



INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

*Modification of UCT with Patterns in Monte-Carlo
Go*

Sylvain Gelly — Yizao Wang — Rémi Munos — Olivier Teytaud

N° 6062

November 2006

Thème COG



*R*apport
de recherche



Modification of UCT with Patterns in Monte-Carlo Go

Sylvain Gelly*, Yizao Wang[†], Rémi Munos[‡], Olivier Teytaud[§]

Thème COG — Systèmes cognitifs
Projet TAO

Rapport de recherche n° 6062 — November 2006 — 21 pages

Abstract: Algorithm UCB1 for multi-armed bandit problem has already been extended to Algorithm UCT (Upper bound Confidence for Tree) which works for minimax tree search. We have developed a Monte-Carlo Go program, MoGo, which is the first computer Go program using UCT. We explain our modification of UCT for Go application and also the intelligent random simulation with patterns which has improved significantly the performance of MoGo. UCT combined with pruning techniques for large Go board is discussed, as well as parallelization of UCT. MoGo is now a top level Go program on 9×9 and 13×13 Go boards.

Key-words: Computer Go, Exploration-exploitation, UCT, Monte-Carlo, Patterns

* projet TAO, INRIA-Futurs, LRI, Batiment 490, Université Paris-Sud 91405 ORSAY CEDEX, France

[†] Centre de Mathématiques Appliquées, École Polytechnique, 91128 PALAISEAU CEDEX, France

[‡] Centre de Mathématiques Appliquées, École Polytechnique, 91128 PALAISEAU CEDEX, France

[§] projet TAO, INRIA-Futurs, LRI, Batiment 490, Université Paris-Sud 91405 ORSAY CEDEX, France

Modification d'UCT avec motifs dans le Monte-Carlo Go

Résumé : L'algorithme UCB1 pour le problème du bandit-manchot a récemment été étendu en l'algorithme UCT (Upper bound Confidence for Tree) pour la recherche arborescente min-max. Nous avons développé un joueur artificiel de Go, MoGo, basé sur des simulations Monte-Carlo, qui est le premier programme de Go utilisant UCT. Nous exposons notre modification de l'algorithme UCT pour l'application au jeu de Go, ainsi que l'utilisation de motifs dans les simulations aléatoires qui ont permis d'augmenter significativement le niveau de MoGo. Nous introduisons d'autre part des techniques d'élagage dans l'algorithme UCT pour les grands Goban, ainsi que de la parallélisation d'UCT. MoGo est maintenant un joueur artificiel de premier plan sur les Gobans de taille 9×9 et 13×13 .

Mots-clés : Joueurs artificiels de Go, Exploration-exploitation, UCT, Monte-Carlo, Motifs

1 Introduction

The history of Go stretches back some 4000 years and the game still enjoys a great popularity all over the world. Although its rules are simple (see <http://www.gobase.org> for a comprehensive introduction), its complexity has defeated the many attempts done to build a good Computer-Go player since the late 70's [4]. Presently, the best Computer-Go players are at the level of weak amateurs; Go is now considered one of the most difficult challenges for AI, replacing Chess in this role.

Go differs from Chess in many respects. First of all, the size and branching factor of the tree are significantly larger. Typically the Go board ranges from 9×9 to 19×19 (against 8×8 for the Chess board); the number of potential moves is a few hundred against a few dozen for Chess. Secondly, no efficient evaluation function approximating the minimax value of a position is available. For these reasons, the powerful alpha-beta search used by Computer-Chess players (see [14]) failed to provide good enough Go strategies.

Recent progress has been done regarding the evaluation of Go positions, based on Monte-Carlo approaches [6] (more on this in section 2). However, this evaluation procedure has a limited precision; playing the move with highest score in each position does not end up in winning the game. Rather, it allows one to restrict the number of relevant candidate moves in each step. Still, the size of the (discrete) search space makes it hardly tractable to use some standard Reinforcement Learning approach [16], to enforce the exploration *versus* exploitation (EvE) search strategy required for a good Go player.

Another EvE setting originated from Game Theory, the multi-armed bandit problem, is thus considered in this paper. The multi-armed bandit problem models the gambler, choosing the next machine to play based on her past selections and rewards, in order to maximize the total reward [2]. The UCB1 algorithm proposed by Auer et al. in the multi-armed bandit framework [1] was recently extended to tree-structured search space by Kocsis et al. (algorithm UCT) [12].

The main contributions of the player we present (named MoGo) are: (i) modification of UCT algorithm for Go, (ii) original use of sequence-like simulations in Monte-Carlo evaluation function. Several algorithmic (dynamic tree structure [9], parallelized implementation) or heuristic (simple pruning heuristics) issues were also tackled. MoGo has reached a comparatively good Go level: MoGo has been ranked as the first Go program out of 142 on 9×9 Computer Go Server (CGOS¹) since August 2006; and it won all the tournaments (9×9 and 13×13) on the international Kiseido Go Server² on October and November 2006.

This paper is organized as follows. Section 2 briefly introduces related work, assuming the reader's familiarity with the basics of Go. Section 3 describes MoGo, focussing on our contributions: the implementation of UCT in large sized search spaces, and the use of prior, pattern-based, knowledge to bias the Monte-Carlo evaluation. Experiment results are reported and discussed in Section 4. The paper concludes with some knowledge- and computer-intensive perspectives for improving MoGo.

¹<http://cgos.boardspace.net/>

²<http://www.weddslist.com/kgs/past/index.html>

2 Previous Related Work

Our approach is based on the Monte-Carlo Go and multi-armed bandit problems, which we present respectively in Section 2.1 and 2.2. UCT, which applies multi-armed bandit techniques to minimax tree search, is presented in Section 2.3. We suppose minimax tree and alpha-beta search is well known for the reader.

2.1 Monte-Carlo Go

Monte-Carlo Go, first appeared in 1993 [6], has attracted more and more attention in the last years. Monte-Carlo Go has been surprisingly efficient, especially on 9×9 game; CrazyStone, developed by Rémi Coulom [9], a program using stochastic simulations with very little knowledge of Go, is the best known³.

Two principle methods in Monte-Carlo Go are also used in our program. First we evaluate Go board situations by simulating random games until the end of game, where the score could be calculated easily and precisely. Second we combine the Monte-Carlo evaluation with minimax tree search. We use the tree structure of CrazyStone [9] in our program.

Remark 1 *We speak of a tree, in fact what we have is often an oriented graph. However, the terminology "tree" is widely used. As to the Graph History Interaction Problem (GHI) explained in [11], we ignore this problem considering it not very serious, especially compared to other difficulties in Computer-Go.*

2.2 Bandit Problem

A K -armed bandit, is a simple machine learning problem based on an analogy with a traditional slot machine (one-armed bandit) but with more than one arm. When played, each arm provides a reward drawn from a distribution associated to that specific arm. The objective of the gambler is to maximize the collected reward sum through iterative plays⁴. It is classically assumed that the gambler has no initial knowledge about the arms, but through repeated trials, he can focus on the most rewarding arms.

The questions that arise in bandit problems are related to the problem of balancing reward maximization based on the knowledge already acquired and attempting new actions to further increase knowledge, which is known as the exploitation-exploration dilemma in reinforcement learning. Precisely, exploitation in bandit problems refers to select the current best arm according to the collected knowledge, while exploration refers to select the sub-optimal arms in order to gain more knowledge about them.

A K -armed bandit problem is defined by random variables $X_{i,n}$ for $1 \leq i \leq K$ and $n \geq 1$, where each i is the index of a gambling machine (i.e., the "arm" of a bandit). Successive plays

³CrazyStone won the gold medal for the 9×9 Go game during the 11th Computer Olympiad at Turin 2006, beating several strong programs including GnuGo, Aya and GoIntellect.

⁴We will use "play an arm" when referring to general multi-armed problems, and "play a move" when referring to Go. In Go application, the "play" will not refer to a complete game but only one move.

of machine i yield rewards $X_{i,1}, X_{i,2}, \dots$ which are independent and identically distributed according to a certain but unknown law with unknown expectation μ_i . Here independence holds also for rewards across machines; i.e., $X_{i,s}$ and $X_{j,t}$ are independent (probably not identically distributed) for each $1 \leq i < j \leq K$ and each $s, t \geq 1$. Algorithms choose the next machine to play depending on the obtained results of the previous plays. Let $T_i(n)$ be the number of times machine i has been played after the first n plays. Since the algorithm does not always make the best choice, its expected loss is studied. Then the regret after n plays is defined by

$$\mu^* n - \sum_{j=1}^K \mu_j E[T_j(n)] \quad \text{where} \quad \mu^* = \max_{1 \leq i \leq K} \mu_i$$

$E[\]$ denotes expectation. In the work of Auer and Al. [1], a simple algorithm UCB1 is given, which ensures the optimal machine is played exponentially more often than any other machine uniformly when the rewards are in $[0, 1]$. Note

$$\bar{X}_{i,s} = \frac{1}{s} \sum_{j=1}^s X_{i,j} \quad , \quad \bar{X}_i = \bar{X}_{i,T_i(n)} ,$$

then we have:

Algorithm 1 *Deterministic policy: UCB1*

- *Initialization: Play each machine once.*
- *Loop: Play machine j that maximizes $\bar{X}_j + \sqrt{\frac{2 \log n}{T_j(n)}}$, where n is the overall number of plays done so far.*

One formula with better experimental results is suggested in [1]. Let

$$V_j(s) = \left(\frac{1}{s} \sum_{\gamma=1}^s X_{j,\gamma}^2 \right) - \bar{X}_{j,s}^2 + \sqrt{\frac{2 \log n}{s}}$$

be an estimated upper bound on the variance of machine j , then we have a new value to maximize:

$$\bar{X}_j + \sqrt{\frac{\log n}{T_j(n)} \min\{1/4, V_j(T_j(n))\}}. \quad (1)$$

According to Auer and Al., the policy maximizing (1) named UCB1-TUNED, considering also the variance of the empirical value of each arms, performs substantially better than UCB1 in all his experiments. This corresponds to our early results and then we use always the policy UCB1-TUNED in our program⁵.

⁵We will however say UCB1 for short.

2.3 UCT: UCB1 for Tree Search

UCT [12] is the extension of UCB1 [1] to minimax tree search. The idea is to consider each node as an independent bandit, with its child-nodes as independent arms. Instead of dealing with each node once iteratively, it plays sequences of bandits within limited time, each beginning from the root and ending at one leaf.

The algorithm UCT is defined in Table 1⁶. The program continues playing one sequence each time, which is defined from line 1 to line 8. Line 9 to line 21 are the function using UCB1 for choosing one arm (one child-node in the UCT case). Line 15 ensures each arm be selected once before further exploration. Line 16 applies the formula of UCB1. After each sequence, the value of played arm of each bandit is updated⁷ iteratively from the father-node of the leaf to the root by formula UCB1, described in function *updateValue* from line 22 to line 29. Here the code deals with the minimax case. In general, the value of each node converges to the real max (min) value as the number of simulations increases.

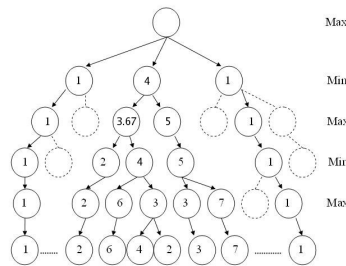


Figure 1: UCT search. The shape of the tree enlarges asymmetrically. Only updated values ($node[i].value$) are shown for each visited nodes.

In the problems of minimax tree search, what we are looking for is often the optimal branch at the root node. It is sometimes acceptable if one branch with a score near to the optimal one is found, especially when the depth of the tree is very large and the branching factor is big, like in Go, as it is often too difficult to find the optimal branch within short time.

In this sense, UCT outperforms alpha-beta search. Indeed we can outlight three major advantages. First, it works in an anytime manner. We can stop at any moment the algorithm, and its performance can be somehow good. This is not the case of alpha-beta search. Figure 2 shows if we stop alpha-beta algorithm prematurely, some moves at first level has even not been explored. So the chosen move can be far from optimal. Of course iterative deepening can be used, and solve partially this problem. Still, the anytime property is stronger for UCT and it is easier to finely control time in UCT algorithm.

⁶Giving the pseudocode in order to have a clear explanation, we do not talk about the optimization here.

⁷Here we use the original formula in Algorithm 1.


```

1: function playOneSequence(rootNode);
2:   node[0] := rootNode; i = 0;
3:   while(node[i] is not leaf) do
4:     node[i+1] := descendByUCB1(node[i]);
5:     i := i + 1;
6:   end while ;
7:   updateValue(node, -node[i].value);
8: end function;

9: function descendByUCB1(node)
10:  nb := 0;
11:  for i := 0 to node.childNode.size() - 1 do
12:    nb := nb + node.childNode[i].nb;
13:  end for;
14:  for i := 0 to node.childNode.size() - 1 do
15:    if node.childNode[i].nb = 0
16:      do v[i] := ∞;
17:    else v[i] := -node.childNode[i].value
18:      /node.childNode[i].nb
19:      +sqrt(2*log(nb))/(node.childNode[i].nb)
20:    end if;
21:  end for;
22:  index := argmax(v[j]);
23:  return node.childNode[index];
24: end function;

25: function updateValue(node,value)
26:  for i := node.size()-2 to 0 do
27:    node[i].value := node[i].value + value;
28:    node[i].nb := node[i].nb + 1;
29:    value := -value;
30:  end for;
31: end function;

```

Table 1: Pseudocode of UCT for minimax tree.

Second, UCT is robust as it automatically handles uncertainty in a smooth way. At each node, the computed value is the mean of the value for each child weighted by the frequency of visits. Then the value is a smoothed estimation of max, as the frequency of visits depends on the difference between the estimated values and the confidence of this estimates. Then, if one child-node has a much higher value than the others, and the estimate is good, this child-node will be explored much more often than the others, and then UCT selects most of the time the 'max' child node. However, if two child-nodes have a similar value, or a low confidence, then the value will be closer to an average.

Third, the tree grows in an asymmetric manner. It explores more deeply the good moves. What is more, this is achieved in an automatic manner. Figure 1 gives an example.

Figure 1 and Figure 2 compares clearly the explored tree of two algorithms within limited time. However, the theoretical analysis of UCT is in progress [13]. We just give some remarks on this aspect at the end of this section. It is obvious that the random variables involved in UCT are not identically distributed nor independent. This complicates the analysis of convergence. In fact we can define the bias for the arm i by:

$$\delta_{i,t} = \left| \mu_i^* - \frac{1}{t} \sum_{s=1}^t X_{i,s} \right|,$$

where μ_i^* is the minimax value of this arm. It is clear that at leaf level $\delta_{i,t} = 0$. We can also prove that

$$\delta_{i,t} \leq K^D \frac{\log t}{t},$$

with K constant and D the depth of the arm (counted from the root down). This corresponds to the fact that the bias is amplified when passing from deep level to the root, which prevents the algorithm from finding quickly the optimal arm at the root node.

An advantage of UCT is that it adapts automatically to the 'real' depth. For each branch of the root, its 'real' depth is the depth from where $\delta_{i,t} = 0$ holds true. For these branches, the bias at the root is bounded by $K^d \frac{\log t}{t}$ with the real depth $d < D$. The values of these branches converging faster than the other, UCT spends more time on other interesting branches.

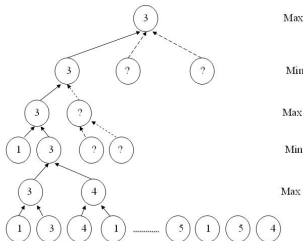


Figure 2: Alpha-beta search with limited time. The node with '??' is the node not explored yet. This happens often during the large-sized tree search where entire search is impossible. Iterative deepening solves partially this problem.

3 Main Work

In this section we present our program MoGo using UCT algorithm. Section 3.1 presents our application of UCT. Then, considering two important aspects for having a strong Monte-

```

1: function playOneSequenceInMoGo(rootNode)
2:   node[0] := rootNode; i := 0;
3:   do
4:     node[i+1] := descendByUCB1(node[i]); i := i + 1;
5:     while node[i] is not first visited;
6:     createNode(node[i]);
7:     node[i].value := getValueByMC(node[i]);
8:     updateValue(node,-node[i].value);
9:   end function;

```

Table 2: Pseudocode of UCT for MoGo

Carlo program: the quality of simulations (then the estimation of score) and the depth of the tree, we show in the two following sections our corresponding improvements. Section 3.2 presents the sequence-like random simulation with patterns. Section 3.3 presents ideas for tree search pruning on large Go board. Section 3.4 presents the modification on the exploring order of non-visited nodes. At last, Section 3.5 presents parallelization.

3.1 Application of UCT for Computer-Go

MoGo contains mainly two parts, namely the tree search part and the random simulation part, as shown in Figure 3. Each node of the tree represents a Go board situation, with child-nodes representing next situations after corresponding move.

The application of UCT for Computer-Go is based on the hypothesis that each Go board situation is a bandit problem, where each legal move is an arm with unknown reward but of a certain distribution. We suppose that there are only two kinds of arms, the winning ones and the losing ones. We set respectively reward 1 and 0. We ignore the case of draw, which is too rare in Go.

In the tree search part, we use a parsimonious version of UCT by introducing the same dynamic tree structure as in CrazyStone [9] in order to economize memory. The tree is then created incrementally by adding one node after each simulation as explained in the following. This is different from the one presented in [12], and is more efficient because less nodes are created during simulations. In other words, only nodes visited more than twice are saved, which economizes largely the memory and accelerates the simulations. The pseudocode is given in Table 2. Again we do not talk about optimization.

During each simulation, MoGo starts from the root of the tree that it saves in the memory. At each node, MoGo selects one move according to the UCB1 formula 1. MoGo then descends to the selected child node and selects a new move (still according to UCB1) until such a node has not yet been created in the tree. This part corresponds to the code from line 1 to line 5. The tree search part ends by creating this new node (in fact one leaf) in the tree. This is finished by *createNode*. Then MoGo calls the random simulation part,

the corresponding function *getValueByMC* at line 7, to give a score of the Go board at this leaf.

In the random simulation part, one random game is played from the corresponding Go board till the end, where score is calculated quickly and precisely according to the rules. The nodes visited during this random simulation are not saved. The random simulation done, the score received, MoGo updates the value at each node passed by the sequence of moves of this simulation⁸.

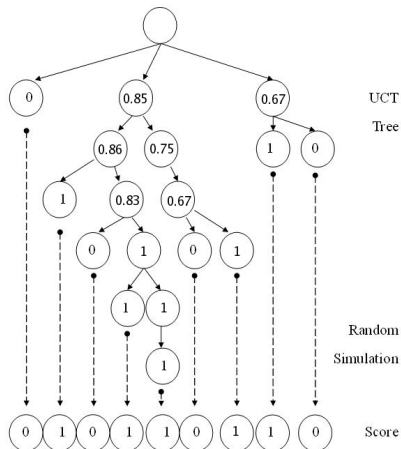


Figure 3: MoGo contains the tree search part using UCT and the random simulation part giving scores. The numbers on the bottom correspond to the final score of the game (win/loss). The numbers in the nodes are the updated values of the nodes ($node[i].value$)

Remark 2 *In the update of the score, we use the 0/1 score instead of the territory score, since the former is much more robust. Then the real minimax value of each node should be either 0 or 1. In practice, however, UCT approximates each node by a weighted average value in $[0, 1]$. This value is usually considered as the probability of winning.*

3.2 Improving simulation with domain knowledge

In this section we introduce our sequence-like random simulation with patterns. Its advantage is obvious compared with the classical random simulation, which we call pure random mode. In the following we talk about our improved random mode as well as our implementation of patterns.

⁸It is possible to arrive at one end game situation during the tree search part. In this case, one score could be calculated immediately and there is no need to create the node nor to call the random simulation part.

In the random simulation part, it is very important to have clever simulations giving credible scores. Using some simple 3×3 patterns inspired by Indigo [3] (similar patterns can also be found in [15]), our random simulation is likely to have more meaningful sequences in random simulations than before, which has improved significantly the level of MoGo.

The essential difference on the usage of patterns between the ours and other pattern-based Monte-Carlo programs ([3]) is that we use patterns to create meaningful sequences in simulations by finding local answers. The moves played are not necessarily globally better moves. It is not obvious that is more important to get better sequences rather than better moves to make the Monte-Carlo evaluation more accurate. However our experiments showed that the main improvement came from the use of local answers. If the same patterns are used to find interesting moves everywhere in the board instead of near the previous moves, the accuracy decreases. We believe that this claim is not obvious, and one of the main contribution of MoGo. We also don't use patterns for pruning in the tree. We have not investigated the more sophisticated patterns equipped by other programs like GnuGo.

In our pure random mode, legal moves are played on the Go board uniformly randomly, with few rules preventing the program from filling its own eyes. We also privilege the moves capturing some stones. On CGOS the bot named MoGo using exactly this mode has achieved rank score 1647 ELO⁹. Currently the rank of MoGo is close to 2200 ELO.

Then, since we were not satisfied by the pure random simulations which gave meaningless games most of the time, local patterns are introduced in order to have some more reasonable moves during random simulations. Our patterns are defined as 3×3 intersections, centered on one free intersection, where one move is supposed to be played. Our patterns consist of several functions, testing if one move in such a local situation (3×3) is interesting. More precisely, we test if one move satisfies some classical forms in Go games, for example cut move, Hane move, etc.

Moreover, we look for interesting moves only around the last played move on the Go board. This is because that local interesting moves look more likely to be the answer moves of the last moves, and thus local sequence appears when several local interesting moves are tested and then played continuously in random simulations.

We describe briefly how the improved random mode generates moves. It first verifies whether the last played move is an Atari; if yes, and if the stones under Atari can be saved (in the sense that it can be saved by capturing stones or increasing liberties), it chooses one saving move randomly; otherwise it looks for interesting moves in the 8 positions around the last played move and plays one randomly if there is any; otherwise it looks for the moves capturing stones on the Go board, plays one if there is any. At last, if still no move is found, it plays one move randomly on the Go board. Surely, the code of MoGo is actually complicated in details, with many small functions equipped with hand-coded Go knowledges. However, we believe the main frame given here is the most essential to have sequence-like simulations.

⁹The ELO (http://en.wikipedia.org/wiki/Elo_rating_system) is a rating system where the probability of winning is related to the difference between the ranks.

Figure 4 shows the first 30 moves of two random games using different modes. Moves generated by the improved random mode are obviously much more meaningful. We now give

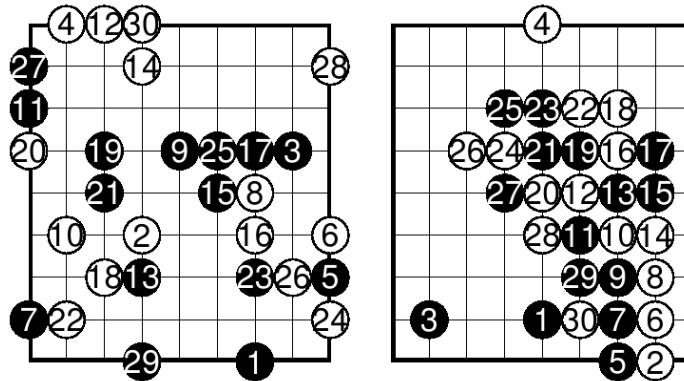


Figure 4: Left: beginning of one random game simulated by pure random mode. Moves are sporadically played with little sense. Right: beginning of one random game simulated by the pattern-based random mode. From move 5 to move 29 one complicated sequence is generated.

the detailed information on our patterns. The patterns are 3×3 intersections centered on an empty position, say p , where is supposed to play the next move. Each pattern is a boolean function, answering the question whether the next move playing on p is an interesting move. True is returned (when pattern is matched), if and only if the state of each position on the Go board is the same as the one on the corresponding position of the pattern, or there is a cross on the corresponding position (which means the situation of this position is ignored). Normally there is no constraints on the color of the next move (one move good for black is supposed to be also good for white). Some special cases are explained when mentioned. The symmetry, rotations and exchange of stone colors of patterns are considered. Moves are tested by patterns only if they are neither illegal moves nor self-Atari moves.

We have tried several patterns during the development of MoGo and implemented finally the ones shown in Figure 5, 6, 7 and 8, where the position with a square is where the next move is supposed to be played. We used hand-coded patterns in our implementation. However, it will be more interesting if this can be achieved by a learning system. Another approach using Bayesian generation can be found in Bouzy's work [5].

Remark 3 *We believe that it is not always better to have more 'good' patterns in the random modes, meanwhile what is more important is whether the random simulation can have some meaningful often. This claim needs more experiments.*

The Table 3 shows clearly how patterns improve the overall performance.

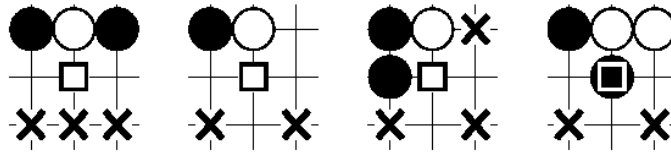


Figure 5: Patterns for Hane. True is returned if any pattern is matched. In the right one, a square on a black stone means true is returned if and only if the eight positions around are matched and it is black to play.

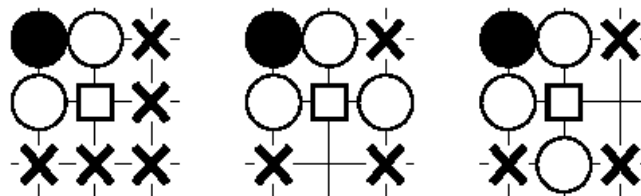


Figure 6: Patterns for Cut1. The Cut1 Move Pattern consists of three patterns. True is returned when the first pattern is matched and the next two are not matched.

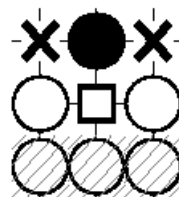


Figure 7: Pattern for Cut2. True is returned when the 6 upper positions are matched and the 3 bottom positions are not white.



Figure 8: Patterns for moves on the Go board side. True is returned if any pattern is matched. In the three right ones, a square on a black (resp. white) stone means true is returned if and only if the positions around are matched and it is black (resp. white) to play.

Random mode	Win. Rate for B. Games	Win. rate for W. Games	Total Win. Rate
Pure	46% (250)	36% (250)	41.2% \pm 4.4%
Sequence-like	77% (400)	82% (400)	80% \pm 2.8%

Table 3: Different modes with 70000 random simulations/move in 9x9.

3.3 UCT with pruning ideas

In this section we show our ideas (heuristics) to reduce the huge tree size, which makes MoGo relatively strong on large Go board. Thus we gain a larger local depth in the search tree by losing the global view of UCT. Obviously pruning heuristics may lead to a sub-optimal solution. First we define group by Go knowledge to reduce the branching factor in tree search. Then zone division is derived from group, which helps to have a more precise score. We use group and zone mode for 13×13 and 19×19 Go board. Figure 9 and 10 will give two convincing examples.

Remark 4 *As we are not very experienced for Go-knowledge-based programming and we had little time working on it, we believe other programs like GnuGo and AyaGo, or Monte-Carlo programs as [7][8] have more clever pruning techniques. Then, due to the space limitation, we do not give the detailed pseudo code of this part. However, our experimental results of combining UCT with pruning techniques are already encouraging.*

First we define one group as a set of strings and free intersections on a Go board according to certain Go knowledge, which gathers for example one big living group and its close enemies. We have implemented Common Fate Graph (CFG) [10] in our program to help the calculation of groups. The method starts from one string and recursively adds close empty intersections and strings close to these empty intersections until no more close strings are found within a distance controlled by a parameter.

In group mode, in the tree search part we search only the moves in the group instead of all over the Go board. In random simulation part there is no more such restriction. Using groups, we reduce the branching factor to less than 50 at the opening period. Then, depth of MoGo's tree could be around 7-8 on large Go board. Table 4 shows MoGo becomes competitive on large Go board by using group pruning technique. However, sophisticated pruning techniques are undoubtedly necessary to improve the level of Computer-Go programs.

Opponents	Win. Rate for B. Games	Win. rate for W. Games	Total Win. Rate
No group vs GnuGo level 0	53.2%(216)	51.8% (216)	52% \pm 4.8%
No group vs GnuGo level 8	24.2%(300)	30% (300)	27% \pm 3.6%
group vs GnuGo level 0	67.5% (80)	61.2% (80)	64.3% \pm 7.5%
group vs GnuGo level 8	51.9% (160)	60% (160)	56% \pm 5.5%

Table 4: MoGo with 70000 simulations per move, on 13×13 Go board, using or not the group mode heuristic.

Then we developed the zone mode. As explained above, group mode limits the selection of moves in the tree search part. It has however no restriction on the random simulation. In zone mode, the random moves is generated only in a certain zone instead of on the whole Go board. This is because on the large Go board, random simulations are too long (of an average length more than 400 for 19×19 Go board) to be able to return credible scores. The score obtained in this zone is then tuned for the global Go board.

The tuning idea is very crucial in our zone mode, otherwise the global evaluation of situation by random simulations, which is very imprecise most of the time for the large Go board, biases too much the local decision. More precisely, before UCT we estimate, by a certain number (5000 in our experiments) of Monte-Carlo simulations, the score \bar{s}_0 for the zone decided by the current situation (including last move of opponent), and the score \bar{s}_1 for the area outside the zone. Then we tune the s_1 by

$$\hat{s}_1 = \alpha(\bar{s}_1 + \bar{s}_0 - komi) - (\bar{s}_0 - komi) \quad (2)$$

and then after each simulation in the zone, with local score s_0 we return $s_0 + \hat{s}_1$ as the final score of the simulation. In this way, the score is always closer to a close game score than it should have been by using $s_0 + \bar{s}_1$. This helps MoGo play more normally locally.

3.4 Modification of exploring order for non-visited nodes

UCT works very well when the node is frequently visited as the trade-off between exploration and exploitation is well handled by UCB1 formula. However, for the nodes far from the root, whose number of simulations is very small, UCT tends to be too much exploratory. This is due to the fact that all the possible moves in one position are supposed to be explored before using the UCB1 formula. Thus, the values associated to moves in deep nodes are

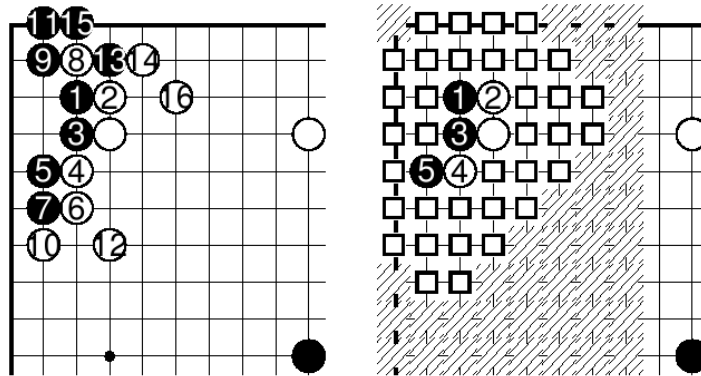


Figure 9: From one game played by MoGo (White) on KGS. In the left figure, it found almost the exact joseki response in the left-up corner using group and zone mode. In the right figure, the union of the two groups related to the last two moves (move 4 and move 5) are marked by squares (only on the free positions). The other area of zone (besides the stones and the group area) at this moment is marked by shadow.

not meaningful, since the child-nodes of these nodes are not all explored yet and, sometimes even worse, the visited ones are selected in fixed order. This can lead to bad predicted sequences.

Two modifications are made to have a better order.

First-play urgency

UCB1 algorithm begins by exploring each arm once, before using the formula (1). This can sometimes be inefficient especially if the number of trials is not large comparing to the number of arms. This is the case for numerous nodes in the tree (number of visits is small comparing to the number of moves). For example if an arm keeps returning 1 (win), there is no good reason to explore other arms. We have set a fixed constant named first-play urgency (FPU) in the algorithm. For each move, we name its urgency by the value of formula (1). The FPU is by default set to ∞ for each legal move before first visit (see line 15 in Table 1). Any node, after being visited at least once, has its urgency updated according to UCB1 formula. We play the move with the highest urgency. Thus, the FPU ∞ ensures the exploration of each move once before further exploitation of any previously visited move. On the other way, smaller FPU ensures earlier exploitations if the first simulations lead to an urgency smaller than FPU (in this case the other unvisited nodes are not selected). This improved the level of MoGo according to our experiment as shown in Table 7.

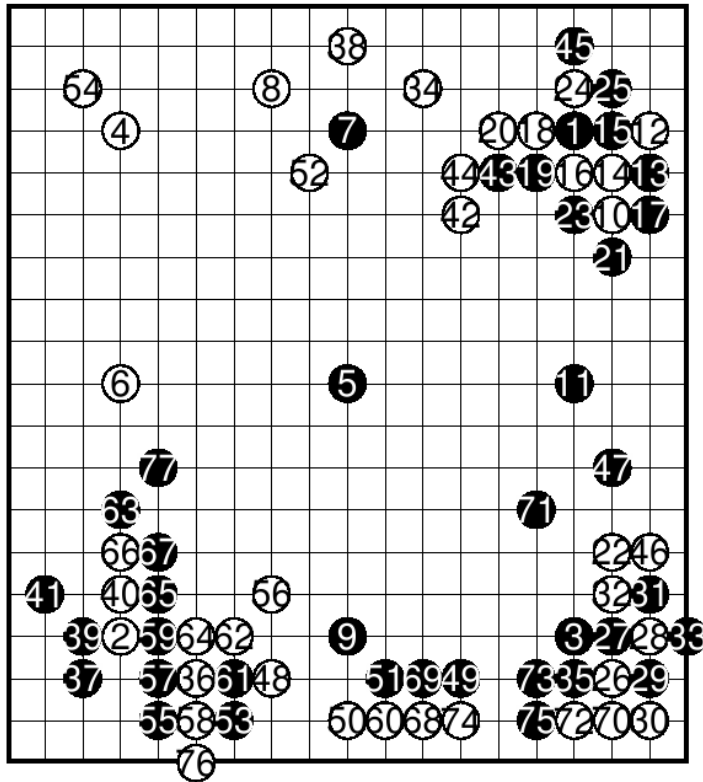


Figure 10: The opening of one game between MoGo and Indigo in the 18th KGS Computer Go Tournament. MoGo (Black) was in advantage at the beginning of the game, however it lost the game at the end.

Use information of the parents

One assumption that can be made in go game is that given a situation, good moves may sometimes still be good ones on the following move. When we encounter a new situation, instead of exploring each move m in any order, we can use the value estimation of m in an earlier position to choose a better order. We typically use the estimated values of the grandfather of the node. We believe this helps MoGo on the large Go board, however we do not have enough experiments to claim significant results.

3.5 Parallelization

As UCT scales well with time, we made MoGo run on a multi-processors machine with shared memory. The modifications to the algorithm are quite straightforward. All the processors share the same tree, and the access to the tree is locked by mutexes. As UCT is deterministic, all the threads could take exactly the same path in the tree, except for the leaf. The behavior of the multithreaded UCT as presented here is then different from the monothreaded UCT. Two experiments has then to be done. First, we can compare the level of MoGo using the monothreaded or the multithreaded algorithms while allowing the same number of simulations per move. All such experiments showed non significant differents in the play level¹⁰. Second, we can compare the level using the same time per move (the multithreaded version will then make more simulations per move). As UCT benefits from the computational power increase, the multithreaded UCT is efficient (+100 ELO on CGOS with 4 processors).

4 Results

We list in this section several experiment results who reflect characteristics of the algorithm. All the tests are made by letting MoGo play against GnuGo 3.6 with default mode. Komi are set to 7.5 points. In the tables, are given, the winning rate when MoGo plays black and white are given along the number of games played in each color (between parenthesis). The confidence interval is 95%.

4.1 Dependence of Time

The performance of our program depends on the given time (equally the number of simulations) for each move. Table 5 shows its level improves as this number increases. The outstanding performance of MoGo on double-processors and quadri-processors also supports this claim.

Seconds per move	Winning Rate for Black Games	Winning rate for White Games	Total Winning Rate
5	26% ± 12% (50)	26% ± 12% (50)	26% ± 8.7%
20	41% ± 6% (250)	42% ± 6% (250)	41.8% ± 4.4%
60	53% ± 7% (200)	50% ± 7% (200)	51.5% ± 5%

Table 5: Pure random mode with different times.

¹⁰we had only access to a 4 processors computer, the behavior can be very different with many more processors.

4.2 Parametrized UCT

We parametrize the UCT implemented in our program by two new parameters, namely p and FPU . First we add one coefficient p to formula UCB1-TUNED (1), which by default is 1. This leads to the following formula: choose j that maximizes:

$$\bar{X}_j + p \sqrt{\frac{\log n}{T_j(n)} \min\{1/4, V_j(n_j)\}}$$

p decides the balance between exploration and exploitation. To be precise, the smaller the p is, the deeper the tree is explored. According to our experiment shown in Table 6, UCB1-TUNED is almost optimal in this sense.

p	Winning Rate for Black Games	Winning rate for White Games	Total Winning Rate
0.05	2% ± 4% (50)	4% ± 5% (50)	3% ± 3.4%
0.55	30% ± 13% (50)	36% ± 13% (50)	33% ± 9.4%
0.80	33% ± 9% (100)	39% ± 10% (50)	36% ± 6.7%
1.0	40% ± 8% (150)	38% ± 8% (150)	39% ± 5.6%
1.1	39% ± 8% (150)	41% ± 8% (150)	40% ± 5.6%
1.2	40% ± 8% (150)	44% ± 8% (150)	42% ± 5.7%
1.5	30% ± 13% (50)	26% ± 12% (50)	28% ± 9%
3.0	36% ± 13% (50)	24% ± 12% (50)	30% ± 9%
6.0	22% ± 11% (50)	18% ± 10% (50)	20% ± 8%

Table 6: Coefficient p decides the balance between exploration and exploitation. (Pure random Mode)

The second is the first-play urgency (FPU) as explained in 3.4. Some results are shown in Table 7. We believe that changing exploring order of non-visited nodes can bring further improvement. We finally use 1.1 as the initiate FPU for MoGo on CGOS.

4.3 Results On CGOS

MoGo is ranked as the first program on 9×9 Computer Go Server since August.

5 Conclusion

The success of MoGo shows the efficiency of UCT compared to alpha-beta search in the sense that nodes are automatically studied with better order, especially in the case when search time is too limited. We have discussed the advantages of UCT relevant to Computer-Go. It is worthy to mention that a growing number of top level Go programs now use UCT.

FPU	Winning Rate for Black Games	Winning rate for White Games	Total Winning Rate
1.4	37% \pm 9% (100)	38% \pm 10% (100)	37.5% \pm 7%
1.2	46% \pm 10% (100)	36% \pm 10% (100)	41% \pm 7%
1.1	45% \pm 6% (250)	41% \pm 6% (250)	43.4% \pm 4.4%
1.0	49% \pm 6% (300)	42% \pm 6% (300)	45% \pm 4%
0.9	47% \pm 8% (150)	32% \pm 8% (150)	40% \pm 5.6%
0.8	40% \pm 14% (50)	32% \pm 13% (50)	36% \pm 9.6%

Table 7: Influence of FPU (70000 simulations/move).

We have discussed improvements that could be made to UCT algorithm. In particular, UCT does not help to choose a good ordering for non-visited moves, nor is it so effective for little explored moves. We proposed some methods adjusting the first-play urgency to solve this problem, and further improvements are expected in this direction.

We have proposed the pattern-based (sequence-like) simulation which has improved significantly the level of MoGo. We implemented 3×3 patterns in random simulations in order to have more meaningful sequences. We believe that significant improvements can still be made in this direction, for example by using larger patterns, perhaps automatically generated ones. It is also possible to implement some sequence-forced patterns to improve the quality of simulations.

We have also shown the possibilities of combining UCT with pruning techniques in order to have a strong program on large Go board. Having had some encouraging results, we believe firmly further improvements in this direction.

A straightforward parallelization of UCT on shared-memory computer is made and has given some positive results. Parallelization on a cluster of computers can be interesting but the way to achieve that is yet to be found.

Acknowledgments

We would like to thank Pierre-Arnaud Coquelin for the help during the development of MoGo. We specially thank Rémi Coulom for sharing his experiences of programming Crazy-Stone. We also appreciate the discussions on the Computer-Go mailing list.

References

- [1] P. Auer, N. Cesa-Bianchi, and P. Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine Learning*, 47(2/3):235–256, 2002.

-
- [2] P. Auer, N. Cesa-Bianchi, Y. Freund, and R. E. Schapire. Gambling in a rigged casino: the adversarial multi-armed bandit problem. In *Proceedings of the 36th Annual Symposium on Foundations of Computer Science*, pages 322–331. IEEE Computer Society Press, Los Alamitos, CA, 1995.
 - [3] B. Bouzy. Associating domain-dependent knowledge and monte carlo approaches within a go program. *Information Sciences, Heuristic Search and Computer Game Playing IV, Edited by K. Chen*, (4):247–257, 2005.
 - [4] B. Bouzy and T. Cazenave. Computer go: An AI oriented survey. *Artificial Intelligence*, 132(1):39–103, 2001.
 - [5] B. Bouzy and G. Chaslot. Bayesian generation and integration of k-nearest-neighbor patterns for 19x19 go. In *G. Kendall and Simon Lucas, editors, IEEE 2005 Symposium on Computational Intelligence in Games, Colchester, UK*, pages 176–181, 2005.
 - [6] B. Bruegmann. Monte carlo go. 1993.
 - [7] T. Cazenave. Abstract proof search. *Computers and Games, Hamamatsu, 2000*, pages 39–54, 2000.
 - [8] T. Cazenave and B. Helmstetter. Combining tactical search and monte-carlo in the game of go. *IEEE CIG 2005*, pages 171–175, 2005.
 - [9] R. Coulom. Efficient selectivity and backup operators in monte-carlo tree search. In *P. Ciancarini and H. J. van den Herik, editors, Proceedings of the 5th International Conference on Computers and Games, Turin, Italy, 2006*, To appear.
 - [10] T. Graepel, M. Goutr  , M. Kr  ger, and R. Herbrich. Learning on graphs in the game of go. *Lecture Notes in Computer Science*, 2130:347–352, 2001.
 - [11] A. Kishimoto and M. M  ller. A general solution to the graph history interaction problem. *Nineteenth National Conference on Artificial Intelligence (AAAI 2004), San jose, CA*, pages 644–649, 2004.
 - [12] L. Kocsis and C. Szepesvari. Bandit based monte-carlo planning. In *15th European Conference on Machine Learning (ECML)*, pages 282–293, 2006.
 - [13] L. Kocsis, C. Szepesv  ri, and J. Willemsen. Improved monte-carlo search. working paper, 2006.
 - [14] M. Newborn. *Computer Chess Comes of Age*. Springer-Verlag, 1996.
 - [15] L. Ralaivola, L. Wu, and P. Baldi. Svm and pattern-enriched common fate graphs for the game of go. *ESANN 2005*, pages 485–490, 2005.
 - [16] R. Sutton and A. Barto. *Reinforcement learning: An introduction*. MIT Press., Cambridge, MA, 1998.

Unité de recherche INRIA Futurs
Parc Club Orsay Université - ZAC des Vignes
4, rue Jacques Monod - 91893 ORSAY Cedex (France)

Unité de recherche INRIA Lorraine : LORIA, Technopôle de Nancy-Brabois - Campus scientifique
615, rue du Jardin Botanique - BP 101 - 54602 Villers-lès-Nancy Cedex (France)

Unité de recherche INRIA Rennes : IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex (France)

Unité de recherche INRIA Rhône-Alpes : 655, avenue de l'Europe - 38334 Montbonnot Saint-Ismier (France)

Unité de recherche INRIA Rocquencourt : Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex (France)

Unité de recherche INRIA Sophia Antipolis : 2004, route des Lucioles - BP 93 - 06902 Sophia Antipolis Cedex (France)

Éditeur
INRIA - Domaine de Voluceau - Rocquencourt, BP 105 - 78153 Le Chesnay Cedex (France)
<http://www.inria.fr>
ISSN 0249-6399