

Monte-Carlo Tree Search for the Simultaneous Move Game Tron

Niek G.P. Den Teuling and Mark H.M. Winands

Games and AI Group, Department of Knowledge Engineering,
Faculty of Humanities and Sciences,
Maastricht University, Maastricht, The Netherlands
{n.denteuling@student., m.winands@}maastrichtuniversity.nl

Abstract. Monte-Carlo Tree Search (MCTS) has been successfully applied to many games, particularly in Go. In this paper, we investigate the performance of MCTS in Tron, which is a two-player simultaneous move game. We try to increase the playing strength of an MCTS program for the game of Tron by applying several enhancements to the selection, expansion and play-out phase of MCTS.

Based on the experiments, we may conclude that Progressive Bias, altered expansion phase and play-out cut-off all increase the overall playing strength, but the results differ per board. MCTS-Solver appears to be a reliable replacement for MCTS in the game of Tron, and is preferred over MCTS due to its ability to search the state space for a proven win. The MCTS program is still outperformed by the best $\alpha\beta$ program A1K0N, which uses a sophisticated evaluation function, indicating that there is quite some room for improvement.

1 Introduction

The classic way of exploring the game tree is using $\alpha\beta$ -search [9] with a domain-specific evaluation function. However, for games that require a complex positional evaluation function, this approach might not be the best way. An alternative approach is Monte-Carlo Tree Search (MCTS) [6, 10]. In contrast to $\alpha\beta$ -search, MCTS does not require a positional evaluation function as it relies on stochastic simulations. MCTS has proven itself to be a viable alternative in, for instance, the board games Go [6], Hex [1], Amazons [11] and Lines of Action [19].

A challenging new game is Tron. It is a two-player game that bears resemblance to Snake, except that in Tron, players leave a wall behind at each move. An interesting aspect of Tron is that it is a simultaneous move game, rather than the usual turn-taking game. In 2010, the University of Waterloo Computer Science Club organized an AI tournament for the game of Tron [18]. Overall, the MCTS programs were outperformed by $\alpha\beta$ programs. In this paper, the performance of MCTS in Tron is investigated, continuing the pioneering work performed by Samothrakis *et al.* [12]. We examine approaches to improve the program's playing strength, by trying out different evaluation functions and




Fig. 1. A Tron game on a board with obstacles after 13 moves. The blue player (1) has cut off the upper part of the board, restricting the space the red player (2) can fill.

MCTS enhancements. Our MCTS program equipped with the enhancements is subsequently matched against the top $\alpha\beta$ program *AIKON*.

This paper is organized as follows. Section 2 gives a brief description of the game of Tron and the difficulties that a program has to be able to handle, to play at a decent skill level. Section 3 explains the MCTS and MCTS-Solver method applied to Tron. The enhancements applied to MCTS regarding the selection strategy are described in Section 4, followed by an enhanced expansion strategy in Section 5. Play-out strategies are described in Section 6. Experiments and results are given in Section 7. Finally, in Section 8, conclusions from the results are drawn and future research is suggested.

2 The Game of Tron

The game of Tron originates from the movie *Tron*, released by Walt Disney Studios in 1982. The movie is about a virtual world where motorcycles drive at a constant speed and can only make 90° angles, leaving a wall behind them as they go. The game of Tron investigated in this paper is a board version of the game played in the movie.

Tron is a two-player board game played on an $m \times n$ grid. It is similar to Snake: each player leaves a wall behind them as they move. In Snake, the player's wall is of a limited length, but Tron does not have such a restriction. At each turn, the red and blue player can only move one square straight ahead, or to the left or right. Both players perform their moves at the same time; they have no knowledge of the other player's move until the next turn. Players cannot move to a square that already contains a wall. If both players move to the same square, it is considered a draw. If a player moves into a wall, he loses and the other player wins. Usually the boards are symmetric, such that none of the players has an advantage over the other player. A typical board size is 13×13 .

The game is won by outlasting your opponent such that the opponent has no moves left other than moving into a wall. At the early stage of the game, it is difficult to find good moves as the number of possible move sequences is quite large and it is difficult to predict what the opponent will do. Boards can contain

obstacles (see Figure 1), further increasing the difficulty of the game because filling the available space becomes a more difficult task. Obstacles can provide opportunities to cut off an opponent, reducing the opponent's free space while maximizing your own.

3 Monte-Carlo Tree Search

Monte-Carlo Tree Search (MCTS) is a best-first search method that constructs a search tree using many simulations (called play-outs) [6, 10]. Play-outs are used to evaluate a certain position. Positions with a high winning percentage are preferred over those with a lower winning percentage. MCTS constructs a search tree consisting of nodes, where each node represents a position of the game. Each node i has a value v_i and a visit count n_i . The search starts from the root node, which represents the current position. The tree is explored at random, but as the number of simulated nodes increases, it gains a better evaluation of the nodes and can focus on the most promising nodes.

Although Tron is a simultaneous move game, it is possible to represent it as a turn-taking game. The player under consideration is always first to move inside the tree, followed by the other player. An issue arises when the players can run into each other. A solution is given in Subsection 3.1.

MCTS is divided into four phases [5]. These phases are executed until time is up. The phases are illustrated in Figure 2. We explain the phases in detail below.

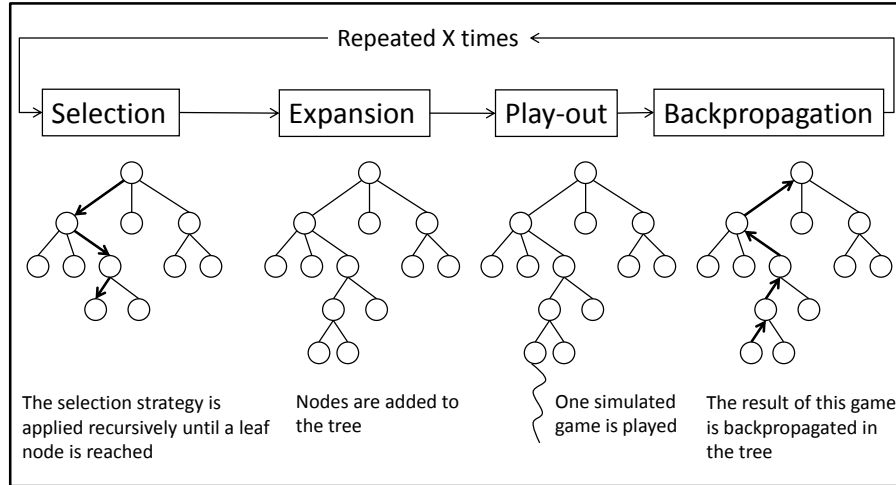


Fig. 2. Outline of the Monte-Carlo Tree Search [5].

Selection. In the selection phase, a child node of a given node is selected according to some strategy until a leaf node is reached. The selection task is an important one, as the goal is to find the best move. Because moves are evaluated by simulation, promising nodes should be played (exploited) more often than unpromising nodes. However, to find these nodes, unvisited nodes have to be tried out as well (exploration). Considering that in Tron, a player has at most 3 different moves at any turn (except for the first turn, where there could be 4 moves), this is not a problem. Because the number of simulations that can be performed is limited, a good balance has to be found between exploring and exploiting nodes. The simplest selection strategy is selecting a child node at random. A selection strategy that provides a balance between exploration and exploitation, is UCT (Upper Confidence Bound applied to Trees) [10]. It is based on the UCB1 algorithm (Upper Confidence Bound) [2]. UCT selects a child node k from node p as follows:

$$k \in \operatorname{argmax}_{i \in I} \left(v_i + C \times \sqrt{\frac{\ln n_p}{n_i}} \right) \quad (1)$$

C is a constant, which has to be tuned experimentally. Generally, UCT is applied after the node has first been visited a certain number of times T , to ensure all nodes have been sufficiently explored to apply UCT. If a node has a visit count less than T , the random-selection strategy is applied [6].

Expansion. In the expansion phase, the selected leaf node p is expanded. Since the number of child nodes is at most 3 in Tron, all nodes are added. The selection strategy is then applied to node p , returning the node from which the play-out starts.

Play-out. In this phase, the game is simulated in self-play, starting from the position of the selected node. Moves are performed until the game ends, or when the outcome can be estimated reliably. In contrast to the selection phase, both players move simultaneously in the play-out phase. The strategy used for selecting the moves to play can either be performing random moves, or using domain-specific knowledge that increases the quality of the simulation. The play-out phase returns a value of 1, 0 or -1 for the play-out node p , depending on whether the simulated game resulted in a win, draw, or loss, respectively. The same values are awarded to terminal nodes in the search tree. If the play-out node belongs to the player under consideration, the other player will first perform a move, such that both players have performed the same number of moves.

Backpropagation. The result of the simulation is backpropagated from the leaf node all the way back to the root node of the search tree. The values of the nodes are updated to match the new average of the play-outs.

After the simulations, the final move to be played by the MCTS program has to be chosen. The move is selected by taking the most ‘secure’ child of the root node [5]. The secureness of a node i is defined as: $v_i + \frac{A}{\sqrt{n_i}}$, where A is a constant. In the experiments, based on trial-and-error for the MCTS program, $A = 1$ is used.

3.1 Handling Simultaneous Moves

Treating Tron as a turn-taking game inside the tree works out quite well in almost every position of the game. However, if a position arises where the MCTS program has the advantage, but the players are at risk of crashing into each other, the program might play the move that leads to a draw. This happens because inside the search tree, the root player is always the first to move (as done in [12]). Because the root player already moved to the square that was accessible to both, the non-root player can no longer move to this square.

This problem is solved by adding an enhancement to the expansion strategy: if a node n belongs to the root player, and the non-root player could have moved to the square the root player is currently at, a terminal node is added to n with the value of a draw (i.e., 0). An example is shown in Figure 3.




Fig. 3. A game tree of Tron. In the left-most branch, both players moved to the same square, resulting in a terminal node that ends in a draw.

3.2 Monte-Carlo Tree Search Solver

Monte-Carlo Tree Search Solver (MCTS-Solver) [19] is an enhancement for MCTS that is able to prove the game-theoretic value of a position. MCTS combined with UCT may require significantly more time to find the best move, because it requires a large number of play-outs to converge to the game-theoretic value. Running MCTS-Solver requires a negligible amount of additional computation time on top of MCTS. Since proven positions do not have to be evaluated again, time can be spent on other positions that have not been proven yet. The original MCTS-Solver only considered win and loss outcomes, because in the test domain Lines of Action draws are exceptional [19]. Since draws occur often in Tron, an enhanced MCTS-Solver is used that does handle draws.

The *Score-Bounded MCTS-Solver* [4] extends MCTS-Solver to games that have more than two game-theoretic outcomes. It attaches an interval to each node, as done in the B* algorithm [3]. The interval is described by a pessimistic and optimistic bound. The pessimistic score represents the lowest achievable outcome for the root player, and the optimistic score represents the best achievable

outcome for the root player. Given sufficient time, the pessimistic and optimistic score of a node n will converge to its true value. An advantage of Score-Bounded MCTS is that the bounds enable pruning as in $\alpha\beta$ search, skipping unpromising branches. The initial bound of a node is set to $[-1.0, 1.0]$. Score-Bounded MCTS has been shown to solve positions considerably faster than MCTS-Solver [4].

3.3 Progressive Bias

Although UCT gives good results compared to other selection strategies that do not use knowledge of the game, there is room for improvement. The random strategy applied when the visit count of the node is small, can be replaced by the more promising play-out strategy. Since the accuracy of the UCT selection strategy increases as the number of visits increases, it is desirable to introduce a so-called ‘progressive strategy’. The progressive strategy provides a soft transition between the two strategies [5]. A popular choice is *progressive bias* (PB) that combines heuristic knowledge with UCT to select the best node [5]. By using Tron knowledge, node selection can be guided in a more promising direction, one that might not have been found by using play-outs only. This domain knowledge can be computationally expensive. A trade-off has to be made between simulating more games and spending more time for computing the domain knowledge.

When few games have been played, the heuristic knowledge has a major influence on the decision. The influence gradually decreases when the node is visited more often. The PB formula is as follows: $\frac{W \times P_{mc}}{l_i + 1}$ [19]. W is a constant (set to 10 in the experiments). P_{mc} is the transitional probability of a move category mc [17]. l_i denotes the number of losses in node i , this way, nodes that do not turn out well are not biased for too long. The formula of UCT and PB combined is:

$$k \in \operatorname{argmax}_{i \in I} \left(v_i + C \times \sqrt{\frac{\ln n_p}{n_i}} + \frac{W \times P_{mc}}{l_i + 1} \right) \quad (2)$$

The transitional probability of each move category is acquired from games played by expert players. Since no such games are available, the probabilities are obtained from self-play experiments of the MCTS program. The transitional probability P_{mc} of a move belonging to the move category mc is given by:

$$P_{mc} = \frac{n_{played}(mc)}{n_{available}(mc)} \quad (3)$$

$n_{played}(mc)$ denotes the number of positions in which a move belonging to move category mc was played. $n_{available}(mc)$ is the number of positions where a move belonging to move category mc could have been played.

We distinguish six move features in Tron:

- Passive: The player follows the wall that it is currently adjacent to.
- Offensive: The player moves towards the other player, when close to each other.

- Defensive: The player moves away from the other player, when close to each other.
- Territorial: The player attempts to close off a space by moving across open space towards a wall.
- Reckless: The player moves towards a square where the other player could have moved to, risking a draw.
- Obstructive: The player moves to a square that contains paths to multiple subspaces, closing off these spaces (at least locally).

4 Heuristic Knowledge in Tron

Estimating the remaining available space of a player is a useful heuristic in Tron because the game is won by filling a larger space than the opponent. *Space estimation* is valuable when, for instance, the program is at a square where it has to choose between two spaces. Biasing the selection towards moving to larger spaces saves time on simulating less-promising nodes that lead towards smaller spaces. We only focus on estimating the number of moves a player can make in a space that is not reachable by the other player.

Counting the number of empty squares does not always give an accurate estimation of the number of moves a program requires to fill the available space. Spaces can contain squares that can be reached, but offer no path back to fill the rest of the space.

One way to get an estimation of the available space is by filling up the space in a single-player simulation, and counting the number of moves [13]. The simulation uses a greedy *wall-following heuristic*. This heuristic works as follows. A move is selected such that the player moves to a square that lies adjacent to one or more walls (excluding the wall at its current square). If any of the moves cuts the player off from the other possible moves, the available space of each move is estimated and the move leading having the largest available space is selected. If there are multiple possible moves of equal score, a move is selected at random. This method does not always give the correct number of moves, but it gives a good lower bound on the amount of available space.

Instead of counting the number of empty squares, a tighter upper bound can be obtained by treating the board as a checker board. The difference in the number of grey and white squares gives an indication of the number of moves that can be performed [7]. The estimated number of moves M is computed by: $M = Z - |c_g - c_w|$. Z is the total available space, c_g is the number of grey tiles (including the one the player is currently standing on), and c_w is the number of white tiles. The estimated number of moves can be substantially off if the space contains large subspaces that offer no way back to other subspaces. Three example spaces are shown in Figure 4. The estimated number of moves for boards a , b and c are 13, 9 and 11, respectively. The true number of moves for the boards are 13, 9 and 10. Note that the estimation of board c is off because of the two separated subspaces. Had Z been solely used as the move estimation, the estimated number of moves would have been 14, 10 and 12.




Fig. 4. Three example boards where the player is isolated from the other player. The player starts at square S .

5 Predictive Expansion Strategy

In Tron, players will often get separated from each other (if not, the game ends in a draw). If a game is in such a position, the outcome of a game can be predicted in a reliable way. The *predictive expansion* strategy uses space estimation to predict the outcome of a position. It has to be noted that only nodes of the non-root player are evaluated, so both players have performed their moves when the position is evaluated.

If the outcome of a node can be predicted with certainty, the node does not have to be simulated, and is treated as a terminal node. This works as follows. The node candidate for expansion is evaluated using the space estimation heuristic. If there is a way for the players to reach each other, the node is expanded in the default way. If the players are isolated from each other and the outcome can be predicted, the node is not expanded and it is treated as a terminal node. The result of the prediction is backpropagated.

This strategy has two advantages over the old expansion strategy. First, applying space estimation is faster than performing a play-out. If a sufficient number of games are cut off, the time spent on space estimation is regained, and more time can be spent on searching through other parts of the tree. Second, the outcome prediction is more reliable than performing multitudes of play-outs. This prevents the program from underestimating positions where one or more players are closed off in a large space.

Once the game reaches the endgame phase, the enhanced expansion strategy is no longer applied since the MCTS program has shown to be capable of efficiently filling up the remaining space.

6 Play-out Strategies

The simplest play-out strategy is the *random-move* strategy. During the play-out, players perform a random move. Moves that result in an immediate defeat are excluded from the selection.

The authors observed that play-outs performed using a random-move strategy give surprisingly good results in Tron, considering that the randomly moving

Tron programs frequently trap themselves. The reliability of the play-outs can be further increased by using a more advanced play-out strategy [8].

We propose six play-out strategies. The resulting playing strength of each of these strategies is determined in the experiments.

Wall-following strategy. This strategy is inspired by the wall-following heuristic described in Section 4. The strategy selects the move leading to the square with the most number of walls (but smaller than 3). If multiple of moves lead to squares of the same number of walls, one of the moves is randomly selected. A problem with the wall-following strategy is that it does not leave much room for a rich variety of simulations. During each play-out, the moves performed will roughly be the same. It means that running more simulations does not necessarily increase the accuracy of the move value.

Offensive strategy. The offensive strategy selects moves that bring the player closer to the opponent player. If more than one move brings the player closer, one of the moves is selected at random. If there is no move that brings it closer to the opponent, a random move is performed.

Defensive strategy. This play-out strategy selects the move that increases the distance to the opponent player. If there is no such move, a random move is performed. If more than one move increases the distance from the opponent player, one of the moves is played at random.

Mixed strategy. The mixed strategy is a combination of the random play-out strategy and the previously mentioned strategies. At each move, a strategy is randomly selected according to a certain probability. The reasoning behind this strategy is that none of the strategies are particularly strong, and combining them may give better results.

The wall-following strategy has a 50% probability of being played, whereas the random-move, defensive and offensive are played 20%, 25% and 5% of the time, respectively.

Move-category strategy. This strategy uses the move category statistics used by the Progressive Bias enhancement to select a move. Moves are selected by roulette-wheel selection.

ϵ -greedy strategy. This strategy has a probability of $1 - \epsilon$ (i.e. 90%) of playing the wall-following strategy, and a probability of ϵ (i.e. 10%) of playing a random other play-out strategy [14, 15].

6.1 Endgame Strategies

The game of Tron can be split into two phases: the phase where players try to maximize their own available free space, and the phase where the players are isolated and attempt to outlast the other player by filling the remaining space as efficiently as possible, referred to as the endgame phase.

During the endgame phase, the same play-out strategies as mentioned above can be used, with the exception of the offensive and defensive strategy since there is no point in biasing the move on the position of the other player.

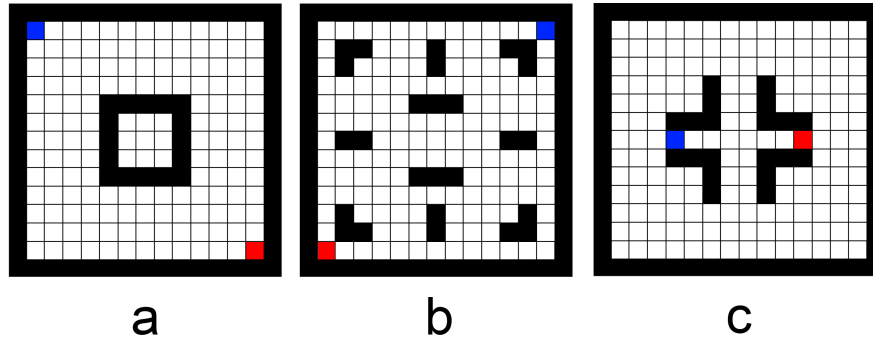


Fig. 5. The three boards used in the experiments.

6.2 Play-out Cut-Off

Although a Tron game is guaranteed to terminate, as each move brings the game moves closer to a terminal position, the number of moves performed during the play-out phase can be reduced. This saves time, leaving room for more simulations. Using heuristic knowledge, the result of a game can be predicted without the need to completely simulate it. A major problem with applying heuristic knowledge is that it costs much computation time compared to playing moves. Therefore, the positions are only evaluated once every 5 moves. An additional advantage of predicting the play-out outcome is that the accuracy of the play-outs is increased, because the player with the largest space can still lose a portion of the simulated games due to the weak play of the play-out strategies. The heuristic to predict the outcome of a game is the same as used in Subsection 4.

7 Experiments and Results

In this section, the proposed enhancements of Section 4, 5 and 6 are tested in our MCTS program [16]. The experiments are conducted on a 2.4 GHz AMD Opteron CPU with 8 GB of RAM. Experiments are conducted for three different boards, each providing different difficulties for MCTS. The boards are shown in Figure 5. Although all boards are symmetric, experiments are run for both colours to eliminate the possibility that the playing strength of the program is affected by its colour.

The following settings are used for the experiments, unless mentioned otherwise. In each experiment, 100 games are played on each board for both setups. In total 600 games are played. The time players can spend on computing the next move is 1 second. The programs have no knowledge about the move the other program will be performing.

Progressive Bias is tested in Subsection 7.1. Next, Subsection 7.2 describes the results of the various play-out strategies. Subsequently, Subsection 7.3 evaluates the play-out cut-off enhancement. The Predictive Expansion strategy is reported in Subsection 7.4. The playing strength of MCTS-Solver is tested in

Subsection 7.5. Finally, MCTS programs combined with the enhancements are tested against the winning program of the Tron Google AI Challenge in Subsection 7.6.

7.1 Progressive Bias

In the first series of experiment the Progressive Bias (PB) strategy is tested for different values of W . The MCTS-PB program is tested against the MCTS-UCT program [16]. Table 1 shows the transitional probabilities, obtained from 600 self-play games by the MCTS-UCT program at the three boards (200 games per board).

Table 2 shows that Progressive Bias does not improve the playing strength at boards a and b for any of the tested values of W , but it noticeably increases the playing strength at board c . Furthermore, the exact value of W does not seem to matter that much.

Table 1. Move categories and their respective transitional probabilities. **Table 2.** Win rates of MCTS-PB vs. MCTS-UCT.

Move category	P_{mc}
Defensive	28.9%
Defensive/Territorial	36.2%
Offensive	29.6%
Offensive/Reckless	0.0%
Offensive/Territorial	21.6%
Passive/Defensive	77.1%
Passive/Defensive/Obstructive	92.8%
Passive/Offensive	78.9%
Passive/Offensive/Obstructive	86.5%
Passive/Offensive/Reckless	6.6%
Passive/Off./Obs./Reck.	15.9%

W	Board a	Board b	Board c	Total
0.5	45%	48%	65%	53 ± 4 %
1	49%	45%	67%	54 ± 4 %
5	55%	45%	63%	54 ± 4 %
10	48%	49%	63%	53 ± 4 %
20	36%	52%	66%	51 ± 4 %

7.2 Play-out Strategy Experiments

In the next of series of experiments the various play-out strategies proposed are tested in the MCTS program. They are matched against each other in a round-robin tournament. The tournament is run at all three boards. Table 3 gives the results for each board. Table 4 shows the averaged results over all three boards.

The results show that the boards have a large influence on the effectiveness of the strategies. As such, it is difficult to select the best strategies based on these results. Overall, the random-move, defensive and wall-following strategies seem to be the best strategies. The random-move strategy performs well for all boards, whereas the defensive strategy stands out for board c . The wall-following strategy works well against the random-move strategy on board b . The random-move strategy is used in the experiments of the next subsections.

Table 3. Play-out strategy results for board *a*, *b* and *c*.

Board <i>a</i>	Rand.	Wall	Off.	Def.	Mixed	Cat.	ϵ -g.
Random		58%	90%	83%	71%	67%	52%
Wall	42%		90%	52%	21%	40%	26%
Offensive	10%	10%		31%	16%	13%	12%
Defensive	17%	48%	69%		55%	35%	32%
Mixed	29%	79%	84%	45%		49%	28%
Category	33%	60%	87%	65%	51%		38%
ϵ -greedy	48%	74%	88%	68%	72%	62%	
Board <i>b</i>	Rand.	Wall	Off.	Def.	Mixed	Cat.	ϵ -g.
Random		15%	100%	70%	56%	53%	58%
Wall	85%		99%	50%	78%	86%	88%
Offensive	0%	1%		0%	0%	0%	0%
Defensive	30%	50%	100%		77%	96%	51%
Mixed	44%	22%	100%	23%		37%	45%
Category	47%	14%	100%	4%	63%		36%
ϵ -greedy	42%	12%	100%	49%	55%	64%	
Board <i>c</i>	Random	Wall	Off.	Def.	Mixed	Cat.	ϵ -g.
Random		91%	98%	18%	43%	50%	49%
Wall	9%		81%	15%	7%	44%	9%
Offensive	2%	19%		9%	20%	12%	15%
Defensive	82%	85%	91%		76%	80%	90%
Mixed	57%	93%	80%	24%		40%	50%
Category	50%	56%	88%	20%	60%		57%
ϵ -greedy	51%	91%	85%	10%	50%	43%	

7.3 Play-out Cut-off

The play-out cut-off (PC) enhancement is tested by matching the MCTS-PC program against the MCTS-UCT program. MCTS-PC runs considerably fewer play-outs (25,000 per second on average) due to the computation time required by the play-out cut-off heuristic.

Table 5 shows the win rate for MCTS-PC against MCTS-UCT. At boards *a* and *b*, 800 games were run to ensure that the observed win rate was not influenced too much by statistical noise. 400 games were run for board *c*. The bad performance at board *c* might have to do with the difficulty for a player to isolate itself on this board.

7.4 Expansion Strategy Experiments

In the subsequent series of experiments are conducted for the predictive expansion (PDE) strategy. The MCTS-PDE program is tested against the MCTS-UCT program. The MCTS-PDE program ran 60,000 play-outs per second on average. The results are shown in Table 6. For each board, 600 games were run.

Similar to the results of the play-out cut-off experiment, the MCTS-PDE program appears to be slightly better than the MCTS-UCT program at board *a*

Table 4. Averaged play-out strategy results for all boards.

	Random	Wall	Off.	Def.	Mixed	Cat.	ϵ -g.
Random		55%	96%	57%	56%	57%	53%
Wall	45%		90%	39%	35%	57%	41%
Offensive	4%	10%		13%	12%	8%	9%
Defensive	43%	61%	87%		69%	71%	58%
Mixed	44%	65%	88%	31%		42%	41%
Category	43%	43%	92%	29%	58%		44%
ϵ -greedy	47%	59%	91%	42%	59%	56%	

Table 5. Win rates of MCTS-PC vs. MCTS-UCT.

	Board <i>a</i>	Board <i>b</i>	Board <i>c</i>	Total
Win	54%	56%	33%	$48 \pm 2\%$

and *b*. The poor win rate at board *c* is likely caused by the behaviour of MCTS for this board. The programs spiral around the centre, leaving the outer edges of the board open. Because the space estimation heuristic used by the predictive expansion strategy is only applicable when the players are isolated from each other, the MCTS-PDE program is squandering computation time on a mostly useless heuristic. Since the MCTS-UCT program spends all of its time on the play-outs, it can look further ahead and therefore has an advantage over the MCTS-PDE program.

7.5 MCTS-Solver Experiments

In this series of experiments the Score-Bounded MCTS-Solver is tested. The MCTS-Solver program ran at the same speed as the MCTS-UCT program. In the experiments of MCTS-Solver, MCTS-Solver-PDE and MCTS-Solver-PDE-PC, 400 games were run on each board.

As shown in Table 7, the MCTS-Solver program shows a slight improvement over MCTS-UCT. MCTS-Solver in combination with PDE or PDE-PC performs poorly at board *a* because it tends to cut off one side of the board, and due to lower number of simulations, cannot look ahead far enough to see the resulting outcome (i.e. loss). PDE and PDE-PC perform well at boards *b* and *c*, probably due to the obstacles and mistakes made by MCTS-UCT.

In terms of computation time and playing style, the MCTS-Solver program is the preferred choice since it requires no additional computations once a move has been proven to lead to a guaranteed win (or draw, when this is the best achievable outcome). Furthermore, the program can look up and play the shortest move sequence leading to a win by searching for the shortest winning path in the tree.

7.6 Playing against an $\alpha\beta$ program

In the final series of experiment, the MCTS program is tested against the winning program of the Tron Google AI Challenge, A1K0N [13]. The A1K0N program uses

Table 6. Win rates of MCTS-PDE vs. MCTS-UCT.

	Board <i>a</i>	Board <i>b</i>	Board <i>c</i>	Total
Win	53%	58%	48%	53 ± 2 %

Table 7. Win rates of MCTS-Solver variants against MCTS-UCT.

Win	Board <i>a</i>	Board <i>b</i>	Board <i>c</i>	Total
Solver	50%	52%	57%	53 ± 3 %
Solver-PDE	32%	74%	53%	53 ± 3 %
Solver-PDE-PC	30%	82%	70%	61 ± 3 %

$\alpha\beta$ -search [9] together with a good evaluation function that is primarily based on the tree of chambers heuristic.

Table 8. Win rates of various MCTS players against A1K0N.

	Board <i>a</i>	Board <i>b</i>	Board <i>c</i>	Total
MCTS-UCT	40%	0%	0%	14 ± 3 %
MCTS-PDE	44%	0%	0%	15 ± 3 %
Solver-PDE	13%	12%	0%	8 ± 3 %
Solver-PDE-PC	28%	10%	16%	18 ± 3 %
Solver-PDE-PC-PB	12%	18%	0%	10 ± 3 %

As can be seen in Table 8, A1K0N is the stronger player by far, achieving a win rate of 82% against MCTS-Solver-PDE-PC and a win rate higher than 85% against the other MCTS programs. Applying PB to MCTS-Solver does not seem to give an improvement in overall playing strength. Although the MCTS-Solver reaches a decent level of play, it still makes mistakes, mainly due to the fact that the reliability of the play-outs rapidly drops as the players get more distant from each other. By the time MCTS-Solver sees that it is in a bad position, it is already too late to correct.

8 Conclusion and Future Research

In this paper we developed an MCTS program for the game of Tron. Several enhancements were made to the selection, expansion and play-out phase. All of the enhancements were tested against an MCTS-UCT program.

The enhancement made to the selection phase, the progressive bias strategy, showed no improvement over UCT at two out of three boards. At board *c* the enhancement scored a consistent win rate of over 63%.

The experiments of the play-out strategies have shown that the board configuration has a large influence on the game and the effectiveness and accuracy of

the play-out strategies. The random-move strategy appeared to be the most robust choice, doing reasonably well on all three boards. The wall-following strategy outperformed the other strategies only at board *b*, whereas the defensive strategy outperformed the other strategies at board *c*.

Applying play-out cut-off showed an increase in playing strength at boards *a* and *b* (54% and 56%, respectively), but significantly decreased the playing strength at board *c*. The bad performance at board *c* may have to do with the difficulty for a player to isolate itself on this board.

Similar to the play-out cut-off enhancement, the predictive expansion strategy showed a slight increase in playing strength at board *a* (53%) and *b* (58%), but not at board *c*. The poor win rate at board *c* is likely caused by the behaviour of MCTS for this board. The MCTS programs keep a large space open behind them, up until late in the game. In such positions, the space estimation heuristic used is not helpful.

The Score-Bounded MCTS-Solver was tested against MCTS, and turned out to be an improvement, although MCTS-Solver with PDE performs poorly at board *a*, in comparison to MCTS-PDE. MCTS-Solver is preferred for its ability to look up the shortest path leading to a win (or draw, when a draw is the best achievable outcome) once one or more moves have been proven. In contrast, MCTS-UCT usually postpones the victory.

Using PDE on MCTS-Solver enables the program to prove a position more quickly, however, the extra time spent on computing the heuristic did not work out well for all boards. PDE and PC in combination with MCTS-Solver further increased the overall playing strength of the program. MCTS-Solver-PDE-PC achieves a surprisingly high win rate on board *c* (70%), where MCTS-PC only scored (33%).

The experiment involving the $\alpha\beta$ program showed that the MCTS programs struggle at evaluating positions where the players are distant from one another (further than 10 steps away). Overall, the Solver-PDE-PC program is the best performing program, winning approximately 1 out of 5 games on average against the $\alpha\beta$ program, and achieving a win rate of 61% against MCTS-UCT.

The experiments show that the board configuration has a large influence on the playing strength of the enhancements tested. Since our goal was to create a Tron program capable of playing on any 13×13 map, the experiments should be conducted for many more boards.

As future research, applying a more sophisticated play-out strategy may increase the playing strength of MCTS in Tron. It would be interesting to see whether the play-out strategy and selection strategy can be improved such that MCTS can correctly look far ahead. The play-out phase might even have to be replaced completely by a sophisticated evaluation function (e.g. tree of chambers), as used by the $\alpha\beta$ program A1K0N.

References

1. B. Arneson, R.B. Hayward, and P. Henderson. Monte Carlo Tree Search in Hex. *IEEE Transactions on Computational Intelligence and AI in Games*, 2(4):251–257,

- 2010.
2. P. Auer, N. Cesa-Bianchi, and P. Fischer. Finite-time Analysis of the Multiarmed Bandit Problem. *Machine learning*, 47(2):235–256, 2002.
 3. H.J. Berliner. The B* Tree Search Algorithm: A Best-first Proof Procedure. *Artificial Intelligence*, 12(1):23–40, 1979.
 4. T. Cazenave and A. Saffidine. Score Bounded Monte-Carlo Tree Search. In *Computers and Games (CG 2010)*, volume 6515 of *Lecture Notes in Computer Science (LNCS)*, pages 93–104, Berlin Heidelberg, Germany, 2011. Springer.
 5. G.M.J-B. Chaslot, M.H.M. Winands, H.J. van den Herik, J.W.H.M. Uiterwijk, and B. Bouzy. Progressive Strategies for Monte Carlo Tree Search. *New Mathematics and Natural Computation*, 4(3):343–357, 2008.
 6. R. Coulom. Efficient Selectivity and Backup Operators in Monte-Carlo Tree Search. In *Computers and Games (CG 2006)*, volume 4630 of *LNCS*, pages 72–83. Springer-Verlag, 2007.
 7. Dmj. Survival Mode. <http://www.ai-contest.com/forums/viewtopic.php?p=1568>, 2010.
 8. S. Gelly, Y. Wang, R. Munos, and O. Teytaud. Modification of UCT with Patterns in Monte-Carlo Go. Technical Report 6062, INRIA, 2006.
 9. D.E. Knuth and R.W. Moore. An Analysis of Alpha-beta Pruning. *Artificial intelligence*, 6(4):293–326, 1975.
 10. L. Kocsis and C. Szepesvári. Bandit Based Monte-Carlo Planning. In *Machine Learning: ECML 2006*, volume 4212 of *LNCS*, pages 282–293. Springer, 2006.
 11. R.J. Lorentz. Amazons Discover Monte-Carlo. In *Computers and Games (CG 2008)*, volume 5131 of *LNCS*, pages 13–24. Springer, 2008.
 12. S. Samothrakis, D. Robles, and S.M. Lucas. A UCT Agent for Tron: Initial Investigations. In *2010 IEEE Conference on Computational Intelligence and Games (CIG)*, pages 365–371. IEEE, 2010.
 13. Sloane, A. Google AI Challenge Post-mortem. <http://alk0n.net/blah/archives/2010/03/>, Mar 2010.
 14. N.R. Sturtevant. An analysis of uct in multi-player games. In *Computers and Games (CG 2008)*, volume 5131 of *LNCS*, pages 37–49. Springer, 2008.
 15. R.S. Sutton and A.G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, Cambridge, MA, USA, 1998.
 16. N.G.P. Den Teuling. Monte-Carlo Tree Search for the Simultaneous Move Game Tron, 2011. B.Sc. thesis, Maastricht University.
 17. Y. Tsuruoka, D. Yokoyama, and T. Chikayama. Game-tree search algorithm based on realization probability. *ICGA Journal*, 25(3):132–144, 2002.
 18. University of Waterloo Computer Science Club. Google AI Challenge. <http://csclub.uwaterloo.ca/contest/>, Feb 2010.
 19. M.H.M. Winands, Y. Björnsson, and J.-T. Saito. Monte Carlo Tree Search in Lines of Action. *IEEE Transactions on Computational Intelligence and AI in Games*, 2(4):239–250, 2010.