

Single-Player Monte-Carlo Tree Search for SameGame

Maarten P.D. Schadd^{a,*}, Mark H.M. Winands^a, Mandy J.W. Tak^a,
Jos W.H.M. Uiterwijk^a

^a*Games and AI Group, Department of Knowledge Engineering, Maastricht University,
Maastricht, The Netherlands*

Abstract

Classic methods such as A* and IDA* are a popular and successful choice for one-player games. However, without an accurate admissible evaluation function, they fail. In this article we investigate whether Monte-Carlo Tree Search (MCTS) is an interesting alternative for one-player games where A* and IDA* methods do not perform well. Therefore, we propose a new MCTS variant, called Single-Player Monte-Carlo Tree Search (SP-MCTS). The selection and back-propagation strategy in SP-MCTS are different from standard MCTS. Moreover, SP-MCTS makes use of randomized restarts. We tested IDA* and SP-MCTS on the puzzle SameGame and used the Cross-Entropy Method to tune the SP-MCTS parameters. It turned out that our SP-MCTS program is able to score a substantial number of points on the standardized test set.

Keywords: Monte-Carlo Tree Search, One-Player Game, Puzzle, SameGame, Cross-Entropy Method

1. Introduction

The traditional approaches to deterministic one-player games with perfect information (Kendall et al., 2008) are applying A* (Hart et al., 1968), IDA* (Korf, 1985) or one of their variants (e.g., Multiobjective A* (Stewart and White, 1991), Limited-Damage A* (Bayili and Polat, 2011)). These methods have been quite successful for solving this type of games. The disadvantage of the methods is that they require an admissible heuristic evaluation function. The construction of such a function can be difficult. Since Monte-Carlo Tree Search (MCTS) (Chaslot et al., 2006b; Coulom, 2007; Kocsis and Szepesvári, 2006) does not require an admissible heuristic and because of its success in two-player games (Lee et al., 2010) and multi-player games (Sturtevant, 2008),

*Corresponding author, Tel: (+31) (0)43-3882005, Fax: (+31) (0)43-3884910

Email addresses: maarten.schadd@maastrichtuniversity.nl (Maarten P.D. Schadd),
m.winands@maastrichtuniversity.nl (Mark H.M. Winands),
m.tak@student.maastrichtuniversity.nl (Mandy J.W. Tak),
uiterwijk@maastrichtuniversity.nl (Jos W.H.M. Uiterwijk)

MCTS seems a promising approach for deterministic one-player games with perfect information as well (Schadd et al., 2008b; Tak, 2010).

So far, MCTS has not been widely applied in one-player games. One example is the Sailing Domain (Kocsis and Szepesvári, 2006), which is a non-deterministic game with perfect information. MCTS has also been used for optimization and planning problems which can be represented as deterministic one-player games. Chaslot et al. (2006a) applied MCTS in production management problems. Mesmay et al. (2009) proposed the MCTS variant TAG for optimizing libraries for different platforms.

This article proposes an MCTS method, called Single-Player Monte-Carlo Tree Search (SP-MCTS), for a one-player game.¹ MCTS for two-player games forms the starting point for this search method. We adapted MCTS by two modifications resulting in SP-MCTS. The modifications are (1) in the selection strategy and (2) in the backpropagation strategy. SP-MCTS is tested in the puzzle² SameGame, because there exists no reliable admissible heuristic evaluation function for this game. This article also investigates whether the Cross-Entropy Method (Rubinstein, 2003) is able to fine tune the parameters of SP-MCTS.

The article is organized as follows. In Section 2 we present the rules, complexity and related work of SameGame. In Section 3 we discuss why classic approaches such as A* and IDA* are not suitable for SameGame. Then, we introduce the SP-MCTS approach in Section 4. Section 5 describes the Cross-Entropy Method which is used for tuning the SP-MCTS parameters. Experiments and results are given in Section 6. Section 7 gives the conclusions and indicates future research.

2. SameGame

SameGame is a puzzle invented by Kuniaki Moribe under the name *Chain Shot!* in 1985. It was distributed for Fujitsu FM-8/7 series in a monthly personal computer magazine called *Gekkan ASCII* (Moribe, 1985). The puzzle was afterwards re-created by Eiji Fukumoto under the name of *SameGame* in 1992.

In this section, we first explain the rules in Subsection 2.1. Subsequently, we analyze the complexity of SameGame in Subsection 2.2. Finally, we present related work in Subsection 2.3.

2.1. Rules

SameGame is played on a vertically oriented 15×15 board initially filled with blocks of 5 colors at random. A move consists of removing a group of (at least two) orthogonally adjacent blocks of the same color. The blocks on top of the removed group fall down. As soon as an empty column occurs, the columns to the right of the empty column are shifted to the left. Therefore, it is impossible

¹Preliminary work has been performed by Schadd et al. (2008b).

²From now on, we call one-player deterministic games with perfect information for the sake of brevity *puzzles* (Kendall et al., 2008).

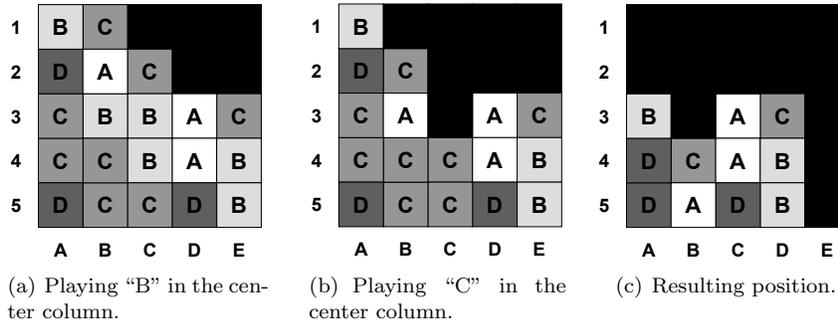


Figure 1: Example SameGame moves.

to create separate subgames. For each removed group points are rewarded. The number of points is dependent on the number of blocks removed and can be computed by the formula $(n - 2)^2$, where n is the size of the removed group.

We show two example moves in Figure 1. When the "B" group in the third column with a connection to the second column of position 1(a) is played, the "B" group is removed from the game. In the second column the "CA" blocks fall down and in the third column the "C" block falls down, resulting in position 1(b). Due to this move, it is now possible to remove a large group of "C" blocks ($n = 6$). Owing to an empty column the two columns at the right side of the board are shifted to the left, resulting in position 1(c).³ The first move is worth 1 point; the second move is worth 16 points.

The game is over if no more blocks can be removed. This happens when either the player (1) has removed all blocks or (2) is left with a position where no adjacent blocks have the same color. In the first case, 1,000 bonus points are rewarded. In the second case, points are deducted. The formula for deducting is similar to the formula for awarding points but now iteratively applied for each color left on the board. Here it is assumed that all blocks of the same color are connected.

There are variations that differ in board size and the number of colors, but the 15×15 variant with 5 colors is the accepted standard. If a variant differs in the scoring function, it is named differently (e.g., Clickomania (Biedl et al., 2002) or Jawbreaker (Julien, 2008)).

2.2. Complexity of SameGame

The complexity of a game indicates a measure of difficulty for solving the game. Two important measures for the complexity of a game are the game-tree complexity and the state-space complexity (Allis, 1994). The game-tree

³Shifting the columns at the left side to the right would not have made a difference in number of points. For consistency, we always shift columns to the left.

complexity is the number of leaf nodes in the solution search tree of the initial position(s) of the game. The state-space complexity is the number of legal game positions reachable from the initial position(s) of the game.

For SameGame these complexities are as follows. The game-tree complexity is approximated by b^d , where b is the average branching factor and d is the average game length. b and d may be obtained by self-play experiments. We used our engine to play 250 puzzles. The average length of the game was estimated to be 62.2 moves and the average branching factor to be 21.1, resulting in a game-tree complexity of 10^{82} . The first step for computing the state-space complexity is to compute the number of possible states for one column. It is possible to calculate the number of states for one column by $C = \sum_{n=0}^r c^n$ where r is the height of the column and c is the number of colors. To compute the state-space complexity we take C^k where k is the number of columns. For SameGame, using $c = 5$, $r = 15$ and $k = 15$, there exist 10^{159} states. This is an over-estimation because a small percentage of the positions are symmetrical.

Furthermore, the difficulty of a game can be described by deciding to which complexity class it belongs (Johnson, 1990). The similar game Clickomania was proven to be NP-complete by Biedl et al. (2002). However, the complexity of SameGame could be different. The more points are rewarded for removing large groups, the more the characteristics of the game may differ from Clickomania. In Clickomania the only goal is to remove as many blocks as possible, whereas in SameGame points are rewarded for removing large groups as well. In the following, we show that SameGame independently from its scoring function is at least NP-complete (Schadd et al., 2008a), such as the 3-SAT problem (Cook, 1971).

Theorem. *SameGame is at least NP-complete.*

Proof. To prove that SameGame is at least NP-complete, it is sufficient to reduce SameGame to a simpler problem which is NP-complete. Instead of proving that finding the optimal path in SameGame is NP-complete, we prove that checking whether a given solution S is optimal is already NP-complete. A solution S is defined as a path from the initial position to a terminal position. Either S (1) has removed all blocks from the game or (2) has finished with blocks remaining on the board. In both cases a search has to be performed to investigate whether a solution exists that improves the score by clearing the board. If finding a solution that clears the board with a higher score is NP-complete, then SameGame is at least NP-complete.

Clickomania is a variant of SameGame where no points are rewarded and the only objective is to clear the board. Finding only one solution to this problem is easier than finding a higher scoring solution (as in SameGame). Therefore, SameGame is at least as difficult as Clickomania. Because Clickomania is NP-complete with 5 colors and at least 2 columns (Biedl et al., 2002), SameGame with 5 colors and 15 columns is thus at least NP-complete. \square

2.3. Related Work

For the game of SameGame some research has been performed. The contributions are benchmarked on a standardized test set of 20 positions.⁴ The first SameGame program was written by Billings (2007). This program applies a non-documented method called *Depth-Budgeted Search* (DBS). When the search reaches a depth where its budget has been spent, a greedy simulation is performed. On the test set his program achieved a total score of 72,816 points with 2 to 3 hours computing time per position. When we first introduced Single-Player Monte-Carlo Tree Search (SP-MCTS), it was able to set a new high score of 73,998 points (Schadd et al., 2008b). This article will describe SP-MCTS in detail and further improve that score. Takes and Kusters (2009) proposed *Monte Carlo with Roulette-Wheel Selection* (MC-RWS). It is a simulation strategy that tries to maximize the size of one group of a certain color and at the same time tries to create larger groups of another color. On the test set their program achieved a total score of 76,764 points with a time limit of 2 hours. In the same year Cazenave (2009) applied Nested Monte-Carlo Search which led to an even higher score of 77,934. Until the year 2010, the top score on this set was 84,414 points, held by the program SPURIOUS AI.⁵ This program applies a method called Simple Breadth Search (SBS), which uses beam search, multiple processors and a large amount of memory (cf. Takes and Kusters, 2009). Further details about this program are not known. Later in 2010 this record was claimed to be broken with 84,718 points by using a method called *Heuristically Guided Swarm Tree Search* (HGSTS) (Edelkamp et al., 2010), which is a parallelized version of MCTS. The scores of HGSTS were not independently verified to be correct.

3. Search Algorithms for Puzzles

The classic approach to puzzles involves methods such as A* (Hart et al., 1968) and IDA* (Korf, 1985). A* is a best-first search where all nodes have to be stored in a list. The list is sorted by an admissible evaluation function. At each iteration the first element is removed from the list and its children are added to the sorted list. This process is continued until the goal state arrives at the start of the list. IDA* is an iterative deepening variant of A* search. It uses a depth-first approach in such a way that there is no need to store the complete tree in memory. The search continues depth-first until the cost of arriving at a leaf node and the value of the evaluation function exceeds a certain threshold. When the search returns without a result, the threshold is increased.

Both methods are strongly dependent on the quality of the evaluation function. Even if the function is an admissible under-estimator, it still has to give an accurate estimation. Classic puzzles where this approach works well are the Eight Puzzle with its larger relatives (Korf, 1985; Sadikov and Bratko, 2007)

⁴The positions can be found at: www.js-games.de/eng/games/samegame.

⁵The exact date when the scores were uploaded to <http://www.js-games.de/> is unknown.

and Sokoban (Junghanns, 1999). Here a good under-estimator is the well-known Manhattan Distance. The main task in this field of research is to improve the evaluation function, e.g., with pattern databases (Culberson and Schaeffer, 1998; Felner et al., 2005).

These classic methods fail for SameGame because it is not straightforward to make an admissible function that still gives an accurate estimation. An attempt to make such an evaluation function is by just awarding points to the current groups on the board. This resembles the score of a game where all groups are removed in a top-down manner. However, if an optimal solution to a SameGame problem has to be found, we may argue that an “over-estimator” of the position is required, because in SameGame the score has to be maximized, whereas in common applications costs have to be minimized (e.g., shortest path to a goal). An admissible “over-estimator” can be created by assuming that all blocks of the same color are connected and would be able to be removed at once. This function can be improved by checking whether there is a color with only one block remaining on the board. If this is the case, the 1,000 bonus points for clearing the board may be deducted because the board cannot be cleared completely. However, such an evaluation function is far from the real score for a position and does not give good results with A* and IDA*. Our tests have shown that using A* and IDA* with the proposed “over-estimator” results in a kind of breadth-first search. The problem is that after expanding a node, the heuristic value of a child can be significantly lower than the value of its parent, unless a move removes all blocks with one color from the board. We expect that other Depth-First Branch-and-Bound methods (Vempaty et al., 1991) suffer from the same problem. Since no good evaluation function has been found yet, SameGame presents a new challenge for puzzle research.

Another alternative is to apply the GRAPHPLAN algorithm (Blum and Furst, 1997). This algorithm constructs a planning graph in which there exist proposition nodes and action nodes. Special edges are used to denote relations, such as preconditions. At each iteration i of the algorithm, a search is run to prove or disprove if a plan using i steps exists. This algorithm is effective if (1) such a planning graph can easily be constructed and (2) planning-graph-based heuristics can be developed. SameGame can be represented as a directed acyclic graph, although the goal criterion is not known (i.e., the maximum number of points achievable in a position is usually unknown). A move may drastically change the board, making the application of heuristics rather difficult.

4. Single-Player Monte-Carlo Tree Search

MCTS is a best-first search method, which does not require a positional evaluation function. MCTS builds a search tree employing Monte-Carlo simulations at the leaf nodes. Each node in the tree represents an actual board position and typically stores the average score found in the corresponding subtree and the number of visits. MCTS constitutes a family of tree-search algorithms applicable to the domain of board games (Chaslot et al., 2006b; Coulom, 2007; Kocsis and Szepesvári, 2006).

In general, MCTS consists of four steps, repeated until time has run out (Chaslot et al., 2008b). (1) A *selection strategy* is used for traversing the tree from the root to a leaf. (2) A *simulation strategy* is used to finish the game starting from the leaf node of the search tree. (3) The *expansion strategy* is used to determine how many and which children are stored as promising leaf nodes in the tree. (4) Finally, the result of the MC evaluation is propagated backwards to the root using a *backpropagation strategy*.

Based on MCTS, we propose an adapted version for puzzles: Single-Player Monte-Carlo Tree Search (SP-MCTS) (Schadd et al., 2008b). We discuss the four steps (selection, play-out, expansion and backpropagation) and point out differences between SP-MCTS and MCTS in Subsection 4.1. SameGame serves as example domain to explain SP-MCTS. The final move selection is described in Subsection 4.2. Subsection 4.3 describes how randomized restarts can improve the score.

4.1. The Four Steps of SP-MCTS

In this subsection we discuss the four steps of SP-MCTS that are applied iteratively until time runs out. These steps consist of (1) the *Selection Step*, (2) the *Play-out Step*, (3) the *Expansion Step*, and (4) the *Backpropagation Step*.

Selection Step

Selection is the strategic task to select one of the children of a given node. It controls the balance between *exploitation* and *exploration*. Exploitation is the task to focus on the moves that led to the best results so far. Exploration deals with the less promising moves that still may have to be explored, due to the uncertainty of their evaluation so far. In MCTS at each node starting from the root, a child has to be selected until a position is reached that is not part of the tree yet. Several strategies have been designed for this task (Chaslot et al., 2006b; Coulom, 2007; Kocsis and Szepesvári, 2006).

Kocsis and Szepesvári (2006) proposed the selection strategy UCT (Upper Confidence bounds applied to Trees). For SP-MCTS, we use a modified UCT version. At the selection of node p with children i , the strategy chooses the move, which maximizes the following formula.

$$v_i + C \times \sqrt{\frac{\ln n_p}{n_i}} + \sqrt{\frac{\sum r^2 - n_i \times v_i^2 + D}{n_i}} \quad (1)$$

The first two terms constitute the original UCT formula. It uses n_i as the number of times that node i was visited where i denotes a child and p the parent to give an upper confidence bound for the average game value v_i . For puzzles, we added a third term, which represents a possible deviation of the child node (Chaslot et al., 2006a; Coulom, 2007). It contains the sum of the squared results ($\sum r^2$) so far achieved at the child node corrected by the expected results $n_i \times v_i^2$. A high constant D is added to ensure that nodes, which have been rarely explored, are considered promising. As this deviation term is not domain specific, it may also be beneficial in other domains where the variation on the

scores is an important feature (e.g., puzzles with the aim to score points). Below we describe two differences between puzzles and two-player games, which may affect the selection strategy.

First, the essential difference between puzzles and two-player games is the *range of values*. In two-player games, the outcome of a game is usually denoted by *loss*, *draw*, or *win*, i.e., $\{-1, 0, 1\}$. The average score of a node always stays within $[-1, 1]$. In a puzzle, an arbitrary score can be achieved that is not by definition within a preset interval. For example, in SameGame there are positions, which result in a value above 5,000 points. As a first solution to this issue we may set the constants C and D in such a way that they are feasible for a certain interval (e.g., $[0, 6000]$ in SameGame). A second solution would be to scale the values back into the above mentioned interval $[-1, 1]$, given a maximum score (e.g., 6,000 for a SameGame position). When the exact maximum score is not known a theoretical upper bound can be used. For instance, in SameGame a theoretical upper bound is to assume that all blocks have the same color. A direct consequence of such an upper bound is that due to the high upper bound, the game scores are located near to zero. It means that the constants C and D have to be set with completely different values compared to two-player games. We have opted for the first solution in our program.

A second difference is that puzzles do not have any *uncertainty on the opponent's play*. It means that the line of play has to be optimized without the hindrance of an opponent (Chaslot, 2010). Due to this, not only the average score but the top score so far of a move can be used as well. Based on manual tuning, we add the top score using a weight W with a value of 0.02 to the average score.

Here we remark that we follow Couloum (2007) in choosing a move according to the selection strategy only if n_p reaches a certain threshold T . Based on preliminary experiments, we set T to 10. This parameter will be tuned in Subsection 6.6. As long as the threshold is not exceeded, the simulation strategy is used. The latter is explained below.

Play-Out Step

The play-out step begins when we enter a position that is not part of the tree yet. Moves are randomly selected until the game ends. This succeeding step is called the play-out. In order to improve the quality of the play-outs, the moves are chosen quasi-randomly based on heuristic knowledge (Bouzy, 2005; Chen and Zhang, 2008; Gelly et al., 2006). For SameGame, several simulation strategies exist.

We have proposed two simulation strategies, called TabuRandom and TabuColorRandom. Both strategies aim at creating large groups of one color. In SameGame, creating large groups of blocks is advantageous. TabuRandom chooses a random color at the start of a play-out. The idea is not to allow playing this color during the play-out unless there are no other moves possible. With this strategy large groups of the chosen color are formed automatically. The new aspect in the TabuColorRandom strategy with respect to the previous strategy is that the chosen color is the color most frequently occurring at the

start of the play-out. This may increase the probability of having large groups during the play-out. We also use the ϵ -greedy policy to occasionally deviate from this strategy (Sutton and Barto, 1998). Before the simulation strategy is applied, with probability ϵ a random move is played. Based on manual tuning, we chose $\epsilon = 0.003$.

An alternative simulation strategy for SameGame is Monte-Carlo with Roulette-Wheel Selection (MC-RWS) (Takes and Kusters, 2009). This strategy not only tries to maximize one group of a certain color, but also tries to create bigger groups of other colors. Tak (2010) showed that MC-RWS does not improve the score in SP-MCTS because it is computationally more expensive than TabuColorRandom.

Expansion Step

An expansion strategy decides on which nodes are stored in memory. Coulom (2007) proposed to expand one child per play-out. With his strategy, the expanded node corresponds to the first encountered position that was not present in the tree. This is also the strategy we used for SP-MCTS.

Backpropagation Step

During the backpropagation step, the result of the play-out at the leaf node is propagated backwards to the root. Several backpropagation strategies have been proposed in the literature (Chaslot et al., 2006b; Coulom, 2007). The best results that we have obtained for SP-MCTS was by using the plain average of the play-outs. Therefore, we update (1) the average score of a node. Additional to this, we also update (2) the sum of the squared results because of the third term in the selection strategy (see Formula 1), and (3) the top score achieved so far.

4.2. Final Move Selection

The four steps are repeated until time runs out.⁶ When this occurs, a final move selection is performed to determine which move should be played. There exist two standard final move selection strategies, which base their choice on the children of the root node. They are (1) choosing the child with *highest average score* and (2) choosing the child with *highest number of play-outs*. For puzzles however, the aim is to find the move with the highest *topscore*. We propose to use (3) the topscore for the final move selection. The highest scoring play-out is saved, and its first move is chosen for the final move selection. We investigate the effectiveness of this final move selection strategy in Subsection 6.5.

An essential difference between deterministic one-player games and two-player games is that there is no uncertainty on the opponent’s play. It is therefore possible to perform one large search from the initial position and then play all moves at once. With this approach all moves at the start are under

⁶In general, there is no time limitation for puzzles. However, a time limit is necessary to make testing possible.

consideration until the time for SP-MCTS runs out. It has to be investigated whether this approach outperforms an approach that allocates search time for every move. These experiments are presented in Subsection 6.3.

4.3. Randomized Restarts

We observed that it is important to generate deep trees in SameGame (see Subsection 6.2). However, by exploiting the most-promising lines of play, the SP-MCTS can be caught in local maxima. So, we randomly restart SP-MCTS with a different seed to overcome this problem. Because no information is shared between the searches, they explore different parts of the search space. This method resembles root parallelization (Chaslot et al., 2008c).

Root parallelization is an effective way of using multiple cores simultaneously (Chaslot et al., 2008c). However, we argue that root parallelization may also be used for avoiding local maxima in a single-threaded environment. Because there is no actual parallelization, we call this *randomized restarts*. Subsection 6.3 shows that randomized restarts are able to increase the average score significantly.

5. The Cross-Entropy Method

Choosing the correct SP-MCTS parameter values is important for its success. For instance, an important parameter is the C constant which is responsible for the balance between exploration and exploitation. Optimizing these parameters manually may be a hard and time-consuming task. Although it is possible to make educated guesses for some parameters, for other parameters it is not possible. Specially hidden dependencies between the parameters complicate the tuning process. Here, a learning method can be used to find the best values for these parameters (Beal and Smith, 2000; Sutton and Barto, 1998).

The *Cross-Entropy Method* (CEM) (Rubinstein, 2003) has successfully tuned parameters of an MCTS program in the past (Chaslot et al., 2008a). CEM is an evolutionary optimization method, related to *Estimation-of-Distribution Algorithms* (EDAs) (Mühlenbein, 1997). CEM is a population-based learning algorithm, where members of the population are sampled from a parameterized probability distribution (e.g., Gaussian, Binomial, Bernoulli, etc.). This probability distribution represents the range of possible solutions.

CEM converges to a solution by iteratively changing the parameters of the probability distribution (e.g., μ and σ for a Gaussian distribution). An iteration consists of three main steps. First, a set S of vectors $x \in X$ is drawn from the probability distribution, where X is some parameter space. These parameter vectors are called samples. In the second step, each sample is evaluated and gets assigned a fitness value. A fixed number of samples within S having the highest fitness are called the *elite samples*. In the third step, the elite samples are used to update the parameters of the probability distribution.

Generally, CEM aims to find the optimal solution x^* for a learning task described in the following form

$$x^* \leftarrow \underset{x}{\operatorname{argmax}} f(x), \quad (2)$$

where x^* is a vector containing all parameters of the (approximately) optimal solution. f is a fitness function that determines the performance of a sample x (for SameGame this is the average number of points scored on a set of positions). The main difference of CEM to traditional methods is that CEM does not maintain a single candidate solution, but maintains a distribution of possible solutions.

There exist two methods for generating samples from the probability distribution, (1) random guessing and (2) distribution focusing (Rubinstein, 2003). *Random guessing* straightforwardly creates samples from the distribution and selects the best sample as an estimate for the optimum. If the probability distribution peaked close to the global optimum, random guessing may obtain a good estimate. If the distribution is rather uniform, the random guessing is unreliable. After drawing a moderate number of samples from a distribution, it may be impossible to give an acceptable approximation of x^* , but it may be possible to obtain a better sampling distribution. To modify the distribution during sampling to form a peak around the currently best samples is called *distribution focusing*. Distribution focusing is the central idea of CEM (Rubinstein, 2003).

When starting CEM, an initial probability distribution is required. Chaslot et al. (2008a) used a Gaussian distribution and proposed that for each parameter, the mean μ of the corresponding distribution is equal to the average of the lower and upper bound of that parameter.⁷ The standard deviation σ is set to half the difference between the lower and upper bound (Chaslot et al., 2008a).

6. Experiments and Results

Subsection 6.1 tests the quality of the two simulation strategies TabuRandom and TabuColorRandom. Thereafter, the results of manual parameter tuning are presented in Subsection 6.2. Subsequently, Subsection 6.3 gives the performance of the randomized restarts on a set of 250 positions. In Subsection 6.4, it is investigated whether it is better to exhaust all available time at the initial position or to distribute the time uniformly for every move. Subsection 6.5 investigates which final move selection strategy produces the best results. Next, in Subsection 6.6 the parameter tuning by CEM is shown. Finally, Subsection 6.7 compares SP-MCTS to other approaches. All experiments use parameter values $W = 0.02$, $T = 10$ and $\epsilon = 0.003$, unless stated otherwise.

⁷The lower and upper bound of a variable are either domain dependant or based on intuition.

Table 1: Effectiveness of the simulation strategies.

	Random	TabuRandom	TabuColorRandom
Average Score	2,069	2,737	3,038
StdDev	322	445	479

6.1. Simulation Strategy

In order to test the effectiveness of the two simulation strategies, we used a test set of 250 randomly generated positions.⁸ We applied SP-MCTS without randomized restarts for each position until 10 million nodes were reached in memory. These runs typically take 5 to 6 minutes per position. The best score found during the search is the final score for the position. The constants C and D were set to 0.5 and 10,000, respectively. For the final move selection, the topscore strategy is used, for which the experiments are given in Subsection 6.5. The results are shown in Table 1.

Table 1 shows that the TabuRandom strategy has a significantly better average score (i.e., 700 points) than plain random. Using the TabuColorRandom strategy the average score is increased by another 300 points. We observe that a low standard deviation is achieved for the random strategy. In this case, it implies that all positions score almost equally low. The proposed TabuColorRandom strategy has also been successfully applied in Nested Monte-Carlo Search (Cazenave, 2009) and HGSTS (Edelkamp et al., 2010).

6.2. Manual Parameter Tuning

This subsection presents the parameter tuning in SP-MCTS. Three different settings were used for the pair of constants (C ; D) of Formula 1, in order to investigate which balance between exploitation and exploration gives the best results. These constants were tested with three different time controls on the test set of 250 positions, expressed by a maximum number of nodes. The short time control refers to a run with a maximum of 10^5 nodes in memory. At the medium time control, 10^6 nodes are allowed in memory, and for a long time control 5×10^6 nodes are allowed. We have chosen to use nodes in memory as measurement to keep the results hardware-independent. Based on preliminary manual tuning, we have selected three parameter pairs for further investigation. The parameter pair (0.1; 32) represents *exploitation*, (1; 20,000) performs *exploration*, and (0.5; 10,000) is a balance between the other two.

Table 2 shows the performances of the SP-MCTS approach for the three time controls. The short time control corresponds to approximately 20 seconds per position. In this table the term *Average Depth* is the average depth of a node and *Average Deepest Node* is the average depth of the deepest leaf node

⁸The test set can be found at <http://www.personeel.unimaas.nl/maartenschadd/TestSet.txt>

added to the tree, measured as the distance to the root node. The best results are achieved by exploitation. The average score is 2,552. With this setting the search is able to build trees that have on average the deepest leaf node at ply 63, implying that a substantial part of the chosen line of play is inside the SP-MCTS tree. Also, we observe that the other two settings are not generating a deep tree.

Table 2: Results of SP-MCTS for different settings.

10^5 nodes (~ 20 seconds)	Exploitation (0.1; 32)	Balanced (0.5; 10,000)	Exploration (1; 20,000)
Average Score	2,552	2,388	2,197
Standard Deviation	572	501	450
Average Depth	25	7	3
Average Deepest Node	63	19	8
10^6 nodes (~ 200 seconds)	(0.1; 32)	(0.5; 10,000)	(1; 20,000)
Average Score	2,674	2,858	2,579
Standard Deviation	607	560	492
Average Depth	36	14	6
Average Deepest Node	71	37	15
5×10^6 nodes ($\sim 1,000$ seconds)	(0.1; 32)	(0.5; 10,000)	(1; 20,000)
Average Score	2,806	3,008	2,901
Standard Deviation	576	524	518
Average Depth	40	18	9
Average Deepest Node	69	59	20

For the medium time control, the best results were achieved by using the balanced setting. It scores 2,858 points. Moreover, Table 2 shows that the average score of the balanced setting increases most compared to the short time control, viz. 470. The balanced setting is able to build substantially deeper trees than at the short time control (37 vs. 19). An interesting observation can be made by comparing the score of the exploration setting for the medium time control to the exploitation score for the short time control. Even with 10 times the amount of time, exploration is not able to achieve a significantly higher score than exploitation.

The results for the long experiment are that the balanced setting again achieves the highest score with 3,008 points. The deepest node for this setting is on average at ply 59. However, the exploitation setting only scores 200 points fewer than the balanced setting and 100 points fewer than exploration.

From the results presented we may draw two conclusions. First, it is important to have a deep search tree. Second, exploiting local maxima can be more advantageous than searching for the global maximum when the search only has a small amount of time.

6.3. Randomized Restarts

This subsection presents the performance tests of the randomized restarts on the set of 250 positions. We remark that the experiments are time constrained. Each experiment could only use 5×10^5 nodes in total and the restarts distributed these nodes uniformly among the number of searches. It means that a single search can take all 5×10^5 nodes, but that two searches can only use 2.5×10^5 nodes each. We used the exploitation setting (0.1; 32) for this experiment. The results are depicted in Figure 2.

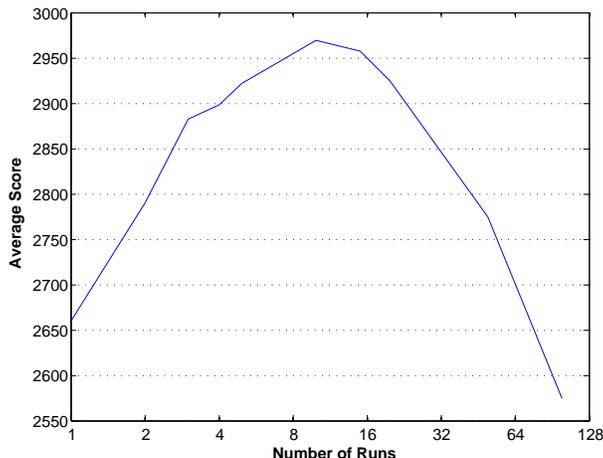


Figure 2: The average score for different settings of randomized restarts.

Figure 2 indicates that already with two searches instead of one, a significant performance increase of 140 points is achieved. Furthermore, the maximum average score of the randomized restarts is at ten searches, which uses 5×10^4 nodes for each search. Here, the average score is 2,970 points. This result is almost as good as the best score found in Table 2, but with the difference that the randomized restarts together used one tenth of the number of nodes. After 10 restarts the performance decreases because the generated trees are not deep enough.

6.4. Time Control

This subsection investigates whether it is better to exhaust all available time at the initial position or to distribute the time uniformly for every move. Table 3 shows the average score on 250 random positions with five different time settings. This experiment uses the exploitation settings for the C and D parameters. When SP-MCTS is applied for every move, this time is divided by the average game length (64.4). It means that depending on the number of moves, the total search time varies. These time settings are exact in the case that SP-MCTS is applied per game.

Table 3: Average score on 250 positions using different time control settings.

Time in seconds	~5	~10	~20	~30	~60
SP-MCTS per game	2,223	2,342	2,493	2,555	2,750
SP-MCTS per move	2,588	2,644	2,742	2,822	2,880

Table 3 shows that distributing the time uniformly for every move is the better approach. For every time setting a higher score is achieved when searching per move. The difference in score is largest for 5 seconds, and smallest for 60 seconds. It is an open question whether for longer time settings it may be beneficial to exhaust all time at the initial position.

6.5. Final Move Selection

In the previous experiments we used the topscore for the final move selection. To verify the correctness of this choice, we investigate in this subsection the effectiveness of the three final move selection strategies (1) highest average, (2) highest number of play-outs and (3) topscore. To measure the effectiveness of these strategies, we apply SP-MCTS per move, using 10 seconds per position and the exploitation setting. Table 4 shows the average score on the test set of 250 positions.

Table 4: Final move selection strategies.

	Highest Average	Highest Number of Play-outs	TopScore
Average Score	784	1,458	2,644

Using the average score for the final move selection produced rather disappointing results (784 points). Especially at the beginning of the game, play-outs are unlikely to achieve high scores or clear the board. This produces too much noise within the average score. Using the highest number of play-outs for the final move selection results in 1,458 points on average and is significantly better than using the average strategy. Using the topscore produced by far the best results (2,644 points). Even if a play-out achieves a high score only once, the score of this play-out gives a lower bound for the final score. To achieve this score, only the moves of this play-out have to be followed. We will use the topscore strategy for the final move selection in all remaining experiments.

6.6. CEM Parameter Tuning

In the next series of experiments we tune SP-MCTS with CEM. The following settings for CEM were used. The sample size is equal to 100, the number of elite samples is equal to 10. Each sample plays 30 games with 1 minute thinking time for each game. The 30 initial positions are randomly generated at the start of each iteration. The fitness of a sample is the average of the scores

of these games. The parameters tuned by CEM are presented in Table 5. C , D , T , W and ϵ were described in Subsection 4.1. The CEM-tuned parameters differ significantly from the manually tuned ones. The only exception is parameter T , which is the number of visits before a node is expanded. CEM found values 11 and 13 which are in the same range as the manually tuned value of 10. The CEM-tuned parameters are more explorative than the manually tuned parameters.

Table 5: Parameter tuning by CEM.

Parameter	Manual	CEM per game	CEM per move
C	0.1	5.96	4.31
D	32	67.98	96.67
T	10	13	11
W	0.02	0.49	0.28
ϵ	0.003	0.00007	0.00007 ⁹

To determine the performance of the parameters found by CEM an independent test set of 250 randomly created positions was used. Five different time settings were investigated. Table 6 shows the results of the CEM experiments.

Table 6: Performance of CEM tuning.

	Time in seconds	~ 5	~ 10	~ 20	~ 30	~ 60
Per Game	Manually tuning	2,223	2,342	2,493	2,555	2,750
	CEM tuned	2,405	2,493	2,598	2,734	2,804
Per Move	Manually tuning	2,588	2,644	2,742	2,822	2,880
	CEM tuned	2,652	2,749	2,856	2,876	2,913

Table 6 shows that both time controls and for every time setting CEM is able to improve the score. This demonstrates the difficulty of finding parameters manually in a high-dimensional parameter space.

6.7. Comparison on the Standardized Test Set

Using two hours per position, we tested IDA* and three different versions of SP-MCTS, subsequently called SP-MCTS(1), SP-MCTS(2), and SP-MCTS(3), on the standardized test set. IDA* applies the evaluation function described in Section 3. SP-MCTS(1) builds one large tree at the start and uses the exploitation setting (0.1; 32) and randomized restarts, which applied 1,000 runs using 100,000 nodes for each search thread. SP-MCTS(2) uses the same parameters as

⁹This parameter was not tuned again because it was obvious that the optimal weight is close to or equal to zero.

SP-MCTS(1), but distributes its time for every move. SP-MCTS(3) distributes its time per move and uses the parameters found by CEM.

Table 7 compares IDA* and the SP-MCTS variants with the other approaches, which were described in Subsection 2.3. The table shows that IDA* does not perform well for SameGame. It plays at human beginner level. Next, SP-MCTS(1) outperformed DBS on 11 of the 20 positions and was able to achieve a total score of 73,998. This was the highest score on the test set at that point of time (Schadd et al., 2008b). SP-MCTS(2) scored 76,352 points, 2,354 more than SP-MCTS(1). This shows that it is important to distribute search time for every move. SP-MCTS(3) achieved 78,012 points, the third strongest method at this point of time. All SP-MCTS versions are able to clear the board for all 20 positions.¹⁰ This confirms that a deep search tree is important for SameGame as shown in Subsection 6.2.

The two highest scoring programs (1) SPURIOUS AI and (2) HGSTS achieved more points than SP-MCTS. We want to give the following remarks on these impressive scores. (1) SPURIOUS AI is memory intensive and it is unknown what time settings were used for achieving this score. (2) HGSTS utilized the graphics processing unit (GPU), was optimized for every position in the standardized test set, and applied our TabuColorRandom strategy. Moreover, the scores of HGTS were not independently verified to be correct.

7. Conclusions and Future Research

In this article we proposed a new MCTS variant called Single-Player Monte-Carlo Tree Search (SP-MCTS). We adapted MCTS by two modifications resulting in SP-MCTS. The modifications are (1) in the selection strategy and (2) in the backpropagation strategy. Below we provide five observations and one conclusion.

First, we observed that our TabuColorRandom strategy significantly increased the score of SP-MCTS in SameGame. Compared to the pure random play-outs, an increase of 50% in the average score is achieved. The proposed TabuColorRandom strategy has also been successfully applied in Nested Monte-Carlo Search (Cazenave, 2009) and HGSTS (Edelkamp et al., 2010). Second, we observed that exploiting works better than exploring at short time controls. At longer time controls a balanced setting achieves the highest score, and the exploration setting works better than the exploitation setting. However, exploiting the local maxima still leads to comparable high scores. Third, with respect to the randomized restarts, we observed that for SameGame combining a large number of small searches can be more beneficial than performing one large search. Fourth, it is better to distribute search time equally over the consecutive positions than to invest all search time at the initial position. Fifth, CEM is able to find better parameter values than manually tuned parameter

¹⁰The best variations can be found at the following address:
<http://www.personeel.unimaas.nl/maarten-schadd/SameGame/Solutions.html>

Table 7: Comparing the scores on the standardized test set.

#	IDA*	DBS	SP-MCTS(1)	SP-MCTS(2)	MC-RWS
1	548	2,061	2,557	2,969	2,633
2	1,042	3,513	3,749	3,777	3,755
3	841	3,151	3,085	3,425	3,167
4	1,355	3,653	3,641	3,651	3,795
5	1,012	3,093	3,653	3,867	3,943
6	843	4,101	3,971	4,115	4,179
7	1,250	2,507	2,797	2,957	2,971
8	1,246	3,819	3,715	3,805	3,935
9	1,887	4,649	4,603	4,735	4,707
10	668	3,199	3,213	3,255	3,239
11	1,073	2,911	3,047	3,013	3,327
12	602	2,979	3,131	3,239	3,281
13	667	3,209	3,097	3,159	3,379
14	749	2,685	2,859	2,923	2,697
15	745	3,259	3,183	3,295	3,399
16	1,647	4,765	4,879	4,913	4,935
17	1,248	4,447	4,609	4,687	4,737
18	2,586	5,099	4,853	4,883	5,133
19	1,437	4,865	4,503	4,685	4,903
20	872	4,851	4,853	4,999	4,649
Sum	22,354	72,816	73,998	76,352	76,764
#	Nested MC	SP-MCTS(3)	SPURIOUS AI	HGSTS	
1	3,121	2,919	3,269	2,561	
2	3,813	3,797	3,969	4,995	
3	3,085	3,243	3,623	2,858	
4	3,697	3,687	3,847	4,051	
5	4,055	4,067	4,337	4,633	
6	4,459	4,269	4,721	5,003	
7	2,949	2,949	3,185	2,717	
8	3,999	4,043	4,443	4,622	
9	4,695	4,769	4,977	6,086	
10	3,223	3,245	3,811	3,628	
11	3,147	3,259	3,487	2,796	
12	3,201	3,245	3,851	3,710	
13	3,197	3,211	3,437	3,271	
14	2,799	2,937	3,211	2,432	
15	3,677	3,343	3,933	3,877	
16	4,979	5,117	5,481	6,074	
17	4,919	4,959	5,003	5,166	
18	5,201	5,151	5,463	6,044	
19	4,883	4,803	5,319	5,019	
20	4,835	4,999	5,047	5,175	
Sum	77,934	78,012	84,414	84,718	

values. The parameters found by CEM resemble a balanced setting. They were tuned for applying SP-MCTS for every move, causing that deep trees are less important.

The main conclusion is that we have shown that MCTS is applicable to a one-player deterministic perfect-information game. Our variant, SP-MCTS, is able to achieve good results in SameGame. Thus, SP-MCTS is a worthy alternative for puzzles where a good admissible estimator cannot be found.

There are three directions of future research for SP-MCTS. The first direction is to test several enhancements in SP-MCTS. We mention two of them. (1) The selection strategy can be enhanced with RAVE (Gelly and Silver, 2007) or progressive widening (Chaslot et al., 2008b; Coulom, 2007). (2) This article demonstrated that combining small searches can achieve better scores than one large search. However, there is no information shared between the searches. This can be achieved by using a transposition table, which is not cleared at the end of a small search. The second direction is to apply SP-MCTS to other domains. For instance, we could test SP-MCTS in puzzles such as Morpion Solitaire and Sudoku (Cazenave, 2009) and Single-Player General Game Playing (Méhat and Cazenave, 2010). Other classes of one-player games, with non-determinism or imperfect information, could be used as test domain for SP-MCTS as well. The third direction of future research is to apply SP-MCTS to optimization and planning problems outside the domain of games (e.g., web service decomposition (Okutan and Cicekli, 2010))

Acknowledgments

This work is funded by the Dutch Organisation for Scientific Research (NWO) in the framework of the project Go4Nature, grant number 612.000.938.

References

- Allis, L. V., 1994. Searching for Solutions in Games and Artificial Intelligence. Ph.D. thesis, Department of Computer Science, Rijksuniversiteit Limburg, Maastricht, The Netherlands.
- Bayili, S., Polat, F., 2011. Limited-Damage A*: A Path Search Algorithm that Considers Damage as a Feasibility Criterion. *Knowledge-Based Systems* 24 (4), 501–512.
- Beal, D. F., Smith, M. C., 2000. Temporal Difference Learning for Heuristic Search and Game Playing. *Information Sciences* 122 (1), 3–21.
- Biedl, T. C., Demaine, E. D., Demaine, M. L., Fleischer, R., Jacobsen, L., Munro, J. I., 2002. The Complexity of Clickomania. In: Nowakowski, R. J. (Ed.), *More Games of No Chance*, Proceedings of the 2002 MSRI Workshop on Combinatorial Games. Vol. 42 of MSRI Publications. Cambridge University Press, Cambridge, United Kingdom, pp. 389–404.
- Billings, D., 2007. Personal Communication, University of Alberta, Canada.

- Blum, A. L., Furst, M. L., 1997. Fast Planning through Planning Graph Analysis. *Artificial Intelligence* 90 (1–2), 281–300.
- Bouzy, B., 2005. Associating Domain-Dependent Knowledge and Monte-Carlo Approaches within a Go Program. *Information Sciences, Heuristic Search and Computer Game Playing IV* 175 (4), 247–257.
- Cazenave, T., 2009. Nested Monte-Carlo Search. In: Boutilier, C. (Ed.), *Proceedings of the Twenty-first International Joint Conferences on Artificial Intelligence (IJCAI-09)*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, pp. 456–461.
- Chaslot, G. M. J.-B., 2010. Monte-Carlo Tree Search. Ph.D. thesis, Department of Knowledge Engineering, Maastricht University, Maastricht, The Netherlands.
- Chaslot, G. M. J.-B., de Jong, S., Saito, J.-T., Uiterwijk, J. W. H. M., 2006a. Monte-Carlo Tree Search in Production Management Problems. In: Schobbens, P.-Y., Vanhoof, W., Schwanen, G. (Eds.), *Proceedings of the 18th BeNeLux Conference on Artificial Intelligence (BNAIC'06)*. University of Namur, Namur, Belgium, pp. 91–98.
- Chaslot, G. M. J.-B., Saito, J., Bouzy, B., Uiterwijk, J. W. H. M., van den Herik, H. J., 2006b. Monte-Carlo Strategies for Computer Go. In: Schobbens, P.-Y., Vanhoof, W., Schwanen, G. (Eds.), *Proceedings of the 18th BeNeLux Conference on Artificial Intelligence (BNAIC'06)*. University of Namur, Namur, Belgium, pp. 83–90.
- Chaslot, G. M. J.-B., Winands, M. H. M., Szita, I., van den Herik, H. J., 2008a. Cross-Entropy for Monte-Carlo Tree Search. *ICGA Journal* 31 (3), 145–156.
- Chaslot, G. M. J.-B., Winands, M. H. M., Uiterwijk, J. W. H. M., van den Herik, H. J., Bouzy, B., 2008b. Progressive Strategies for Monte-Carlo Tree Search. *New Mathematics and Natural Computation* 4 (3), 343–357.
- Chaslot, G. M. J.-B., Winands, M. H. M., van den Herik, H. J., 2008c. Parallel Monte-Carlo Tree Search. In: van den Herik, H. J., Xu, X., Ma, Z., Winands, M. H. M. (Eds.), *Computers and Games (CG 2008)*. Vol. 5131 of *Lecture Notes in Computer Science (LNCS)*. Springer-Verlag, Berlin, Germany, pp. 60–71.
- Chen, K.-H., Zhang, P., 2008. Monte-Carlo Go with Knowledge-Guided Simulations. *ICGA Journal* 31 (2), 67–76.
- Cook, S. A., 1971. The Complexity of Theorem-Proving Procedures. In: *STOC '71: Proceedings of the Third Annual ACM Symposium on Theory of Computing*. ACM Press, New York, NY, USA, pp. 151–158.

- Coulom, R., 2007. Efficient Selectivity and Backup Operators in Monte-Carlo Tree Search. In: van den Herik, H. J., Ciancarini, P., Donkers, H. H. L. M. (Eds.), *Computers and Games (CG 2006)*. Vol. 4630 of *Lecture Notes in Computer Science (LNCS)*. Springer-Verlag, Heidelberg, Germany, pp. 72–83.
- Culberson, J. C., Schaeffer, J., 1998. Pattern Databases. *Computational Intelligence* 14 (3), 318–334.
- Edelkamp, S., Kissmann, P., Sulewski, D., Messerschmidt, H., 2010. Finding the Needle in the Haystack with Heuristically Guided Swarm Tree Search. In: Schumann, M., Kolbe, L. M., Breitner, M. H., Frerichs, A. (Eds.), *Multikonferenz Wirtschaftsinformatik 2010*. Universitätsverlag Göttingen, Göttingen, Germany, pp. 2295–2308.
- Felner, A., Zahavi, U., Schaeffer, J., Holte, R. C., 2005. Dual Lookups in Pattern Databases. In: Kaelbling, L. P., Saffiotti, A. (Eds.), *Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence (IJCAI-05)*. Professional Book Center, Denver, CO, USA, pp. 103–108.
- Gelly, S., Silver, D., 2007. Combining Online and Offline Knowledge in UCT. In: Ghahramani, Z. (Ed.), *Proceedings of the 24th International Conference on Machine Learning (ICML '07)*. ACM Press, New York, NY, USA, pp. 273–280.
- Gelly, S., Wang, Y., Munos, R., Teytaud, O., 2006. Modification of UCT with Patterns in Monte-Carlo Go. Tech. Rep. 6062, INRIA, Orsay Cedex, France.
- Hart, P. E., Nielson, N. J., Raphael, B., 1968. A Formal Basis for the Heuristic Determination of Minimum Cost Paths. *IEEE Transactions on Systems Science and Cybernetics SSC-4* (2), 100–107.
- Johnson, D. S., 1990. A Catalog of Complexity Classes. In: van Leeuwen, J. (Ed.), *Handbook of Theoretical Computer Science*. Vol. A: Algorithms and Complexity. MIT Press, Cambridge, MA, USA, pp. 67–161.
- Julien, D., 2008. Pocket PC Jawbreaker Game. PDA Game Guide.com, <http://www.pdagameguide.com/jawbreaker-game.html>.
- Junghanns, A., 1999. Pushing the Limits: New Developments in Single Agent Search. Ph.D. thesis, Department of Computing Science, University of Alberta, Edmonton, Alberta, Canada.
- Kendall, G., Parkes, A. J., Spoerer, K., 2008. A Survey of NP-Complete Puzzles. *ICGA Journal* 31 (1), 13–34.
- Kocsis, L., Szepesvári, C., 2006. Bandit based Monte-Carlo Planning. In: Fürnkranz, J., Scheffer, T., Spiliopoulou, M. (Eds.), *Proceedings of the 17th European Conference on Machine Learning (ECML 2006)*. Vol. 4212 of *Lecture Notes in Computer Science (LNCS)*. Springer-Verlag, Berlin Heidelberg, Germany, pp. 282–293.

- Korf, R. E., 1985. Depth-First Iterative Deepening: An Optimal Admissible Tree Search. *Artificial Intelligence* 27 (1), 97–109.
- Lee, C.-S., Müller, M., Teytaud, O., 2010. Editorial of the Special Issue on Monte Carlo Techniques and Computer Go. *IEEE Transactions on Computational Intelligence and AI in Games* 2 (4), 225–228.
- Méhat, J., Cazenave, T., 2010. Combining UCT and Nested Monte-Carlo Search for Single-Player General Game Playing. *IEEE Transactions on Computational Intelligence and AI in Games* 2 (4), 271–277.
- Mesmay, F., Rimmel, A., Voronenko, Y., Püschel, M., 2009. Bandit-Based Optimization on Graphs with Application to Library Performance Tuning. In: Danyluk, A. P., Bottou, L., Littman, M. L. (Eds.), *Proceedings of the 26th Annual International Conference on Machine Learning (ICML'09)*. ACM, New York, NY, USA, pp. 729–736.
- Moribe, K., 1985. Chain Shot! *Gekkan ASCII* (November issue), in Japanese.
- Mühlenbein, H., 1997. The Equation for Response to Selection and its Use for Prediction. *Evolutionary Computation* 5 (3), 303–346.
- Okutan, C., Cicekli, N. K., 2010. A Monolithic Approach to Automated Composition of Semantic Web Services with the Event Calculus. *Knowledge-Based Systems* 5 (23), 440–454.
- Rubinstein, R. Y., 2003. The Cross-Entropy Method for Combinatorial and Continuous Optimization. *Methodology and Computing in Applied Probability* 1 (2), 127–190.
- Sadikov, A., Bratko, I., 2007. Solving 20×20 Puzzles. In: van den Herik, H. J., Uiterwijk, J. W. H. M., Winands, M. H. M., Schadd, M. P. D. (Eds.), *Proceedings of the Computer Games Workshop 2007 (CGW 2007)*. Maastricht University, Maastricht, The Netherlands, pp. 157–164.
- Schadd, M. P. D., Winands, M. H. M., van den Herik, H. J., Aldewereld, H., 2008a. Addressing NP-Complete Puzzles with Monte-Carlo Methods. In: *Proceedings of the AISB 2008 Symposium on Logic and the Simulation of Interaction and Reasoning*. Vol. 9. The Society for the Study of Artificial Intelligence and Simulation of Behaviour, Brighton, United Kingdom, pp. 55–61.
- Schadd, M. P. D., Winands, M. H. M., van den Herik, H. J., Chaslot, G. M. J.-B., Uiterwijk, J. W. H. M., 2008b. Single-Player Monte-Carlo Tree Search. In: van den Herik, H. J., Xu, X., Ma, Z., Winands, M. H. M. (Eds.), *Computers and Games (CG 2008)*. Vol. 5131 of *Lecture Notes in Computer Science (LNCS)*. Springer-Verlag, Berlin, Germany, pp. 1–12.
- Stewart, B., White, C., 1991. Multiobjective A*. *Journal of the Association for Computing Machinery* 38 (4), 93–105.

- Sturtevant, N. R., 2008. An Analysis of UCT in Multi-Player Games. In: van den Herik, H. J., Xu, X., Ma, Z., Winands, M. H. M. (Eds.), *Computers and Games (CG 2008)*. Vol. 5131 of *Lecture Notes in Computer Science (LNCS)*. Springer-Verlag, Berlin, Germany, pp. 37–49.
- Sutton, R. S., Barto, A. G., 1998. *Reinforcement Learning: An Introduction*. MIT Press, Cambridge, MA, USA.
- Tak, M. J. W., 2010. *The Cross-Entropy Method Applied to SameGame*. Bachelor thesis, Maastricht University, Maastricht, The Netherlands.
- Takes, F. W., Kusters, W. A., 2009. Solving SameGame and its Chessboard Variant. In: Calders, T., Tuyls, K., Pechenizkiy, M. (Eds.), *Proceedings of the 21st Benelux Conference on Artificial Intelligence (BNAIC'09)*. Technische Universiteit Eindhoven, Eindhoven, The Netherlands, pp. 249–256.
- Vempaty, N. R., Kumar, V., Korf, R. E., 1991. Depth-First versus Best-First Search. In: *Proceedings of the 9th National Conference on Artificial Intelligence (AAAI-91)*. AAAI Press, Menlo Park, CA, USA, pp. 434–440.