

Combining Monte Carlo Tree Search and Heuristic Search for Weighted Vertex Coloring

Cyril Grelier¹, Olivier Goudet¹ and Jin-Kao Hao^{1*}

¹LERIA, Université d'Angers, 2 Boulevard Lavoisier, Angers,
49045, France.

*Corresponding author(s). E-mail(s): jin-cao.hao@univ-angers.fr;
Contributing authors: cyril.grelier@univ-angers.fr;
olivier.goudet@univ-angers.fr;

Abstract

This work investigates the Monte Carlo Tree Search (MCTS) method combined with dedicated heuristics for solving the Weighted Vertex Coloring Problem. In addition to the basic MCTS algorithm, we study several MCTS variants where the conventional random simulation is replaced by other simulation strategies including greedy and local search heuristics. We conduct experiments on well-known benchmark instances to assess these combined MCTS variants. We provide empirical evidence to shed light on the advantages and limits of each simulation strategy. This is an extension of the work [1] presented at EvoCOP2022.

Keywords: Monte Carlo Tree Search, local search, graph coloring, weighted vertex coloring

1 Introduction

The well-known Graph Coloring Problem (GCP) is to color the vertices of a graph using as few colors as possible such that no adjacent vertices share the same color (*legal* or *feasible* solution). The GCP can also be considered as partitioning the vertex set of the graph into a minimum number of color groups such that no vertices in each color group are adjacent. The GCP has numerous practical applications in various domains [2] and has been studied

intensively since the 19th century in mathematics and for over 50 years in computer science.

The Weighted Vertex Coloring Problem (WVCP), a variant of the GCP, has recently attracted much interest in the literature [3–6]. In this problem, each vertex of the graph has a weight and the objective is to find a *legal* coloring such that the sum of the weights of the heaviest vertex of each color group is minimized. Formally, given a weighted graph $G = (V, E)$ with vertex set V ($n = |V|$) and edge set E , and let W be the set of weights $w(v)$ associated to each vertex v in V , the WVCP consists in finding a partition of the vertices in V into k color groups $S = \{V_1, \dots, V_k\}$ ($1 \leq k \leq n$) such that no adjacent vertices belong to the same color group and such that the score $\sum_{i=1}^k \max_{v \in V_i} w(v)$ is minimized. Note that the value of k is not predetermined for a WVCP instance and may vary during the search as a solution with more colors may have a better score than a solution with less colors. One can notice that when all the weights $w(v)$ ($v \in V$) are equal to one, finding an optimal solution of this problem with a minimum score corresponds to solving the GCP. The WVCP can be seen as a more general problem than the GCP and is therefore NP-hard.

The WVCP is a relevant model for several applications such as matrix decomposition [7], buffer size management, and scheduling of jobs into batches in a multiprocessor environment [8]. Let us consider the last application as illustrated in Figure 1. The objective of this scheduling problem is to execute a set of jobs in a minimum total amount of time. There is no constraint on the number of jobs that can be run in parallel in this environment. However, each job requires a specific execution time and exclusive access to certain resources. Therefore, the time required to complete a batch of jobs in parallel is the time required to complete the longest job in that batch, and two jobs requiring the same resource cannot be launched in the same batch. Solving this problem within the WVCP modeling framework can be done in five steps as displayed in Figure 1: (i) a bipartite graph is used to represent the jobs and the resources required for each job; (ii) this bipartite graph is projected onto the resources to obtain a weighted graph where each vertex is a job and two jobs requiring the same resources are linked by an edge; (iii) a weight corresponding to the time needed to complete a job is set on the corresponding vertex of this graph; (iv) after solving the WVCP associated to this graph, a legal solution is found with an optimal score of 25, corresponding to the sum of the weights of the heaviest vertex of each color group; (v) this partition of vertices allows to set up a job schedule in four batches, which respects the resource constraints, and whose minimum total execution time is 25 seconds.

Different methods have been proposed in the literature to solve the WVCP. First, this problem has been tackled with exact methods: a branch-and-price algorithm [9], two ILP models proposed in [10] and [11] with a transformation of the WVCP into a maximum weight independent set problem, and constraint programming in [12]. These exact methods can prove the optimality on small instances but tends to fail on graphs with more than 250 vertices.

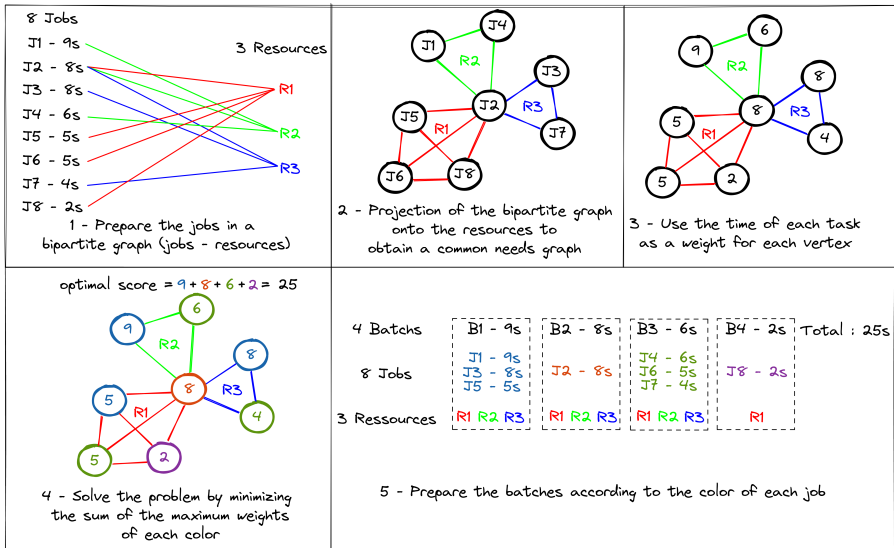


Fig. 1 This figure shows an application of the WVCP for scheduling jobs into batches in a multiprocessor environment with restricted access to certain resources.

To handle large graphs, several heuristics have been introduced to solve the problem approximately [4–7]. The first category of heuristics is based on the local search framework, which iteratively makes transitions from the current solution to a neighbor solution. Three different approaches have been considered to explore the search space: legal, partial legal, or penalty strategies. The legal strategy starts from a *legal* solution and minimizes the score by performing only legal moves so that no color conflict is created in the new solution [7]. The partial legal strategy allows only legal coloring and keeps a set of uncolored vertices to avoid conflicts [4]. The penalty strategy considers both legal and illegal solutions in the search space [5, 6], and uses a weighted evaluation function to minimize both the WVCP objective function and the number of conflicts in the illegal solutions. To escape local optima traps, these local search algorithms incorporate different mechanisms such as perturbation strategies [5, 7], tabu list [4, 5] and constraint reweighting schemes [6].

The second category of existing heuristics for the WVCP relies on the population-based memetic framework that combines local search with crossovers. The DLMCOL algorithm [3] of this category uses a deep neural network to learn an invariant by color permutation regression model, useful to select the most promising crossovers at each generation. The AHEAD algorithm [13] involves automatic selection of crossover and local search operators for the WVCP and GCP.

Research on combining such learning techniques and heuristics has received increasing attention in the past years for graph coloring problems [14, 15]. In these new frameworks, after each search trajectory, a matrix that specifies the

96 probability of a vertex belonging to each color group is updated. This matrix
97 is then used to guide the local search algorithm for subsequent iterations.

98 This study continues in that vein and investigates the potential benefits of
99 combining Monte Carlo Tree Search (MCTS) and sequential coloring or local
100 search algorithms for solving the WVCP. MCTS is a heuristic search algorithm
101 that generated considerable interest due to its spectacular success for the game
102 of Go [16], and in other domains (see the survey [17] on this topic). It has
103 been recently revisited in combination with modern deep learning techniques
104 for difficult two-player games (cf. AlphaGo [18]). MCTS has also been applied
105 to combinatorial optimization problems seen as a one-player game such as the
106 traveling salesman problem [19] or the knapsack problem [20]. An algorithm
107 based on MCTS has recently been implemented with some success for the GCP
108 in [21]. In this work, we investigate for the first time the MCTS approach for
109 solving the WVCP.

110 In MCTS, a tree is built incrementally and asymmetrically. For each iter-
111 ation, a tree policy balancing exploration and exploitation is used to find the
112 most critical node to expand. A simulation is then run from the expanded node
113 and the search tree is updated with the result of this simulation. Its incre-
114 mental and asymmetric properties make MCTS a promising candidate for the
115 WVCP because in this problem only the heaviest vertex of each color group
116 has an impact on the objective score. Therefore learning to color the heaviest
117 vertices of the graph before coloring the rest of the graph seems particularly
118 relevant for this problem.

119 Unlike backtracking algorithms, the main aim of MCTS is not to exhaust-
120 ively test all solutions as fast as possible. Instead, MCTS prioritizes the most
121 promising branches of the search tree using a heuristic that balances exploita-
122 tion and exploration. This approach can potentially find high-quality solutions
123 faster than backtracking algorithms. However, MCTS can revisit the same
124 solution multiple times during the search process. Nevertheless, as explained
125 later, MCTS, like backtracking algorithms, can provide a proof of optimality
126 for any given instance if it is given sufficient time.

127 The contributions of this work are summarized as follows.

128 First, we present a MCTS algorithm dedicated to the WVCP, which consid-
129 ers the problem from the perspective of sequential coloring with a predefined
130 vertex order. The exploration of the tree is accelerated with the use of spe-
131 cific pruning rules, which offer the possibility to explore the whole tree in
132 a reasonable amount of time for small instances and to obtain optimality
133 proofs. Secondly, for large instances, when obtaining an exact result is impos-
134 sible in a reasonable time, we study how this MCTS algorithm can be tightly
135 coupled with other heuristics. Specifically, we investigate the integration of dif-
136 ferent greedy coloring strategies and local search procedures within the MCTS
137 algorithm.

138 The rest of the paper is organized as follows. Section 2 introduces the
139 weighted vertex coloring problem and the constructive approach with a tree.
140 Section 3 describes the MCTS algorithm devised to tackle the problem. Section

141 4 presents the coupling of MCTS with local search. Section 5 reports com-
 142 putational results of different versions of MCTS. Section 6 discusses the
 143 contributions and presents research perspectives.

144 2 Constructive approach with a tree for the 145 weighted graph coloring problem

146 This section presents a tree-based approach for the WVCP, which aims to
 147 explore the partial and legal search space of this problem.

148 2.1 Partial and legal search space

149 The search space Ω studied in our algorithm concerns legal, but potentially
 150 partial, k -colorings. A partial legal k -coloring S is a partition of the set of
 151 vertices V into k disjoint independent sets V_i ($1 \leq i \leq k$), and a set of uncolored
 152 vertices $U = V \setminus \bigcup_{i=1}^k V_i$. A independent set V_i is a set of mutually non adjacent
 153 vertices of the graph: $\forall u, v \in V_i, (u, v) \notin E$. For the WVCP, the number of
 154 colors k that can be used is not known in advance. Nevertheless, it is not lower
 155 than the chromatic number of the graph $\chi(G)$ and not greater than the number
 156 of vertices n of the graph. A solution of the WVCP is denoted as partial if $U \neq$
 157 \emptyset and complete otherwise. The objective of the WVCP is to find a complete
 158 solution S with a minimum score $f(S)$ given by: $f(S) = \sum_{i=1}^k \max_{v \in V_i} w(v)$.

159 2.2 Tree search for weighted vertex coloring

160 Backtracking-based tree search is a popular approach for the graph coloring
 161 problem [2, 22, 23]. In our case, a tree search algorithm can be used to explore
 162 the partial and legal search space of the WVCP previously defined.

163 Starting from a solution where no vertex is colored (i.e., $U = V$) and that
 164 corresponds to the root node R of the tree, child nodes C are successively
 165 selected in the tree, consisting of coloring one new vertex at a time. This process
 166 is repeated until a terminal node T is reached (all the vertices are colored).
 167 A complete solution (i.e., a legal coloring) corresponds thus to a branch from
 168 the root node to a terminal node. The maximum depth of the tree is n , the
 169 number of vertices in the graph.

170 The selection of each child node corresponds to applying a move to the
 171 current partial solution being constructed. A move consists of assigning a
 172 particular color i to an uncolored vertex $u \in U$, denoted as $\langle u, U, V_i \rangle$.
 173 Applying a move to the current partial solution S , results in a new solution
 174 $S \oplus \langle u, U, V_i \rangle$. This tree search algorithm only considers legal moves to stay
 175 in the partial legal space. For a partial solution $S = \{V_1, \dots, V_k, U\}$, a move
 176 $\langle u, U, V_i \rangle$ is said legal if no vertex of V_i is adjacent to the vertex u . At each
 177 level of the tree, there is at least one possible legal move that applies to a ver-
 178 tex a new color that has never been used before (or putting this vertex in a
 179 new empty set V_i , $k + 1 \leq i \leq n$).

180 Applying a succession of n legal moves from the initial solution results in
 181 a legal coloring of the WVCP and reaches a terminal node of the tree. During
 182 this process, at the level t of the tree ($0 \leq t < n$), the current legal and
 183 partial solution $S = \{V_1, \dots, V_k, U\}$ has already used k colors and t vertices
 184 have already received a color. Therefore $|U| = n - t$.

185 At this level, a first naive approach could be to consider all the possible
 186 legal moves, corresponding to choosing a vertex in the set U and assigning to
 187 the vertex a color i , with $1 \leq i \leq n$. This kind of choice can work with small
 188 graphs but with large graphs, the number of possible legal moves becomes
 189 huge. Indeed, at each level t , the number of possible legal moves can go up to
 190 $(n - t) \times n$.

191 To reduce the set of move possibilities, we consider the vertices of the graph
 192 in a predefined order (u_1, \dots, u_n) . When choosing a color for the next vertex,
 193 we only consider the colors already used in the partial solution plus one new
 194 color, as long as this number is less than the vertex degree plus one. This
 195 approach bounds the number of colors needed for a vertex [24]. Thus, for the
 196 current partial and legal solution $S = \{V_1, \dots, V_k, U\}$, at most $d(u) + 1$ moves
 197 are considered for the next vertex u to be colored. If k colors have already
 198 been used, the set of legal moves to place u is, if $k < d(u) + 1$

$$\mathcal{L}(S) = \{ \langle u, U, V_i \rangle, 1 \leq i \leq k, \forall v \in V_i, (u, v) \notin E \} \cup \{ \langle u, U, V_{k+1} \rangle \}, \quad (1)$$

199 or, if $k \geq d(u) + 1$

$$\mathcal{L}(S) = \{ \langle u, U, V_i \rangle, 1 \leq i \leq d(u) + 1, \forall v \in V_i, (u, v) \notin E \} \quad (2)$$

200 This decision cuts symmetries in the tree while reducing the number of branching
 201 factors at each level of the tree. The potential number of leaf nodes in the
 202 tree is, in the worst case, equal to $\prod_{i=1, \dots, n} \min(i, d_{u_i} + 1)$.

203 2.3 Predefined vertex order

204 We propose to consider a predefined ordering of the vertices, sorted by weight
 205 and then by degree. Vertices with higher weights are placed first. If two vertices
 206 have the same weight, then the vertex with the higher degree is placed first.
 207 This order is intuitively relevant for the WVCP because it is more important
 208 to place first the vertices with heavy weights which have the most impact on
 209 the score as well as the vertices with the highest degree because they are the
 210 most constrained decision variables. Such ordering has already been shown
 211 to be effective with greedy constructive approaches for the GCP [22] and the
 212 WVCP [4].

213 Moreover, this vertex ordering allows a simple score calculation while build-
 214 ing the tree. Indeed, as the vertices are sorted by descending order of their
 215 weights, and the score of the WVCP only counts the maximum weight of each
 216 color group, with this vertex order, the score only increases by the value $w(v)$
 217 when a new color group is created for the vertex v .

3 Monte Carlo Tree Search for weighted vertex coloring

The search tree presented in the last subsection can be huge, in particular for large instances. Therefore, in practice, it is often impossible to perform an exhaustive search of this tree, due to expensive computing time and memory requirements. We turn now to an adaptation of the MCTS algorithm for the WVCP to explore this search tree. MCTS keeps in memory a tree (hereinafter referred to as the MCTS tree) that only corresponds to the already explored nodes of the search tree presented in the last subsection. In the MCTS tree, a leaf is a node whose children have not yet all been explored while a terminal node corresponds to a complete solution. MCTS can guide the search toward the most promising branches of the tree, by balancing exploitation and exploration and continuously learning at each iteration.

3.1 General framework

The MCTS algorithm for the WVCP is shown in Algorithm 1. The algorithm takes a weighted graph as input and tries to find a legal coloring S with the minimum score $f(S)$. The algorithm starts with an initial solution where the first vertex is placed in the first color group. This is the root node of the MCTS tree. Then, the algorithm repeats several iterations until a stopping criterion is met. At every iteration, one legal solution is completely built, which corresponds to walking along a path from the root node to a leaf node of the MCTS tree and performing a simulation (or playout/rollout) until a terminal node of the search tree is reached (when all vertices are colored).

Each iteration of the MCTS algorithm involves the execution of 5 steps to explore the search tree with legal moves (cf. Section 2):

1. **Selection** From the root node of the MCTS tree, successive child nodes are selected until a leaf node is reached. The selection process balances the exploration-exploitation trade-off. The exploitation score is linked to the average score obtained after having selected this child node and is used to guide the algorithm to a part of the tree where the scores are the lowest (the WVCP is a minimization problem). The exploration score is linked to the number of visits to the child node and will incite the algorithm to explore new parts of the tree, which have not yet been explored.
2. **Expansion** The MCTS tree grows by adding a new child node to the leaf node reached during the selection phase.
3. **Simulation** From the newly added node, the current partial solution is completed with legal moves, randomly or by using heuristics.
4. **Update** After the simulation, the average score and the number of visits of each node on the explored branch are updated.
5. **Pruning** If a new best score is found, some branches of the MCTS tree may be pruned if it is not possible to improve the best current score with it.

The algorithm continues until one of the following conditions is reached:

Algorithm 1 MCTS algorithm for the WVCP

```

1: Input: Weighted graph  $G = (V, W, E)$ 
2: Output: The best legal coloring  $S^*$  found
3:  $S^* = \emptyset$  and  $f(S^*) = MaxInt$ 
4: while stop condition is not met do
5:    $C \leftarrow R$   $\triangleright$  Current node corresponding to the root node of the tree
6:    $S \leftarrow \{V_1, U\}$  with  $V_1 = \{v_1\}$  and  $U = V \setminus V_1$   $\triangleright$  Current solution
   initialized with the first vertex in the first color group
7:   /* Selection */  $\triangleright$  Section 3.2
8:   while  $C$  is not a leaf do
9:      $C \leftarrow \text{select\_best\_child}(C)$  with legal move  $\langle u, U, V_i \rangle$ 
10:     $S \leftarrow S \oplus \langle u, U, V_i \rangle$ 
11:   end while
12:   /* Expansion */  $\triangleright$  Section 3.3
13:   if  $C$  has a potential child, not yet open then
14:      $C \leftarrow \text{open\_first\_child\_not\_open}(C)$  with legal move  $\langle u, U, V_i \rangle$ 
15:      $S \leftarrow S \oplus \langle u, U, V_i \rangle$ 
16:   end if
17:   /* Simulation */  $\triangleright$  Section 3.4
18:   complete_partial_solution( $S$ )
19:   /* Update */  $\triangleright$  Section 3.5
20:   while  $C \neq R$  do
21:     update( $C, f(S)$ )
22:      $C \leftarrow \text{parent}(C)$ 
23:   end while
24:   if  $f(S) < f(S^*)$  then
25:      $S^* \leftarrow S$ 
26:     /* Pruning */  $\triangleright$  Section 3.6
27:     apply pruning rules
28:   end if
29: end while
30: return  $S^*$ 

```

- 260 • there are no more child nodes to expand, meaning the search tree has been
261 fully explored. In this case, the best score found is proven to be optimal.
262 • a cutoff time is attained. The minimum score found so far is returned. It
263 corresponds to an upper bound of the optimal score for the given instance.

3.2 Selection

264
265 The selection starts from the root node of the MCTS tree and selects children
266 nodes until a leaf node is reached. At every level t of the MCTS tree, if the
267 current node C_t corresponds to a partial solution $S = \{V_1, \dots, V_k, U\}$ with t
268 vertices already colored and k colors used, there are l possible legal moves, with

269 $1 \leq l \leq k + 1$. Therefore, from the node C_t , l potential children $C_{t+1}^1, \dots, C_{t+1}^l$
 270 can be selected.

If $l > 1$, the selection of the most promising child node can be seen as a multi-armed bandit problem [25] with l levers. The problem of choosing the next node can be solved with the UCT algorithm for Monte Carlo tree search by selecting the child with the maximum value of the following expression [20]:

$$\text{normalized_score}(C_{t+1}^i) + c \times \sqrt{\frac{2 * \ln(\text{nb_visits}(C_t))}{\text{nb_visits}(C_{t+1}^i)}}, \text{ for } 1 \leq i \leq l. \quad (3)$$

Here, $\text{nb_visits}(C)$ corresponds to the number of times the node C has been chosen to build a solution. c is a real positive coefficient allowing to balance the compromise between exploitation and exploration, which is set by default to one.¹ $\text{normalized_score}(C_{t+1}^i)$ corresponds to a normalized score of the child node C_{t+1}^i ($1 \leq i \leq l$) given by:

$$\text{normalized_score}(C_{t+1}^i) = \frac{\text{rank}(C_{t+1}^i)}{\sum_{i=1}^l \text{rank}(C_{t+1}^i)}$$

271 where $\text{rank}(C_{t+1}^i)$ is defined as the rank between 1 and l of the nodes
 272 C_{t+1}^i obtained by sorting from bad to good according to their average val-
 273 ues $\text{avg_score}(C_{t+1}^i)$ (nodes that seem more promising get a higher score).
 274 $\text{avg_score}(C_{t+1}^i)$ is the mean score on the sub-branch with the node C_{t+1}^i
 275 selected obtained after all previous simulations.

276 Equation 3 can be compared with the probabilities of choice at each level
 277 of the tree by an Ant Colony Optimisation (ACO) algorithm, except that in
 278 our case, the choice of the child node is deterministic, and explicitly involves
 279 a trade-off between an intensification term, the normalized score of each child
 280 node, and an exploration term, depending on its number of visits.

281 3.3 Expansion

282 From the node C of the MCTS tree reached during the selection procedure,
 283 one new child of C is open and its corresponding legal move is applied to the
 284 current solution. Among the unopened children, the node associated with the
 285 lowest color number i is selected. Therefore the child node needing the creation
 286 of a new color (and increasing the score) will be selected last.

287 3.4 Simulation

288 The simulation takes the current partial and legal solution found after the
 289 expansion phase and colors the remaining vertices. In the original MCTS algo-
 290 rithm, the simulation consists in choosing random moves in the set of all legal

¹A sensitivity analysis of this important hyperparameter is shown in Section 5.3.

291 moves $\mathcal{L}(S)$, defined by Equation 1 and Equation 2, until the solution is com-
 292 pleted. We call this first version *MCTS+Random* (MCTS+R). As shown in
 293 the experimental section, this version is not very efficient as the number of
 294 colors grows rapidly. Therefore, we propose three other simulation procedures:

- 295 • a constrained greedy algorithm that chooses a legal move prioritizing the
 296 moves which do not locally increase the score of the current partial solution.
 297 The move applied for each vertex is randomly selected among the legal moves
 298 in $\mathcal{L}(S)$. It only chooses the move $\langle u, U, V_{k+1} \rangle$, consisting in opening a
 299 new color group and increasing the current score by $w(u)$, only if $\mathcal{L}^g(S) = \emptyset$.
 300 We call this version *MCTS+Greedy-Random* (MCTS+GR).
- 301 • a greedy deterministic procedure which always chooses a legal move in
 302 $\mathcal{L}(S)$ with the first available color i . We call this version *MCTS+Greedy*
 303 (MCTS+G).
- 304 • a greedy deterministic procedure based on DSatur [22], which colors in pri-
 305 ority the heaviest and most saturated vertices. The saturation of a vertex
 306 is the number of colors used by its adjacent vertices. We call this version
 307 *MCTS+DSatur* (MCTS+DS).

308 3.5 Update

309 Once the simulation is over, a complete solution S of the WVCP is obtained.
 310 If this solution is better than the best recorded solution found so far S^* (i.e.,
 311 $f(S) < f(S^*)$), S becomes the new global best solution S^* .

312 Then, a backpropagation procedure updates each node C of the whole
 313 branch of the MCTS tree which has led to this solution:

- the running average score of each node C of the branch is updated with the
 score $f(S)$:

$$avg_score(C) \leftarrow \frac{avg_score(C) \times nb_visits(C) + f(S)}{nb_visits(C) + 1} \quad (4)$$

- 314 • the counter of visits $nb_visits(C)$ of each node of the branch is increased by
 315 one.

316 3.6 Pruning

317 During an iteration of MCTS, three pruning rules are applied:

- 318 1. During expansion, if the score $f(S)$ of the partial solution associated with
 319 a node visited during this iteration of MCTS is equal to or higher than the
 320 current best-found score $f(S^*)$, the node is deleted as the score of such a
 321 partial solution cannot decrease when more vertices are colored.
- 322 2. When the best score $f(S^*)$ is found, the tree is *cleaned*. A heuristic goes
 323 through the whole tree and deletes children and possible children associated
 324 with a partial score $f(S)$ equal or superior to the best score $f(S^*)$.

325 3. If a node is *completely explored*, it is deleted and will not be explored in the
 326 MCTS tree anymore. A node is said *completely explored* if it is a leaf node
 327 without children, or if all of its children have already been opened once
 328 and have all been deleted. Note that this third pruning step is recursive as
 329 a node deletion can result in the deletion of its parent if it has no more
 330 children, and so on.

331 These three pruning rules and the fact that the symmetries are cut in the
 332 tree by restricting the set of legal moves considered at each step (see Section
 333 2.2) offer the possibility to explore the whole tree in a reasonable amount of
 334 time for small instances. This peculiarity of the algorithm makes it possible to
 335 obtain an optimality proof for such instances.

336 3.7 Toy example

337 Figure 2 displays one iteration of MCTS for the WVCP on a small graph
 338 composed of seven vertices named A–G with different weights between 2 and 9.
 339 On each diagram is displayed the current state of the partial coloring solution
 340 being constructed (right) and the current state of the search tree (left). In the
 341 search tree, each square represents a node and the number on the bottom right
 342 of a square is the score of the corresponding partial solution. On top of each
 343 square are written the average score and the number of visits of each node. In

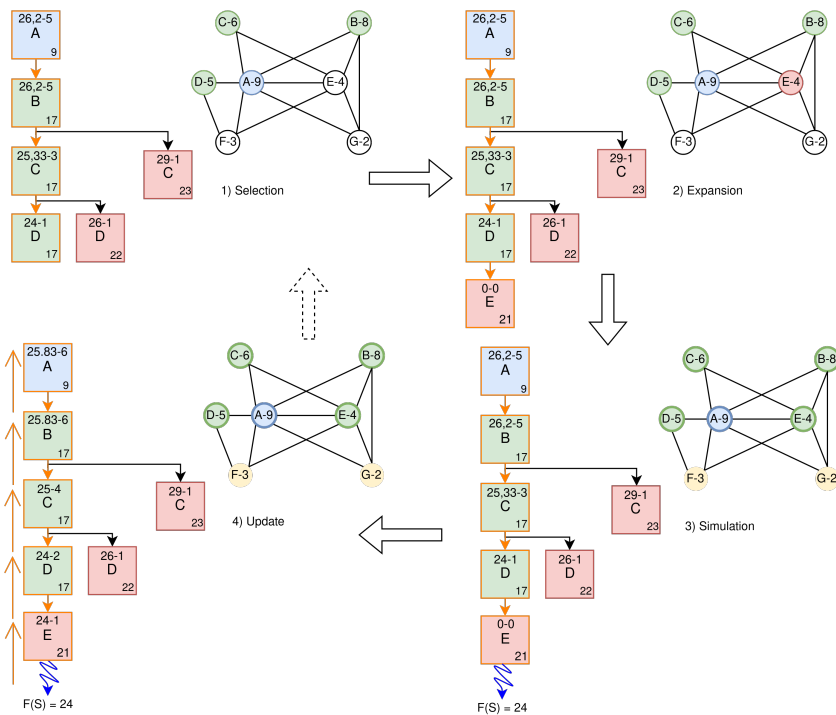


Fig. 2 Toy example of one MCTS iteration

344 addition to the root node (vertex A colored in blue), five nodes have already
 345 been opened in the search tree (five iterations of MCTS). The sixth iteration
 346 of MCTS proceeds as follows.

- 347 • **Selection:** From the root node, the only possible child corresponding to
 348 vertex B in green is selected. From there, there are two options as vertex C
 349 can be colored in green or red. The most interesting option is chosen (vertex
 350 C in green) regarding the score and the number of visits of each child (cf.
 351 equation (3)). Then, the most promising leaf is selected (D in green).
- 352 • **Expansion:** From the node D in green, a new node is added to the tree. It
 353 corresponds to E in red (as it cannot take the color blue nor green).
- 354 • **Simulation:** From there, the solution is completed with a greedy algorithm
 355 to obtain a complete legal solution with a score of 24.
- 356 • **Update:** This score of 24 is back-propagated on the explored branch (update
 357 of the average score and the number of visits of each node in the branch).
- 358 • **Pruning:** Figure 3 presents the state of the tree after some iterations. As
 359 the best-found score is 24, every branch of the tree with a score greater than
 360 or equal to 24 is deleted (indicated with a red cross).

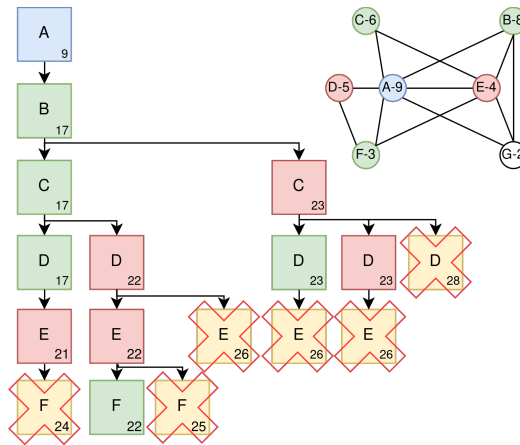


Fig. 3 Toy example of the search tree pruning

361 4 Combining MCTS with Local Search

362 We now explore the possibility of improving the MCTS algorithm with local
 363 search. Coupling MCTS with a local search algorithm is motivated by the fact
 364 that after the simulation phase, the complete solution obtained can be close
 365 to a still better solution in the search space that could be discovered by local
 366 search. In this work, we present the coupling of MCTS with a baseline tabu

367 search (TW) created for this work, as well as three *state of the art* local search
368 algorithms, dedicated for the WVCP: AFISA [5], RedLS [6] and ILS-TS [4].

369 During the simulation phase, the solution is first completed with a greedy
370 algorithm and then improved by the local search procedure. Note that in the
371 first version of this work published in [1], to stay consistent with the search
372 tree learned by MCTS, we allowed the local search procedure to only move the
373 vertices of the complete solution S which are still uncolored after the selection
374 and expansion phases. However, we have realized in the meantime that blocking
375 vertices for the local search can lead to a lot of time spent checking for blocked
376 vertices in the complex neighborhood explored by the various local search
377 procedures. It also leads to missing good opportunities to move in the search
378 space. Therefore, in the new version of the algorithm presented in this paper,
379 a more efficient version of the algorithm is presented where the vertices are not
380 frozen during the local search. In this new version, when coupling MCTS with
381 a local search algorithm, the resulting heuristic can be seen as an algorithm
382 that attempts to learn a good starting point for the local search procedure,
383 by selecting different best promising backbones of partial solutions in every
384 iteration during the selection phase.

385 As one iteration does not have the same meaning for each local search, we
386 use a time limit of $t = 0.02 \times n$ seconds to perform the search, depending on
387 the number of vertices n in the given instance. Once the local search procedure
388 has reached the time limit, the score corresponding to the best legal solution
389 obtained by the local search procedure is used to update all the nodes of the
390 branch which has led to the simulation initiation. In the following subsections,
391 the four different local search procedures used in this work are presented.

392 4.1 Basic tabu search

393 The first local search algorithm tested is a simple tabu search, named
394 TabuWeight (TW), inspired by the classical TabuCol algorithm for the GCP
395 [26]. Starting from a legal solution, TW improves it iteratively by using the
396 *one move* operator, which consists in moving a vertex from its color group to
397 another color group, without creating conflicts. At each iteration, the best *one*
398 *move* which is not forbidden by the tabu list is selected. Each time a move is
399 performed, the reverse move is added to the tabu list and forbidden for the
400 next tt iterations where tt is a parameter called tabu tenure. A tabu move can
401 still be applied exceptionally if it leads to a solution, which is better than the
402 best solution found so far (aspiration criterion).

403 4.2 Adaptive feasible and infeasible tabu search

404 AFISA [5] is an advanced tabu search algorithm, which explores the candidate
405 solutions by oscillating between illegal and legal search spaces². To prevent
406 the search from going too far from legal boundaries, AFISA uses a controlling

²The illegal search space consists of solutions with conflicts (some adjacent vertices in the solution have the same color), while the legal search space consists of solutions without any conflicts.

407 coefficient to adaptively make the algorithm go back and forth between illegal
 408 and legal spaces. The controlling coefficient encourages the algorithm to handle
 409 in priority the vertices in conflicts before trying to reduce the WVCP score.
 410 AFISA uses the popular *one move* operator to explore the search space.

411 **4.3 Local search with multiple operators**

412 Like AFISA, RedLS [6] explores the illegal and legal search spaces. This
 413 algorithm uses the configuration checking strategy [27] that applies multiple
 414 improvements and perturbation strategies to explore the search space. At each
 415 iteration, RedLS perturbs the solution by moving all the heaviest vertices from
 416 one color group to another group, before minimizing the number of conflicts
 417 to recover a new legal solution. It uses different variants of the *one move* oper-
 418 ator to reduce the number of conflicts while keeping the WVCP score as low
 419 as possible. Each conflicting edge has a weight that is increased each time it
 420 is not resolved, to give priority to its resolution for the next iterations.

421 **4.4 Iterated local search with tabu search**

422 ILS-TS [4] explores the legal and partial search spaces. From a complete
 423 solution, the ILS-TS algorithm iteratively performs 2 steps: (i) it deletes the
 424 heaviest vertices from 1 to 3 color groups V_i and places them in the set of uncol-
 425 ored vertices U ; (ii) it improves the solution (i.e., minimizes the score $f(S)$) by
 426 applying different variants of the *one move* operator and the so-called *grenade*
 427 operator until the set of uncolored vertices U becomes empty. The *grenade*
 428 operator $grenade(u, V_i)$ moves a vertex u to V_i and relocates each adjacent
 429 vertex of u in V_i to another color group or in U to keep a legal solution.

430 **5 Experimentation**

431 This section first describes the experimental settings used in this work. Sec-
 432 ondly, we experimentally verify the impacts of the different greedy coloring
 433 strategies used during the MCTS simulation phase. Thirdly, an analysis of
 434 exploration versus exploitation is performed. Lastly, the relevance of coupling
 435 MCTS with a local search procedure is studied.

436 **5.1 Experimental settings and benchmark instances**

437 A total of 188 instances are used for the experimental studies: 30 rxx graphs
 438 and 35 pxx graphs from matrix decomposition [7] and 123 from the DIMACS
 439 and COLOR competitions. The instances are used in a reduced version
 440 presented in [12]. The original and reduced instances are available online.³

441 All presented algorithms are coded in C++, compiled, and optimized with
 442 the g++ 12.1 compiler. Differences in the results may occur compared to [1]
 443 as some optimisations have been done and more reduced instances have been

³https://github.com/Cyril-Grelier/gc_instances

444 used. The source code of our algorithm (and reproduced local searches) is avail-
 445 able online⁴ with complete spreadsheets of the results. To solve each instance,
 446 20 independent runs were performed on a computer equipped with an Intel
 447 Xeon ES 2630, 2,66 GHz CPU with a time limit of one hour, except for the
 448 exploration vs. exploitation coefficient tests where 5 to 15 hours were used.
 449 Running the DIMACS Machine Benchmark procedure dfmax⁵ on our com-
 450 puter took 8.94 seconds to solve the instance r500.5 using gcc 12.1 without
 451 optimization flag.

452 In the following subsections, summary tables allowing general comparisons
 453 between the different versions of the algorithms are presented. Detailed results
 454 on each specific instance are reported with the source code.⁴

455 The 188 instances have been separated into four sets: (i) **pxx**, with the
 456 35 pxx instances from [7], (ii) **rxx**, with the 30 rxx instances from [7],
 457 (iii) **DIMACS_easy**, corresponding to the *easy* 75 DIMACS and COLOL
 458 instances with 72 among them which were solved optimally by exact algo-
 459 rithms [11, 12], and (iv) **DIMACS_hard**, the 48 *hard* DIMACS instances
 460 which have never been solved optimally in the literature, except for 5 which
 461 are really difficult to solve.

462 For all the different versions of the MCTS algorithm, the coefficient c ,
 463 allowing to balance the compromise between exploitation and exploration is set
 464 to the value of one (cf. equation (3)). A sensitivity analysis of this important
 465 hyperparameter is conducted in Section 5.3.

466 5.2 Monte Carlo Tree Search with greedy strategies

467 Table 1 summaries the results of MCTS with greedy heuristics for its
 468 simulation (cf. Section 3.4).

469 Columns 2-5 present results for four instance sets: pxx, rxx, DIMACS_easy,
 470 and DIMACS_hard. Column 6 shows global results for all 188 benchmark
 471 instances. The table header lists the set name, number of instances, and num-
 472 ber of instances with proven optimality (marked with a star). For each method
 473 in each column, two numbers are given: the number of instances where the
 474 method matches the Best Known Scores (BKS) from the literature⁶, and the
 475 number of instances solved to proven optimality (marked with a star).

476 First, we observe that all the MCTS variants dominate the baseline MCTS
 477 with greedy simulation in terms of the number of BKS obtained, highlighting
 478 the relevance of combining the MCTS framework and search heuristics.

479 With the MCTS variants, almost all the pxx instances are optimally
 480 solved. The rxx instances are more difficult to solve except for the ver-
 481 sion MCTS+Greedy-Random. The instances from DIMACS_easy are partially
 482 solved by each MCTS variant. The instances from DIMACS_hard show a real
 483 challenge for all the MCTS variants.

⁴<https://github.com/Cyril-Grelier/gc-wvcp-adaptive-mcts>

⁵<http://archive.dimacs.rutgers.edu/pub/dsj/cliique/>

⁶Note that some of these BKS have been found in the literature with extended search time from several hours to several days, while our methods are only run during one hour.

Table 1 Summary of the number of times the Best Known Score (BKS) is reached by each algorithm. The values with a star indicate the number of times a score has been proved optimal. Values in bold highlight the best results for each line.

	pxx 35*	rxr 30*	DIMACS easy 75 - 72*	DIMACS hard 48 - 5*	Total 188 - 142*
R	2	0	3	0	5
MCTS+R	34 - 26*	0	41 - 22*	0	75 - 48*
GR	14	0	15	0	29
MCTS+GR	35 - 26*	22	57 - 23*	0	114 - 49*
G	14	3	10	0	27
MCTS+G	35-26*	11	46 - 22*	0	92 - 48*
DS	14	2	11	0	27
MCTS+DS	35-25*	11	47 - 22*	0	93 - 47*

484 To better compare these different algorithms, and not only relying on
 485 the number of best-known scores achieved (that can sometimes be found by
 486 "chance"), we performed pairwise comparisons between the algorithms based
 487 on the average scores obtained on each instance as displayed in Table 2.

488 In Table 2, the numbers in each row correspond to the number of instances
 489 for which the method is significantly better than another (with a maximum of
 490 188 instances). A method is said significantly better than another on a given
 491 instance if its average score measured over 20 runs is significantly better (t-test
 492 with a p-value below 0.001). The column *Total* corresponds to the number of
 493 times a method is better than another.

494 Table 2 shows the ranking of the different methods. Unsurprisingly, the
 495 pure random heuristic is completely dominated by all methods. The variant
 496 MCTS+DS is significantly better compared to the others. In particular, it
 497 stays significantly better 48 times out of 188, versus 32 times in favor of the
 498 MCTS+Greedy-Random. Indeed, it seems that for the WVCP, the DSatur pro-
 499 cedure, which colors the heaviest vertices that have the most colored neighbors,
 500 leads to better organization of color groups.

Table 2 Comparison between all greedy and MCTS variants. As an example, the row for MCTS+Random means that the method is better for 186 instances compared to the random procedure (R), is better for 169 instances compared to the Greedy-Random procedure (GR), and is never better compared to MCTS+GR. Values in bold highlight the highest value between two methods. As an example, the variant MCTS+G is more often significantly better (44 times) than MCTS+GR (which is better than MCTS+G 23 times).

/188 instances	R	MCTS+R	GR	MCTS+GR	G	MCTS+G	DS	MCTS+DS
R	-	0	0	0	0	0	0	0
MCTS+R	186	-	169	0	92	1	90	1
GR	186	10	-	0	4	0	7	0
MCTS+GR	186	122	178	-	151	23	152	32
G	186	36	133	10	-	0	37	1
MCTS+G	186	121	178	44	158	-	160	30
DS	186	40	139	9	47	0	-	0
MCTS+DS	186	121	178	48	157	35	158	-

Table 3 Results of the MCTS with greedy simulation on a part of the 188 instances of the literature. At the foot of the table, for each method, we report the number of BKS achieved, the number of best scores, the number of average best scores compared to other methods, and the number of instances solved to optimality. The mean is not shown if it is equal to the best score. A star (*) indicates that the score has been proven to be optimal.

instance	BKS	MCTS+R			MCTS+GR			MCTS+G			MCTS+DSatur		
		best	mean	time	best	mean	time	best	mean	time	best	mean	time
C2000.5	2144	3178	3198.3	2912	2505	2537.7	3422	2385		1685	2397	2398.8	3390
C2000.9	5477	7022	7166.3	3554	6233	6272.9	3575	6125	6147.8	3238	6275		163
DSJC125.1g	23*	27	29.4	2332	25	25.4	22	25		0	25		0
DSJC125.5g	71	78	80.8	755	74	75.3	52	77		0	74		67
DSJC125.9g	169*	173	176.8	17	170	172.7	25	171		1630	171		21
DSJC250.1	127	159	165.2	68	134	141.4	13	141		4	139		1151
DSJC250.5	392	444	457.8	2941	422	429.4	14	427		92	421		705
DSJC250.9	934*	1002	1026.2	20	973	988.7	2793	986		16	984		559
DSJC500.1	184	237	245.9	498	203	208.3	148	203		638	203		2943
DSJC500.5	685	809	824.2	899	754	765.6	136	755		635	775	780.9	3542
DSJC500.9	1662	1831	1858.5	2645	1771	1787.4	113	1794		171	1795	1797.1	3453
DSJC1000.1	300	392	440.8	3480	334	337.4	1618	333		2823	340		2074
DSJC1000.5	1185	1385	1399.2	2699	1271	1293.6	1668	1318		1437	1338		243
DSJC1000.9	2836	3164	3200.8	1750	3040	3070.4	978	3078		2863	3172		3338
DSJR500.1	169*	178	184.4	154	169		46	177		0	176		21
flat1000.50_0	924	1337	1358.5	3022	1236	1255.8	1452	1251		1251.1	3037	1303	1843
flat1000.60_0	1162	1375	1400	3400	1275	1295.8	1478	1260	1260.7	3424	1343		311
flat1000.76_0	1165	1349	1370.3	2910	1252	1269.7	1477	1244	1248.2	3557	1313		1798
GEO120a	105*	107	111.5	14	106	106.6	330	105		1129	109		5
GEO120b	35*	37	38	10	37		0	37		498	37		110
GEO120c	72*	74	76.8	5	72	73	0	72		0	72		36
inithx.i.1	569*	569	569.5	63	569		0	569		0	569		0
inithx.i.2	329*	330	334.6	69	329	329.9	1720	330		0	331		0
inithx.i.3	337*	338	340.9	116	337	337.7	1449	339		0	339		0
latin.square_10	1480	1861	1888.6	2092	1721	1757.1	966	1726		1726.8	3418	1805	2360
le450.25a	306	315	320.2	879	307	310	3364	312		5	310		1919
le450.25b	307*	309	313.4	948	309	309.1	993	309		0	309		224
le450.25c	342	390	397.6	1501	365	372.7	2844	369		950	376		1572
le450.25d	330	384	391.8	1830	354	358.8	2802	364		1657	369		1411
myciel7gb	109*	120	133.8	18	117	118.3	476	109		42	109		38
myciel7g	29*	32	36.1	670	29	29.8	578	29		0	29		0
queen_10.10	162	174	178.2	692	165	169.1	986	171		0	169		8
queen_11.11	172	185	189.8	384	178	180.4	941	180		80	179		50
queen_12.12	185	198	208.2	547	189	196.8	2	193		1102	197		1823
queen_13.13	194	209	219.8	43	204	207.6	1125	205		5	199		71
queen_14.14	215	237	241.4	247	224	226.1	38	224		13	225		2351
queen_15.15	223	252	261.6	415	237	241.3	216	239		49	238		315
queen_16.16	234	268	274	1799	250	252.7	127	247		86	248		420
R75_1gb	70*	73	77.5	1171	70*	73.6	2679	76		184	75		9
R75_1g	18*	18*	19.8	2579	18*	18.8	1164	18*		956	18*		1348
R75_5gb	186*	195	200.7	91	186	191.7	237	192		0	192		2234
R75_5g	51*	54	55.2	608	51	52.7	725	52		203	51		1956
R75_9gb	396*	398	400.9	1281	396	397.9	1902	396		208	396		191
R75_9g	110*	110	111.5	850	110		39	110		7	110		202
R100_1gb	81*	89	97.1	1326	84	87	0	84		0	84		1
R100_1g	21*	24	26.2	1	22	22.6	728	23		0	23		0
R100_5gb	220	232	238.9	28	224	230.7	4	230		0	225		20
R100_5g	59	63	64.7	1389	62	62.4	604	63		0	61		7
R100_9gb	518*	530	538.5	43	518	521	273	526		2696	528		46
R100_9g	141*	144	146.7	0	143	143.9	29	143		5	143		301
wap01a	545	1039	1053.3	3450	657	659.8	2965	599		3272	595		2904
wap02a	538	1021	1038	3533	645	648.7	1301	588		3265	590		262
wap03a	562	1416	1445.4	2767	746	749	803	653		1096	647		1805
wap04a	563	1480	1493.9	1750	755	757.8	3085	649		603	643		828
wap05a	541	713	724.9	3536	583	591.1	1867	566		172	565		2882
wap06a	516	747	753.8	3359	591	595.4	2280	566		9	561		664
wap07a	555	1018	1030.2	2630	693	697.2	3186	635		165	639		986
wap08a	529	984	998.2	1245	664	672	3539	612		2401	604		2583
p40	4984*	4984	4992.8	45	4984		0	4984		0	4984		0
p41	2688*	2688	2700.7	14	2688		0	2688		0	2688		0
p42	2466*	2480	2515.1	16	2466	2466.8	33	2466		0	2466		4
r28	9407*	9435	9532	3599	9407	9409.9	1530	9407		0	9407		0
r29	8693*	8750	8999.5	155	8693	8695.5	1316	8694		0	8694		3
r30	9816*	9825	9876.5	384	9816	9819.6	11	9818		0	9818		0
#BKS			75			114			92			93	
#Best			79			166			111			117	
#Best Avg			56			104			131			138	
#Optimal			48			49			48			47	

501 This is particularly true for the largest instances, where choosing random
 502 moves in the set of all legal moves is not very efficient as the number of color
 503 groups grows rapidly. It explains also why the variant MCTS+Random per-
 504 forms badly on larger or denser instances such as the rxx instances or some
 505 difficult DIMACS instances.

506 However, with the deterministic simulation of the MCTS+Greedy or
 507 MCTS+DSatur variants, there is no sampling of the legal moves like in the
 508 MCTS+Greedy-Random variant allowing greater exploration of the search
 509 space and a better estimation of the most promising branches of the search
 510 tree. This particularity of the MCTS+Greedy-Random variant allows us to
 511 find the BKS for more instances. Moreover, when exploring Table 3, which
 512 presents the results for a part of the 188 instances of the literature, one can
 513 see that the R75_1gb instance from the DIMACS_easy set is proven optimal
 514 by MCTS+Greedy-Random but not by MCTS+Greedy or MCTS+DSatur.
 515 With stochastic help, the MCTS+Greedy-Random version can reach the best
 516 known score of 70, which leads to an early pruning of the tree and allows to
 517 prove the optimality earlier.

518 5.3 Exploitation vs exploration coefficient analysis

519 One key element of the MCTS algorithm is the coefficient c balancing the
 520 compromise between exploration and exploitation in Equation (3). In this sub-
 521 section, we investigate the importance of this coefficient by varying it and
 522 presenting the evolution of the score over time. For this experimentation, the
 523 coefficient c varied from 0 (no exploration) to 5 (encourage exploration). For
 524 each coefficient value, 20 runs of the MCTS+GR variant per instance are
 525 performed during 5h per run (15h for the very large C2000.x instances)⁷.

526 Figure 4 displays 6 plots showing the evolution of the mean of the best
 527 scores over time for the instances DSJC500.5, latin_square_10, le450_25a,
 528 wap01a, C2000.5 and C2000.9 for the different values of the coefficient c . These
 529 6 instances come from the set of DIMACS_hard instances and can be consid-
 530 ered as very difficult. Four typical patterns also seen for other instances are
 531 observed:

- 532 • P1: instances requiring a lot of exploration,
- 533 • P2: instances requiring more exploration than exploitation,
- 534 • P3: instances requiring more exploitation than exploration,
- 535 • P4: instances requiring a lot of exploitation.

536 The first pattern P1 is observed for the instance DSJC500.5 and also queen
 537 instances. For these instances, the lack of exploration leads to poor results,
 538 and better results are reached as the coefficient c increases. The pattern P2
 539 is observed for the instance latin_square_10, but also for other instances such
 540 as flat1000 where the best score obtained in function of the coefficient c has a
 541 U-shape, with an optimal value of c between 1 and 2. This phenomenon can

⁷This longer execution time explains some differences with the sensitivity analysis of this parameter made in [1] with only 1h of computation time.

542 also be observed on instances such as C2000.5 and C2000.9 where it becomes
 543 quickly more interesting to explore up to a certain point.

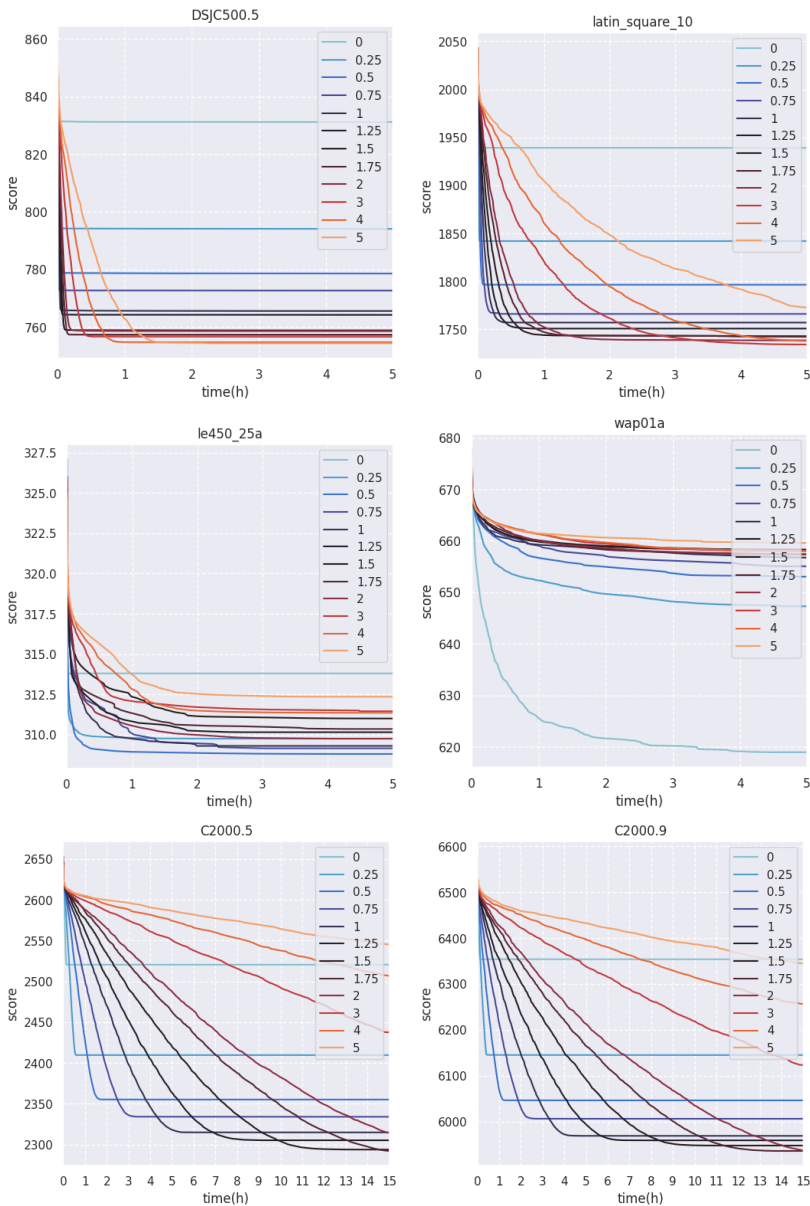


Fig. 4 Plots of the evolution of the means of the scores over time for different values of the coefficient c between 0 and 5, for the instances DSJC500.5, latin_square_10, le450_25a, wap01a, C2000.5, and C2000.9. For each configuration, 20 runs are launched with the MCTS+Greedy-Random variant for 5h and 15h for the C2000.x instances.

The pattern P3 found for the le450 instance shows the best results when there is only weak exploration, but the results are worse when c is set to zero.

In general, for the patterns P1, P2 and P3, having no exploration at all rapidly leads to a local minimum trap and it seems better to secure a minimum of diversity to reach a better score, while, for the pattern P4, found for the wap instances (very large instances), giving a chance to the exploration leads the algorithm to be lost in the search space. For very large instances, as the search tree is huge and cannot be sufficiently explored due to the time limit, it seems more beneficial for the algorithm to favor more intensification to better search for a good solution in a small part of the tree. To sum, the most suitable exploration vs exploitation coefficient thus depends on the instance considered. Finding the right general coefficient is a challenging task. In this work, We adopted the coefficient $c = 1$ for all other experiments.

5.4 Monte Carlo Tree Search with local search

This section studies the effects of the combination of MCTS with a local search heuristic. Table 4 summarizes the results of the local search procedures presented in Section 4 followed by the combination of MCTS with each of these local search procedures during the simulation phase. Each line presents the number of times a BKS is reached with the method. Note that the objective of these MCTS variants coupled with a local search procedure is not to prove optimality. For instances where the optimal score is known, MCTS and local search stop when this score is reached.

We observe from this table that the local search procedure allowing to reach the highest number of BKS is ILS-TS, followed by RedLS. When coupled with the MCTS framework, TW, AFISA, and RedLS improve their results. In particular, the variant MCTS+RedLS finds the BKS for 17 difficult instances of the DIMACS_hard set and 152 instances over 188 in total. It highlights that the MCTS framework proposed in this paper can help a local search algorithm such as RedLS to continuously find new promising starting points in the search space.

Table 4 Summary of the number of times the BKS is reached for the local search procedures and MCTS combined with the different local search procedures. Values in bold highlight the best results for each column.

	pxx 35*	rxr 30*	DIMACS easy 75(72*)	DIMACS hard 48 (5*)	Total 188 (142*)
AFISA	35	5	71	3	114
MCTS+AFISA	34	3	73	5	115
TW	29	2	61	7	99
MCTS+TW	35	16	72	9	132
RedLS	31	3	70	8	112
MCTS+RedLS	35	25	75	17	152
ILS-TS	35	30	75	19	159
MCTS+ILS-TS	35	30	75	15	155

574 However, when coupled with the ILS-TS algorithm (variant MCTS+ILS-
 575 TS), it does not seem to improve the results. It may be explained by the fact
 576 that ILS-TS is an iterated local search algorithm already integrating various
 577 perturbation mechanisms allowing to escape local optima. Therefore, finding
 578 new good starting points in the search space guided by MCTS seems less
 579 interesting for ILS-TS than for other local search procedures such as TW,
 580 AFISA, and RedLS.

581 Table 5 shows pairwise comparisons between all the MCTS variants and
 582 all the local search procedures.

583 When comparing the performances of the local search algorithms between
 584 each other by looking at the number of times they are significantly better than
 585 the others RedLS is at the end of the ranking after TabuWeight and AFISA
 586 while ILS-TS is far more often better.

587 The MCTS+local_search variants always improve the number of BKS found
 588 and are more often significantly better than the corresponding local search
 589 only, except for ILS-TS, which is better on its own.

590 The variant MCTS+RedLS is more often significantly better than the other
 591 methods, even compared to ILS-TS. These results indicate that combining
 592 MCTS with a local search is of interest to improve the underlying local search
 593 procedures such as AFISA, TW, and RedLS.

594 Table 6 shows the results for each local search and MCTS+LS on a portion
 595 of the 188 instances in the literature, with the most difficult instances more
 596 represented than the easiest. The values in bold show the best results among
 597 the different methods for the best score or the mean score.

Table 5 Comparison between all the MCTS variants and all the local search procedures. Each row corresponds to the number of times the method of the row is significantly better than the method on the column on the 188 instances. For example, out of 188 instances, TW is better than MCTS+TW on 25 instances, 48 times against RedLS and 3 times against MCTS+RedLS. Values in bold means that method *a* is more often better than method *b* when we look at how many times *b* is significantly better than *a*. The Total column corresponds to the number of times a method is more often significantly better than the others.

/188 instances	MCTS+GR	MCTS+DSatur	AFISA	MCTS+AFISA	TabuWeight	MCTS+TW	RedLS	MCTS+RedLS	ILSTS	MCTS+ILSTS	Total
MCTS+GR	-	32	77	63	77	33	85	9	0	1	4/9
MCTS+DSatur	48	-	93	72	89	42	107	22	0	2	5/9
AFISA	33	50	-	35	49	17	53	2	0	0	2/9
MCTS+AFISA	40	39	40	-	73	10	86	2	1	0	3/9
TabuWeight	56	54	40	40	-	25	48	3	0	3	1/9
MCTS+TW	52	64	74	73	78	-	101	8	1	0	6/9
RedLS	51	57	41	35	47	29	-	5	14	15	0/9
MCTS+RedLS	88	87	99	75	105	65	103	-	23	27	9/9
ILSTS	101	95	104	82	106	75	96	22	-	18	8/9
MCTS+ILSTS	101	93	102	82	103	73	92	18	2	-	7/9

Table 6 Results of the LS alone and MCTS + LS variants on a collection of the 188 instances of the table, for each method, we report the number of average best scores compared to other methods, and the number of instances solved to optimality. The mean is not shown if it is equal to the best score.

Instance	BKS	AFISA		MCTS+AFISA		TabuWaltz		MCTS+TW		RedLS		ILSTS		MCTS+ILSTS												
		best	mean	time	best	mean	time	best	mean	time	best	mean	time	best	mean	time										
C2000.5	2144	2384	2403.4	3601	2035	2048.9	2665	2118	2322.2	3477	2110	2123.6	1578	2167	2193.8	2403	2199	2213.2	2867	2237	2266.4	3408	2188	2381.2	697	
C2000.9	5477	6282	6650.1	0	6547	6073.2	1071	6049	6073.4	2819	6242	6293.6	1845	5502	5521.8	3303	5734	5742.2	1604	5910	5969.9	3587	6103	6119.9	430	5
DS1C125.5g	21	23	24.4	13	23	23.5	1203	24	24.6	68	23	23.1	1684	23	23.4	0	23	23	0	23	23	0	2	23	5	2
DS1C125.5g	71	72.2	1360	72	71	71.5	1603	71	71.8	1686	72	71.8	1686	72	71.4	207	71	71.2	1462	71	71	128	71	0	60	0
DS1C125.5g	169	170	174.1	2	169	169.1	1047	169	169.7	747	169	169.8	959	169	0	0	169	0	0	169	182	169	169	171	0	0
DS1C20.1	127	129	133.6	3294	133	133.8	1764	134	136.8	29	132	133.5	780	130	131.6	0	127	128.4	867	127	127.8	1608	127	128	612	0
DS1C20.5	392	411	424.2	541	421	428.8	2586	397	406	2717	406	405.6	2349	398	401.2	103	394	398.4	1980	395	397.6	2567	397	399.8	160	0
DS1C20.9	934*	949	976.1	232	962	978.8	1380	959	963.2	1766	950	956	2292	934	935.6	718	935	935.5	1279	938	943.1	3003	948	943.1	2003	948
DS1C50.1	184	188	201.3	1837	189	192	228.6	185	191	1526	201	202.2	1312	187	188.1	2087	188	188.8	1744	187	188.6	2168	187	188.6	2168	0
DS1C50.5	685	702	778.1	1307	748	863.5	1474	733	743.1	3071	751	756	2233	706	716.1	2840	710	713.6	2626	724	733.5	1744	729	734.7	3410	0
DS1C50.9	1662	1744	1775.5	652	1909	1919.3	2002	1733	1755.1	228	1764	1775	1765	1706	1676.1	945	1674	1683.1	404	1720	1742	3009	1762	1773.7	77	0
DS1C100.1	300	319	325	2182	396	401.9	357	313	316.2	2943	365	368.8	2004	305	307.2	1235	304	306.5	2814	305	307.4	1574	303	304.9	1176	0
DS1C100.5	1185	1308	1310	3531	1463	1475	2226	1271	1282.3	2260	1304	1311.1	2562	1198	1213.7	2381	1214	1218.7	2278	1245	1269.2	231	1257	1274.2	1890	0
DS1C100.9	2856	3066	3107	3901	3303	3321.1	2793	3034	3061.1	1384	3135	3132.9	156	2840	2853.5	2953	2884	2892.8	211	3026	3068.8	3580	2998	3123.7	3903	0
DS1R60.1	169	169	169	54	169	169	157	171	174.2	21	169	169.2	1273	111	114.5	0	169	169.2	1270	169	0	169	5	0	0	0
Bat1000_50.0	924	1267	1293.3	2796	1427	1431.3	1512	1238	1247.5	3290	1271	1274.9	2194	1155	1173.8	3238	1171	1179.1	1005	1222	1235	1271	1221	1231.3	42	0
Bat1000_40.0	1162	1309	1323.2	3584	1465	1476.8	1375	1275	1288	2275	1305	1313.5	21	1191	1205.7	1890	1208	1216	201	1250	1270	125	1259	1266.2	3360	0
Bat1000_76.0	1165	1288	1304.4	3278	1442	1449.8	2341	1237	1263	2975	1285	1292.8	693	1176	1194	1107	1189	1196.6	201	1232	1247.5	1198	1237	1246.5	1008	0
GEDM120b	105*	105	106.6	136	105	105.1	515	105	106.1	437	105	105	0	105	109.2	9	105	0	0	0	0	0	105	0	0	0
GEDM120b_35	35	35	36.2	42	35	35	36.2	42	35	35.3	2051	35	35.3	0	0	0	35	35	0	0	0	35	0	0	0	0
GEDM120b_72*	72	72	73	0	72	72	73	0	72	73	2097	72	75.2	0	0	0	72	72	0	0	0	72	0	72	0	0
infix_b.1	569*	569	573.1	41	569	572.4	0	569	572.4	0	569	572.4	0	569	572.4	0	569	572.4	0	0	0	569	16	569	0	0
infix_b.2	329*	334	337.6	448	331	333.9	150	338	342.0	0	332	334.1	2510	329	329.1	311	329	329	0	0	0	329	127	329	8	0
infix_b.3	337*	340	342.1	1384	339	341	98	342	347.1	0	338	340.4	3380	337	337.6	73	337	337	0	0	0	337	29	337	0	0
latin_square.10	1480	1607	1652.4	2257	2037	2055.9	2384	1816	1855.2	45	1774	1789.2	247	1505	1523.2	360	1518	1527.1	543	1559	1581.2	1398	1572	1585.8	3458	0
le49_25a	306	311	316.3	2699	316	318.9	2891	321	325.4	83	313	315.9	1456	306	307.4	503	306	307	0	0	0	306	174	306	736	0
le49_25b	307*	308	312.6	1891	308	310.1	2114	306	310.9	219	307	308.5	1491	307	311.4	56	307	307	0	0	0	307	10	307	0	0
le49_25c	342	355	364.4	3436	368	398.2	3375	361	366	677	371	373.6	1908	351	354.8	43	348	349.6	2165	348	351.8	3207	351	353.1	1689	0
le49_25d	320	351	357.8	1630	383	391.8	1494	351	356.4	1955	363	363.8	3415	332	338.9	154	334	342.5	2477	339	343.5	1999	342	348.8	2330	0
mysq17g	109*	109	111.8	1129	109	109.1	1343	114	115.8	738	112	113.5	2646	109	116.4	0	109	109.1	955	109	0	109	5	109	0	0
mysq187g	29*	29	30.7	2542	29	30.7	497	29	29.2	324	29	29	115	29	29.8	244	29	29	0	0	0	29	4	29	0	0
open10.1d	162	166.7	169.4	524	162	164.1	2964	163	163.9	1629	162	163.3	1249	162	163.6	504	162	163.6	162	163.6	162	163.6	162	163.6	162	0
open11.1	172	179	181.1	172	176	178.8	2859	172	179.2	1731	172	173.2	1071	174	171.7	161	172	173.8	2575	172	172.6	1820	172	172.8	1703	0
open11.1d	185	193	196.7	2062	191	193.9	3218	187	188.1	1883	187	188.4	2635	188	189.2	7	186	189.9	2484	185	186.1	1281	186	186.1	1572	0
open11.1d	194	203	206.7	2062	191	205.2	1684	198	201.1	2285	199	200.5	2765	195	198.7	1528	194	194.7	2243	194	195.7	1875	195	196.6	1600	0
open14.1d	215	224	230.5	3113	226	229.2	1813	217	218.4	1878	219	220.1	2014	217	222.4	18	216	217.3	3040	216	217.3	2018	217	218.4	1514	0
open15.1d	223	236	240.1	1876	241	242.9	1745	233	236.8	2417	233	235.5	1435	227	229.7	1160	225	226.6	1784	227	229.4	2167	227	229.2	1925	0
open16.1d	234	248	255.2	1055	255	258.2	3012	239	242.3	2745	243	244.8	1062	237	239.9	91	235	237.4	1300	238	240.2	2508	240	241.2	1204	0
R75_1g	70*	70	73.7	0	70	70	0	0	0	0	70	70.5	1432	70	78.3	0	70	70	0	0	0	70	0	70	0	0
R75_1g	18*	18	18	94	18	18	15	19	19.6	0	18	18	164	18	19.4	0	18	18	0	0	0	18	0	18	0	0
R75_2g	186*	186	188.6	28	186	186	45	186	186.3	135	186	186	0	186	192.1	0	186	186	0	0	0	186	72	186	1	186
R75_2g	51*	51	51.9	17	51	51	0	0	0	0	51	51.3	118	51	51.3	168	51	51	0	0	0	51	0	51	0	0
R75_2g	396*	396	396.4	8	396	396	12	396	396.6	1911	396	396.1	1380	396	0	0	396	0	0	0	0	396	10	396	16	0
R75_2g	110*	110	110.2	3	110	110	0	110	0	0																

613 When the instance is large and when a time limit is imposed, MCTS does
614 not have the time to learn promising areas in the search space and it seems
615 more beneficial to favor more intensification, which can be done in three dif-
616 ferent ways: (i) by lowering the coefficient which balances the compromise
617 between exploitation and exploration during the selection phase, (ii) by using
618 a dedicated heuristic exploiting the specificity of the problem (grouping in pri-
619 ority the heaviest vertices in the first groups of colors), (iii) and by using a
620 local search procedure to improve the complete solution.

621 Conversely, for small instances, it seems more beneficial to encourage more
622 exploration, to avoid getting stuck in local optima. It can be done, by increas-
623 ing the coefficient, which balances the compromise between exploitation and
624 exploration, and by using a simulation strategy with more randomness, which
625 favors more exploration of the search tree and also allows a better evaluation
626 of the most promising branches of the MCTS tree. For these small instances,
627 the MCTS algorithm can provide some optimality proofs.

628 For medium instances, it seems important to find a good compromise
629 between exploration and exploitation. For such instances, coupling the MCTS
630 algorithm with a local search procedure allows finding better solutions, which
631 cannot be reached by the MCTS algorithm or the local search alone.

632 Other future works could be envisaged. For example, an interesting study
633 would be to automatically choose the balance coefficient between exploitation
634 and exploration on the fly when solving each specific instance. It could also be
635 interesting to use a more adaptive approach to trigger the local search, or to
636 use a machine-learning algorithm to guide the search toward more promising
637 branches of the search tree.

638 Acknowledgements

639 We would like to thank Dr. Wen Sun [5], Dr. Yiyuan Wang, [6] and Pr. Bruno
640 Nogueira [4] for sharing their codes. We also thank the organizers and editors
641 of Evostar 2022 to give us the opportunity to present this extended version
642 of our conference paper [1] for this Topical Issue of the SN Computer Science
643 journal. We thank the reviewer for the useful comments, which helped us to
644 improve the work. This work was granted access to the HPC resources of
645 IDRIS (Grant No. 2020-A0090611887) from GENCI and the Centre Régional
646 de Calcul Intensif des Pays de la Loire (CC IPL).

647 Authors' contributions

648 C. Grelier developed the code, prepared the reduced instances, and performed
649 the tests. O. Goudet and J.K. Hao planned and supervised the work. All the
650 authors contributed to the analysis and the writing of the manuscript.

651 Conflict of Interest

652 The authors declare that they have no conflicts of interest.

Availability of data set, code, and results

The instances used in this work, in reduced version, are available at https://github.com/Cyril-Grelier/gc_instances. The source code of the tested methods and their detailed results are available at https://github.com/Cyril-Grelier/gc_wvcp_adaptive_mcts.

References

- [1] Grelier, C., Goudet, O., Hao, J.-K.: On monte carlo tree search for weighted vertex coloring. In: Pérez Cáceres, L., Verel, S. (eds.) *Evolutionary Computation in Combinatorial Optimization. Lecture Notes in Computer Science*, vol. 13222, pp. 1–16 (2022)
- [2] Lewis, R.: *Guide to Graph Colouring: Algorithms and Applications*. Springer, (2021)
- [3] Goudet, O., Grelier, C., Hao, J.-K.: A deep learning guided memetic framework for graph coloring problems. arXiv:2109.05948 [cs] (2021)
- [4] Nogueira, B., Tavares, E., Maciel, P.: Iterated local search with tabu search for the weighted vertex coloring problem. *Computers & Operations Research* **125**, 105087 (2021)
- [5] Sun, W., Hao, J.-K., Lai, X., Wu, Q.: Adaptive feasible and infeasible tabu search for weighted vertex coloring. *Information Sciences* **466**, 203–219 (2018)
- [6] Wang, S., Yiyuanand Cai, Pan, S., Li, X., Yin, M.: Reduction and local search for weighted graph coloring problem. In: *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, New York, NY, USA, February 7-12, 2020*, pp. 2433–2441 (2020)
- [7] Prais, M., Ribeiro, C.C.: Reactive grasp: An application to a matrix decomposition problem in tdma traffic assignment. *INFORMS Journal on Computing* **12**(3), 164–176 (2000)
- [8] Pemmaraju, S.V., Raman, R.: Approximation algorithms for the max-coloring problem. In: Caires, L., Italiano, G.F., Monteiro, L., Palamidessi, C., Yung, M. (eds.) *Automata, Languages and Programming. Lecture Notes in Computer Science*, pp. 1064–1075 (2005)
- [9] Furini, F., Malaguti, E.: Exact weighted vertex coloring via branch-and-price. *Discrete Optimization* **9**(2), 130–136 (2012)
- [10] Malaguti, E., Monaci, M., Toth, P.: Models and heuristic algorithms for a weighted vertex coloring problem. *Journal of Heuristics* **15**(5), 503–526

- 688 (2009)
- 689 [11] Cornaz, D., Furini, F., Malaguti, E.: Solving vertex coloring problems as
690 maximum weight stable set problems. *Discrete Applied Mathematics* **217**,
691 151–162 (2017)
- 692 [12] Goudet, O., Grelier, C., Lesaint, D.: New bounds and constraint program-
693 ming models for the weighted vertex coloring problem. Proceedings of the
694 Thirty-Second International Joint Conference on Artificial Intelligence,
695 IJCAI 2023, 19th-25th August 2023, Macao, SAR, China, 1927–1934
696 (2023)
- 697 [13] Grelier, C., Goudet, O., Hao, J.-K.: A memetic algorithm with adaptive
698 operator selection for graph coloring. In: Stützle, T., Wagner, M. (eds.)
699 *Evolutionary Computation in Combinatorial Optimization*, pp. 65–80.
700 Springer, Cham (2024)
- 701 [14] Zhou, Y., Hao, J.-K., Duval, B.: Reinforcement learning based local search
702 for grouping problems: A case study on graph coloring. *Expert Systems*
703 *with Applications* **64**, 412–422 (2016)
- 704 [15] Zhou, Y., Duval, B., Hao, J.-K.: Improving probability learning based
705 local search for graph coloring. *Applied Soft Computing* **65**, 542–553
706 (2018)
- 707 [16] Gelly, S., Wang, Y., Munos, R., Teytaud, O.: Modification of uct with
708 patterns in monte-carlo go. Technical report, RR-6062, INRIA (2006)
- 709 [17] Browne, C.B., Powley, E., Whitehouse, D., Lucas, S.M., Cowling, P.I.,
710 Rohlfshagen, P., Tavener, S., Perez, D., Samothrakis, S., Colton, S.:
711 A survey of monte carlo tree search methods. *IEEE Transactions on*
712 *Computational Intelligence and AI in Games* **4**(1), 1–43 (2012)
- 713 [18] Silver, D., Huang, A., Maddison, C.J., Guez, A., Sifre, L., van den Driess-
714 che, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot,
715 M., et al.: Mastering the game of go with deep neural networks and tree
716 search. *Nature* **529**(7587), 484–489 (2016)
- 717 [19] Edelkamp, S., Greulich, C.: Solving physical traveling salesman prob-
718 lems with policy adaptation. In: 2014 IEEE Conference on Computational
719 Intelligence and Games, pp. 1–8 (2014). IEEE
- 720 [20] Jookan, J., Leyman, P., Wauters, T., De Causmaecker, P.: Exploring
721 search space trees using an adapted version of Monte Carlo tree search for
722 combinatorial optimization problems. *Computers & Operations Research*
723 **150**, 106070 (2023)

- 724 [21] Cazenave, T., Negrevergne, B., Sikora, F.: Monte carlo graph coloring. In:
725 Monte Carlo Search 2020, IJCAI Workshop (2020)
- 726 [22] Brélaz, D.: New methods to color the vertices of a graph. *Communications*
727 *of the ACM* **22**(4), 251–256 (1979)
- 728 [23] Kubale, M., Jackowski, B.: A generalized implicit enumeration algorithm
729 for graph coloring. *Communications of the ACM* **28**(4), 412–418 (1985)
- 730 [24] Demange, M., Werra, D.d., Monnot, J., Paschos, V.T.: Time slot schedul-
731 ing of compatible jobs. *Journal of Scheduling* **10**(2), 111–127 (2007)
- 732 [25] Lai, T.L., Robbins, H.: Asymptotically efficient adaptive allocation rules.
733 *Advances in Applied Mathematics* **6**(1), 4–22 (1985)
- 734 [26] Hertz, A., de Werra, D.: Using tabu search techniques for graph coloring.
735 *Computing* **39**, 345–351 (1987)
- 736 [27] Cai, S., Su, K., Sattar, A.: Local search with edge weighting and configu-
737 ration checking heuristics for minimum vertex cover. *Artificial Intelligence*
738 **175**(9), 1672–1696 (2011)