

Expected Runtime Comparisons Between Breadth-First Search and Constant-Depth Restarting Random Walks

Daniel Platnick and Richard Valenzano

Toronto Metropolitan University, Toronto, Canada
Vector Institute for Artificial Intelligence, Toronto, Canada
{daniel.platnick, rick.valenzano}@torontomu.ca

Abstract

When greedy search algorithms encounter a local minima or plateau, the search typically devolves into a breadth-first search (BrFS), or a local search technique is used in an attempt to find a way out. In this work, we formally analyze the performance of BrFS and constant-depth restarting random walks (RRW) — two methods often used for finding exits to a plateau/local minima — to better understand when each is best suited. In particular, we formally derive the expected runtime for BrFS in the case of a uniformly distributed set of goals at a given goal depth. We then prove RRW will be faster than BrFS on trees if there are enough goals at that goal depth. We refer to this threshold as the crossover point. Our bound shows that the crossover point grows linearly with the branching factor of the tree, the goal depth, and the error in the random walk depth, while the size of the tree grows exponentially in branching factor and goal depth. Finally, we discuss the practical implications and applicability of this bound.

1 Introduction

Greedy algorithms like *Greedy Best First Search (GBFS)* (Doran and Michie 1966) and *Enhanced Hill-Climbing (EHC)* (Hoffmann 2001) have been shown to be very effective at solving graph search and planning problems. However, when using a flawed heuristic function, the search performed by these algorithms often devolves into *breadth-first search (BrFS)*. This is explicitly by design in the case of EHC, which performs a sequence of BrFS instances, each in search of a new lower heuristic value. In GBFS, it often happens as a natural consequence of *plateaus*, where the heuristic function does not provide enough information to distinguish between states, and so the search is blind.

The *restarting random walk (RRW)* approach (Nakhost and Müller 2009; Xie, Müller, and Holte 2014, 2015) is an alternative to BrFS for escaping local minima. This method generates a sequence of random walks until one is found that exits the local minima. It is low memory, but does not have the strong termination guarantees that the systematic search of BrFS leads to.

The main objective of this work is to further our understanding of BrFS and RRW to better understand when each is best for dealing with local minima and plateaus, when

these are encountered as part of other search algorithms. As an example of why this is needed, we have performed preliminary experiments testing the performance of an EHC algorithm which replaces BrFS with RRWs. These experiments were performed using Pyperplan (Alkharaji et al. 2020). In particular, we ran each method 10 times per problem, each time with a 10 minute time limit on an M3 MacBook with 16 GB of RAM¹. The results show that different methods for escaping local minima seem best suited for different domains. Using RRWs in EHC instead of BrFS increases coverage in some domains, for example, from 60% to 95% in TPP, and 81% to 91% in Transport. In other domains, we see that using RRW decreases performance, such as the drop in coverage from 93% to 76% in Scanalyzer and 100% to 84% in BlocksWorld.

In this paper, we focus on formally analyzing BrFS and RRW to better understand where each approach is best suited. Our main result shows that the expected runtime of RRW on directed trees is less than BrFS if the number of goals (*ie.* exit states of the local minima) is high enough. Interestingly, the number of goals at this *crossover point* is linear in the goal depth and branching factor, meaning it grows much slower than the overall size of the tree itself. A consequence of this is that the density of the number of goals (*ie.* the percentage of goals to total states) needed for RRW to outperform BrFS actually decreases as the depth of the goals increases. We then conclude by discussing the practical implications of this result as it relates to when each algorithm is most appropriate.

2 Preliminaries

In this section, we introduce the terminology used and briefly summarize the algorithms covered in this work.

A graph search task is defined by the tuple $\mathcal{S} = \langle G, v_{\mathcal{I}}, V_{\mathcal{G}} \rangle$, where $G = \langle V, E \rangle$ is a graph with vertices V and edges E , $v_{\mathcal{I}} \in V$ is the *start* or *initial* vertex and $V_{\mathcal{G}} \subseteq V$ is the *goal* set. The objective is to find a *path* of vertices $\langle v_0, \dots, v_k \rangle$ where $v_0 = v_{\mathcal{I}}$, $v_k \in V_{\mathcal{G}}$, and for every $0 \leq i < k$, $\langle v_i, v_{i+1} \rangle \in E$. Given our specific focus

¹While Pyperplan is not a state-of-the-art planner, we consider it sufficient for demonstrating that the competing methods of using BrFS or RRW for escaping local minima each are best suited for different problems.

algorithms, we ignore edge costs in this work.

For any $v \in V$, we refer to the number of edges in the shortest path to v as the *level* of v . For example, $v_{\mathcal{I}}$ has a level of 0, any vertex that can be reached from $v_{\mathcal{I}}$ by one edge has a level of 1, etc. The *goal level* d^* is defined as the minimum level of any vertex in V_G .

The search algorithms we focus on perform two main operations on vertices. The first is a *goal test* (ie. if v is in V_G). The second is a *successor generation step*. This means finding and generating the successor set of v , defined as $\{v' | \langle v, v' \rangle \in E\}$. When both of these operations are performed on a vertex, we say it has been *expanded*. As will be described below, for both our focus algorithms, the number of goal tests will be exactly 1 larger than the number of successor generation steps. Thus, we can measure runtime equally well using the count of the number of times either operation is performed. We therefore focus on the number of goal tests in the remainder of this work.

Note that while our formal analysis is given in terms of finding a goal in a given graph search task, the analysis equally applies to the case of escaping from a single local minima in a given state-space. In such a situation, the goal test involves simply checking if the heuristic value of a given state (ie. vertex) is less than the heuristic value of the initial state. Thus, the statements below also provide expected runtime results for the time needed to escape from local minima of different sizes.

Breadth-First Search (BrFS)

For BrFS, we assume the reader’s familiarity with the use of open and closed lists in best-first search algorithms like A^* . At every iteration, BrFS will select one of the vertices v in the open list with the lowest level and perform a goal test. If the test succeeds, the search immediately terminates. If the test does not succeed, then the successors of v are generated, those that have been generated for the first time are added to the open list, and v is moved to the closed list. This process then repeats until a goal test succeeds or the open list is emptied. Thus, when we refer to BrFS, we are considering an algorithm equivalent to uniform-cost search (Felner 2011) — or A^* with a heuristic that always returns 0 — on a unit-cost graph.

By definition, BrFS expands all vertices in a level before proceeding to the next level (assuming no goal is found), and BrFS does not perform a goal test on any vertex with a level larger than d^* . Since every failed goal test is immediately followed by a successor generation step and BrFS terminates after the test succeeds once, the number of goal tests is exactly one more than the number of successor generation steps. Moreover, the number of goal tests will depend on how many vertices in the goal level are expanded before a goal vertex, which will depend entirely on how the goal vertices are distributed at that level and tiebreaking. Thus, the number of goal tests made by BrFS on a search task \mathcal{S} is a random variable, which we denote by $B(\mathcal{S})$.

Constant-Depth Restarting Random Walks (RRW)

The constant-depth RRW algorithm is shown in Algorithm 1. It generates a series of random paths called *random walks*,

Algorithm 1: The Constant-Depth RRW Algorithm

Input: task $\mathcal{S} = \langle G = \langle V, E \rangle, v_{\mathcal{I}}, V_G \rangle$, max depth t

- 1: **if** $v_{\mathcal{I}} \in V_G$ **then**
- 2: **return** path $\langle v_{\mathcal{I}} \rangle$
- 3: **while true do**
- 4: $P \leftarrow \langle v_{\mathcal{I}} \rangle, v \leftarrow v_{\mathcal{I}}, d \leftarrow 0$
- 5: **while** $d < t$ **do**
- 6: $v \leftarrow$ vertex sampled from $\{v' | \langle v, v' \rangle \in E\}$
- 7: Append v to $P, d+ = 1$
- 8: **if** $v \in V_G$ **then**
- 9: **return** P

each starting from $v_{\mathcal{I}}$ (lines 4 to 9). At every step of the walk, the algorithm terminates if a goal is found (line 8). Otherwise, the successor of the last vertex is generated, and it is added to the path. The algorithm restarts to the initial vertex when the maximum depth is reached (lines 5 and 4).

Notice that the goal test is performed on $v_{\mathcal{I}}$ exactly once (line 1). Doing otherwise is clearly unnecessary. This approach also ensures that the algorithm performs exactly one more goal test than successor generation steps, since every successor generation step (line 6) is immediately followed by a goal test (line 8).

Clearly, if the maximum depth t is less than the goal level, then constant-depth RRW will not terminate. Thus, we assume a reasonable “guess” of a lower bound on the correct depth can be made. That is, we assume that $t = ed^*$ where $e \geq 1$ and ed^* is a non-negative integer. We refer to e as the *depth error*, and denote the random variable for the number of goal tests by a constant-depth RRW with depth error e on search task \mathcal{S} as $R_e(\mathcal{S})$. Note that for simplicity, we assume that every vertex with a level less than ed^* has at least one successor (ie. a random walk never terminates early).

3 General Expected Runtime Analysis

In this section, we characterize the expected runtime of BrFS and RRW in general settings. The first result is for BrFS when the goal vertices are uniformly distributed at the goal level. This theorem will equivalently hold if random tiebreaking is used to order the vertices for expansion.

Theorem 3.1. *If \mathcal{S} has $N_{\mathcal{O}}$ unique vertices with a level less than d^* , and $g \geq 1$ goal vertices uniformly distributed amongst the $N_{d^*} \geq g$ unique vertices at the goal level, then*

$$\mathbb{E}[B(\mathcal{S})] = N_{\mathcal{O}} + (N_{d^*} + 1)/(g + 1)$$

Proof. Let X be the number of goal tests that BrFS performs on vertices at level d^* . Since BrFS does not test any vertices with a level deeper than d^* , $\mathbb{E}[B(\mathcal{S})] = \mathbb{E}[N_{\mathcal{O}} + X] = N_{\mathcal{O}} + \mathbb{E}[X]$ since $N_{\mathcal{O}}$ is a constant.

Note, because the goal vertices are uniformly distributed, $\mathbb{E}[X]$ is equivalent to the expected number of selections needed when randomly picking vertices from the goal level without replacement, until one of the g vertices is picked. Let v_i be any one of the $(N_{d^*} - g)$ non-goal vertices at level d^* , and let Z_i be an indicator random variable for the event that v_i is picked before any one of the g goal vertices. Therefore $\mathbb{E}[X] = \mathbb{E}[Z_1 + \dots + Z_{N_{d^*} - g}] + 1$, since $\mathbb{E}[X]$ is the

number of non-goal vertices tested plus one for the final goal vertex tested.

Now notice that $\mathbb{P}[Z_i] = 1/(g+1)$ since there are $(g+1)!$ ways of ordering the $(g+1)$ vertices in the set containing v_i and the g goal vertices, and $g!$ of these orderings start with v_i . Importantly, this probability does not change when some other non-goal vertex $v_j \neq v_i$ is picked. Therefore, the Z_i 's are independent and so

$$\begin{aligned} \mathbb{E}[X] &= 1 + \mathbb{E}\left[\sum_{i=1}^{N_{d^*}-g} Z_i\right] = 1 + \sum_{i=1}^{N_{d^*}-g} \mathbb{E}[Z_i] & (1) \\ &= 1 + (N_{d^*} - g)/(g+1) = (N_{d^*} + 1)/(g+1) & (2) \end{aligned}$$

Line 2 holds since the Z_i are indicator variables and so $\mathbb{E}[Z_i] = \mathbb{P}[Z_i]$, and since we are summing over $(N_{d^*} - g)$ of them with the same expectation. Adding this to $N_{\mathcal{O}}$ yields the desired result. \square

We now consider the runtime of constant depth RRW. For this result, we assume that each random walk has an independent probability of $0 \leq s \leq 1$ of reaching a goal vertex. We call s the *success probability*.

Theorem 3.2. *If \mathcal{S} is a search task with goal level d^* such that $s > 0$ and all g goal vertices are at level d^* , then*

$$\mathbb{E}[R_e(\mathcal{S})] = \frac{ed^*}{s} - (e-1)d^* + 1$$

Proof. Let Y be the random variable for the *number of random walks* it takes to find a goal (ie. the goal is first seen on random walk Y). Notice that

$$\mathbb{E}[R_e(\mathcal{S})] = \mathbb{E}[ed^*(Y-1) + d^* + 1] \quad (3)$$

$$= ed^*\mathbb{E}[Y] - (e-1)d^* + 1 \quad (4)$$

Line 3 holds because the first Y random walks all perform ed^* goal tests before restarting and exactly d^* tests on the last random walk (when it finds a goal at depth d^*). The added 1 comes from the single goal test of $v_{\mathcal{T}}$.

Since every random walk has an independent and identically distributed success probability of s , Y follows the well-known geometric distribution. Thus, $E[Y] = 1/s$. Substituting this into Line 4 gives the desired result. \square

4 Comparing BrFS and RRW on Trees

We now consider the performance of BrFS and RRW on directed trees with a constant branching factor. We begin by using Theorems 3.1 and 3.2 to find more specific expressions for the expected number of goal tests in this case.

Corollary 4.1. *If $\mathcal{S}_{\mathcal{T}}$ is a search task on a directed tree with constant branching factor $b \geq 2$ and $g \geq 1$ goals uniformly distributed at goal level d^* , then*

$$\mathbb{E}[B(\mathcal{S}_{\mathcal{T}})] = (b^{d^*} - 1)/(b-1) + (b^{d^*} + 1)/(g+1)$$

This follows since there are b^ℓ vertices at every level $\ell \geq 0$ of such a tree. Thus, we can apply Theorem 3.1 using $N_{d^*} = b^{d^*}$ and $N_{\mathcal{O}} = b^0 + b^1 + \dots + b^{d^*-1} = (b^{d^*} - 1)/(b-1)$, which follows from the formula for the sum of a geometric series.

Next, we consider RRW on such a tree:

Corollary 4.2. *Let $\mathcal{S}_{\mathcal{T}}$ be a search task on a directed tree with constant branching factor $b \geq 2$ and $g \geq 1$ goals all at the goal level d^* . If RRW samples edges uniformly at each step, then*

$$\mathbb{E}[R_e(\mathcal{S}_{\mathcal{T}})] = \frac{ed^*b^{d^*}}{g} - (e-1)d^* + 1$$

This follows immediately from Theorem 3.2 since the success probability is g/b^{d^*} because all vertices at the goal level are equally likely to be visited on any random walk.

Figure 1a uses these corollaries to show the expected number of goal tests for BrFS and RRW on a tree with $b = 4$, $d^* = 6$, and different numbers of goals. BrFS performs significantly fewer goal tests for a small number of goals, but RRW can overtake BrFS quite quickly depending on the error value (note that there are 4096 vertices at the goal level). We refer to the number of goals at which RRW matches or surpasses BrFS as the *crossover point*. The next theorem shows that the crossover point is linear in each of the branching factor, goal depth, and depth error. The theorem only covers the case where the $d^* \geq 2$ and $ed^* > 2$. This requires that either the goal level $d^* > 2$ or the random walk depth error is $e > 1$. We discuss the remaining cases after the theorem.

Theorem 4.1. *Let $\mathcal{S}_{\mathcal{T}}$ be a search task on a directed tree with constant branching factor $b \geq 2$, $d^* \geq 2$, and $g \geq 1$ goals all uniformly distributed at the goal level. If RRW samples edges uniformly at each step and $ed^* > 2$, then $\mathbb{E}[B(\mathcal{S}_{\mathcal{T}})] \geq \mathbb{E}[R(\mathcal{S}_{\mathcal{T}})]$ if either the following holds:*

1. $g = b^{d^*}$
2. $g < b^{d^*}$ and $g \geq (ed^* - 1)(b-1) + 1$

Proof. Let x denote $\mathbb{E}[B(\mathcal{S}_{\mathcal{T}})] - \mathbb{E}[R(\mathcal{S}_{\mathcal{T}})]$. By Corollaries 4.1 and 4.2

$$x = \frac{b^{d^*} - 1}{b-1} + \frac{b^{d^*} + 1}{g+1} - \frac{ed^*b^{d^*}}{g} + (e-1)d^* - 1 \quad (5)$$

We will now show $x \geq 0$ in case 1, when $g = b^{d^*}$. Notice that this simplifies line 5 to $x = (b^{d^*} - 1)/(b-1) - d^*$. Since $b \geq 2$ and $d^* \geq 2$, x is clearly positive here.

Let us now consider Case 2. For readability, we let $x = y + z$, where $z = (e-1)d^* - 1$, and y is the remaining terms in line 5. Since $x \geq 0$ if and only if $Cx \geq 0$ for some positive constant C , we can show the required statement by showing $x' \geq 0$ where $x' = cx$ and $c = (b-1)(g+1)g$. This is done to avoid the denominators. Similarly, we define $y' = cy$ and $z' = cz$ and derive the following:

$$y' = (b^{d^*} - 1)(g+1)g + (b^{d^*} + 1)(b-1)g \quad (6)$$

$$- ed^*b^{d^*}(b-1)(g+1)$$

$$= b^{d^*}[(g+1)g + (b-1)g - ed^*(b-1)(g+1)] + v \quad (7)$$

where $v = (b-1)g - (g+1)g$. Here, we have just separated and factored out the terms with b^{d^*} in them. Now notice the lower bound in Case 2 is equivalent to $ed(b-1) \leq g+b+2$. By substituting this into Line 7, we get




Figure 1: Comparing BrFS and RRW with different depth errors (e) on a directed tree with a branching factor of 4. Figure 1a shows the expected goal tests when the goal depth is 6 for different numbers of goals. Figure 1b shows how the crossover point changes at different goal levels. The crossover point as computed empirically is shown as a dashed line, while the bound from Theorem 4.1 is a solid line. Figure 1c shows how the goal density crossover changes with the goal level.

$$y' \geq b^{d^*} [(g+1)g + (b-1)g - (g+b+2)(g+1)] + v \quad (8)$$

$$\geq b^{d^*} (g-b-2) + v \geq b^{d^*} (b-1) + v \quad (9)$$

Line 8 merely involves expanding and cancelling out terms. The last line holds by the assumption that $ed^* - 1 \geq 3$ which, alongside the lower bound in Case 2 implies that $g \geq 2b - 1$.

$$y' + z' \geq b^{d^*} (b-1) + (b-1)g - (g+1)g \quad (10)$$

$$+ (b-1)(g+1)g[(e-1)d^* - 1] \quad (11)$$

$$\geq b^{d^*} (b-1) - g \quad (12)$$

$$\geq b^{d^*} (b-1) - b^{d^*} \quad (13)$$

Line 11 simply involves expanding and cancelling out terms. Line 12 holds because $(b-1)(g+1)(e-1) \geq 0$ since $b \geq 2$, $g \geq 1$, and $e \geq 1$. The last line holds since $g < b^{d^*}$.

Since $b \geq 2$, line 13 implies that $x' = y' + z' \geq 0$, which in turn implies $x \geq 0$, which completes the proof. \square

Let us now consider cases not handled by Theorem 4.1. When $d^* = 1$, it can be shown that $\mathbb{E}[B(\mathcal{S}_T)] < \mathbb{E}[R(\mathcal{S}_T)]$ for any $g \leq b^{d^*} = b$, and the two expectations are exactly equal when $g = b$. If $d^* = 2$ and $e = 1$, the situation is almost identical to that in Theorem 4.1 except the lower bound on g in Case 2 is exactly one higher than in Theorem 4.1. Namely, the bound is $g \geq (ed^* - 1)(b-1) + 2$.

Figure 1b shows how the goal crossover changes with goal level and depth error. The bound from Theorem 4.1 is shown as a solid line and the actual crossover points — computed based on Corollaries 4.1 and 4.2 — are shown as dashed lines. The figure clearly depicts that the crossover increases linearly with the goal level and depth error as suggested by Theorem 4.1. The figure also shows that the bound is quite accurate, as it only slightly over-estimates the true crossover for small goal levels.

Finally, we note that while the crossover increases linearly with the goal level, the density of goals needed at the goal level in order for RRW to outperform BrFS actually decreases with d^* . For example, Figure 1c shows the *density crossover*, which we define as the crossover point (determined using the lower bound from Theorem 4.1) divided by the number of goals at the goal level (namely b^{d^*}). The figure clearly shows that the required percentage of goals for RRW to be faster actually decreases with the goal level.

5 Discussion and Future Work

We have shown that BrFS performs fewer goal tests when the number of goals is very small, but RRW is faster in expectation as the number of goals grows even slightly. However, RRW may outperform this result in practice for several reasons. First, RRW can benefit from goals at levels between d^* and ed^* as these will increase the success probability and thereby decrease expected runtime. BrFS cannot benefit from such goals in any way.

Second, while we have focused on the number of goal tests performed, the amount of time needed per goal test differs between the algorithms. This is because BrFS has the additional overhead of maintaining the open and closed list (*ie.* duplicate checking, adding and removing from the open list, etc.). BrFS will also generate all children of a given vertex, not just a single one as RRW does. While this overhead may just be constant time, that may often be enough to further decrease the crossover point when expected runtime and not number of goal tests are considered. The open and closed lists can also have enormous memory requirements making BrFS impractical in certain situations.

In contrast, the duplicate detection performed by open and closed lists means BrFS is much better equipped to handle graphs with *cycles* or *transpositions*. In practice, this effectively means RRW is working on a larger search tree. For example, even if the search task is an undirected tree with constant branching factor b , the duplicate detection of BrFS means it will be able to prune the tree to one with a branching factor $b - 1$. While this particular case is easy to handle with RRW, this algorithm will struggle relative to BrFS in a similar way to iterative deepening methods (Korf 1985) if

the transpositions are more complex than that. Along with our results in Theorem 4.1, we therefore expect RRW to have the advantage in large combinatorial state-spaces with a low number of transpositions and an increasing amount of goals, and BrFS to better handle cases with very few goals or many transpositions. It may therefore be important to develop hybrid BrFS and RRW methods — analogous to A* + IDA* (Bu and Korf 2019) — that get some of the advantages of both. We leave such an investigation as future work.

We now review some related work and identify several avenues of future work. Local best-first searches have been empirically shown to be useful for escaping local minima and plateaus in planning search systems (Xie, Müller, and Holte 2015). Nakhost and Müller (2014) formally analyze the expected runtime of a single random walk and restarting random walks on classes of graphs characterized by the probability of getting closer or farther from a goal on every step. Their model for restarting random walks assumed a constant restart probability at every step instead of a constant restart depth, and they also do not compare to BrFS. Extending our results to their graph classes is left as future work. Everitt and Hutter (2015) performed a similar analysis to ours when comparing BrFS and depth-first search (DFS) on bounded depth-trees. They characterized the expected runtime of BrFS, but their formulation considers the case where every vertex at the goal level is equally likely to be a goal (see Proposition 3). Our use of a fixed number of goals g considerably simplifies the resulting expression for the expected runtime, and we also compare against RRW instead of DFS, which is not generally suitable for escaping local minima/plateaus.

Finally, non-constant restart policies have been introduced that do not rely on an initial guess of the goal depth. These methods dynamically adjust the restart depth over time while still maintaining strong theoretical guarantees. For example, the general policy of Luby, Sinclair, and Zuckerman (1993) has an expected runtime of $O(T \log T)$ on any search task \mathcal{S} , where T is the runtime of the best possible restart policy for \mathcal{T} . As on directed trees, the best possible policy is exactly the constant-depth one with $e = 1$, we hope to extend our analysis to compare BrFS to this general policy. We leave this as future work.

References

Alkhazraji, Y.; Frorath, M.; Grütznert, M.; Helmert, M.; Liebetraut, T.; Mattmüller, R.; Ortlieb, M.; Seipp, J.; Sprinzenberg, T.; Stahl, P.; and Wülfing, J. 2020. Pyperplan. <https://doi.org/10.5281/zenodo.3700819>.

Bu, Z.; and Korf, R. E. 2019. A*+IDA*: A Simple Hybrid Search Algorithm. In Kraus, S., ed., *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019*, 1206–1212. ijcai.org.

Doran, J. E.; and Michie, D. 1966. Experiments with the Graph Traverser Program. *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 294(1437): 235–259.

Everitt, T.; and Hutter, M. 2015. Analytical Results on

the BFS vs. DFS Algorithm Selection Problem. Part I: Tree Search. In Pfahringer, B.; and Renz, J., eds., *AI 2015: Advances in Artificial Intelligence - 28th Australasian Joint Conference, Canberra, ACT, Australia, November 30 - December 4, 2015, Proceedings*, volume 9457 of *Lecture Notes in Computer Science*, 157–165. Springer.

Felner, A. 2011. Position Paper: Dijkstra’s Algorithm versus Uniform Cost Search or a Case Against Dijkstra’s Algorithm. In Borrajo, D.; Likhachev, M.; and López, C. L., eds., *Proceedings of the Fourth Annual Symposium on Combinatorial Search, SOCS 2011, Castell de Cardona, Barcelona, Spain, July 15.16, 2011*, 47–51. AAAI Press.

Hoffmann, J. 2001. FF: The Fast-Forward Planning System. *AI Magazine*, 22(3): 57.

Korf, R. E. 1985. Depth-First Iterative-Deepening: An Optimal Admissible Tree Search. *Artificial Intelligence*, 27(1): 97–109.

Luby, M.; Sinclair, A.; and Zuckerman, D. 1993. Optimal Speedup of Las Vegas Algorithms. In *Second Israel Symposium on Theory of Computing Systems, ISTCS 1993, Natanya, Israel, June 7-9, 1993, Proceedings*, 128–133. IEEE Computer Society.

Nakhost, H.; and Müller, M. 2009. Monte-Carlo Exploration for Deterministic Planning. In Boutilier, C., ed., *IJCAI 2009, Proceedings of the 21st International Joint Conference on Artificial Intelligence, Pasadena, California, USA, July 11-17, 2009*, 1766–1771.

Nakhost, H.; and Müller, M. 2014. Towards a theory of random walk planning: Regress factors, fair homogeneous graphs and extensions. *AI Communications*, 27(4): 329–344.

Xie, F.; Müller, M.; and Holte, R. 2014. Adding Local Exploration to Greedy Best-First Search in Satisficing Planning. In Brodley, C. E.; and Stone, P., eds., *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence, July 27 -31, 2014, Québec City, Québec, Canada*, 2388–2394. AAAI Press.

Xie, F.; Müller, M.; and Holte, R. 2015. Understanding and Improving Local Exploration for GBFS. In Brafman, R. I.; Domshlak, C.; Haslum, P.; and Zilberstein, S., eds., *Proceedings of the Twenty-Fifth International Conference on Automated Planning and Scheduling, ICAPS 2015, Jerusalem, Israel, June 7-11, 2015*, 244–248. AAAI Press.