approximate novelty search. In *Int'l Conference on Automated Planning and Scheduling (ICAPS)*, volume 31, 349–357.