Bandit-Based Genetic Programming

J.-B. Hoock and O. Teytaud

TAO (Inria), LRI, UMR 8623(CNRS - Univ. Paris-Sud),
bat 490 Univ. Paris-Sud 91405 Orsay, France, hoock@lri.fr

Abstract. We consider the validation of randomly generated patterns in a Monte-Carlo Tree Search program. Our bandit-based genetic programming (BGP) algorithm, with proved mathematical properties, outperformed a highly optimized handcrafted module of a well-known computer-Go program with several world records in the game of Go.

1 Introduction

Genetic Programming (GP) is the automatic building of programs for solving a given task. In this paper, we investigate a bandit-based approach for selecting fruitful modifications in genetic programming, and we apply the result to our program MoGo.

When testing a large number of modifications in a stochastic algorithm with limited resources in an uncertain framework, there are two issues:

- which modifications are to be tested now ?
- when we have no more resources (typically no more time), we must decide which modifications are accepted.

The second issue is often addressed through statistical tests. However, when many modifications are tested, it is a problem of multiple simultaneous hypothesis testing: this is far from being straightforward; historically, this was poorly handled in many old applications. Cournot stated that if we consider a significance threshold of 1% for differences between two sub-populations of a population, then, if we handcraft plenty of splittings in two sub-populations, we will after a finite time find a significant difference, whenever the two populations are similar. This was not for genetic programming, but the same thing holds in GP: if we consider 100 random mutations of a program, all of them being worst than the original program, and if we have a 1% risk threshold in the statistical validation of each of them, then with probability $(1 - 1/100)^{100} \simeq 37\%$ we can have a positive validation of at least one harmful mutation. Cournot concluded, in the 19th century, that this effect was beyond mathematical analysis; nonetheless this effect is clearly understood today, with the theory of multiple hypothesis testing - papers cited below clearly show that mathematics can address this problem.

The first issue is also non trivial, but a wide literature has been devoted to it: so-called bandit algorithms. This is in particular efficient when no prior information on the modifications is available, and we can only evaluate the quality of a modification through statistical results.

Usually the principles of a Bernstein race are as follows:

- decide a risk threshold δ_0 ;
- then, modify the parameters of all statistical tests so that all confidence intervals are *simultaneously* true with probability $\geq 1 - \delta_0$;
- then, as long as you have computational resources, apply a *bandit* algorithm for choosing which modification to test, depending on statistics; typically, a bandit algorithm will choose to spend computational resources on the modification which has the best statistical upper bound on its average efficiency;
- at the end, select the modifications which are significant.

A main reference, with theoretical justifications, is [22]. A main difference here is that we will not assume that all modifications are cumulative: here, whenever two modifications A and B are statistically good, we can't select both modifications - maybe, the baseline + A + B will be worse than the baseline, whenever both baseline+A and baseline+B are better than the baseline.

In section 2, we present non-asymptotic confidence bounds. In section 3 we present racing algorithms. Then, section 4 presents our algorithm and its theoretical analysis. Section 5 is devoted to experiments.

2 Non-asymptotic confidence bounds

In all the paper, we consider fitness values between 0 and 1 for simplifying the writing. The most classical bound is Hoeffding's bound. Hoeffding's bound states that with probability at least $1 - \delta$, the empirical average $\hat{r}$ verifies $|\hat{r} - \mathbb{E}r| \leq deviation_{\text{Hoeffding}}(\delta, n)$ where n is the number of simulations and where

$$deviation_{\text{Hoeffding}}(\delta, n) = \sqrt{\log(2/\delta)/n}. \quad (1)$$

[1, 22] has shown the efficiency of using Bernstein's bound instead of Hoeffding's bound, in some settings. The bound is then $deviation_{\text{Bernstein}} = \hat{\sigma} \sqrt{2 \log(3/\delta)/n} + 3 \log(3/\delta)/n$, where $\hat{\sigma}$ is the empirical standard deviation. Bernstein's version will not be used in our experiments, because the variance is not small in our case; nonetheless, all theoretical results also hold with Bernstein's variant.

3 Racing algorithms

Racing algorithms are typically (and roughly, we'll be more formal below) as follows:

```

Let  $S$  be equal to  $S_0$ , some given set of admissible modifications.
while  $S \neq \emptyset$  do
  Select  $s = select() \in S$  with some algorithm
  Perform one Monte-Carlo evaluation of  $s$ .
  if  $s$  is statistically worse than the baseline then
     $S \leftarrow S \setminus \{s\}$  //  $s$  is discarded
  else if  $s$  is statistically better than the baseline then
```

```

    Accept  $s$ ;  $S \leftarrow S \setminus \{s\}$   $s$  is accepted
  end if
end while

```

With relevant statistical tests, we can ensure that this algorithm will select all “good” modifications (to be formalized later), reject all bad modifications, and stop after a finite time if all modifications have a non-zero effect. We refer to [22] for more general informations on this, or [18, 15] for the GP case; we will here focus on the most relevant (relevant for our purpose) case. In genetic programming, it’s very clear that even if two modifications are, independently, good, the combination of these two modifications is not necessarily good. We will therefore provide a different algorithm in section 4 with a proof of consistency.

4 Theoretical analysis for genetic programming

We will assume here that for a modification s , we can define:

- $e(s)$, the (of course unknown) expected value of the reward when using modification s . This expected value is termed the efficiency of s . We will assume in the sequel that the baseline is 0.5 - an option is good if and only if it performs better than 0.5, and the efficiency is the average result on experiments.
- $n(s)$, the number of simulations of s already performed.
- $r(s)$ the total reward of s , i.e. the sum of the rewards of the $n(s)$ simulations with modification s .
- $ub(s)$, an upper bound on the efficiency of s , to be computed depending on the previous trials ($ub(s)$ will be computed thanks to Bernstein bounds or Hoeffding bounds).
- $lb(s)$, a lower bound on the efficiency of s (idem).

The two following properties will be proved for some specific functions lb and ub ; the results around our BGP (bandit-based genetic programming) algorithm below hold whenever lb and ub verify these assumptions.

- **Consistency:** with probability at least $1 - \delta_0$, for all calls to ub and lb , the efficiency of s is between $lb(s)$ and $ub(s)$:

$$e(s) \in [lb(s), ub(s)]. \quad (2)$$

- **Termination:** when the number of simulations of s goes to infinity, then

$$ub(s) - lb(s) \rightarrow 0. \quad (3)$$

These properties are exactly what is ensured by Bernstein’s bounds or Hoeffding’s bounds. They will be proved for some variants of ub and lb defined below (Lemma 1, using Hoeffding’s bound); they will be assumed in results about the BGP algorithm below. Therefore, our results about BGP (Theorem 1) will hold for our variants of lb and ub . Our algorithm and proof do not need a specific function ub or lb , provided that these assumptions are verified. However, we precise below a classical form of ub and lb , in order to point out that there exists

such ub and lb ; moreover, they are easy to implement. lb and ub are computed by a function with a memory (*i.e.* with static variables):

Function *computeBounds*(s) (variant 1)

Static internal variable: $nbTest(s)$, initialized at 0.

Let n be the number of times s has been simulated.

Let r be the total reward over those s simulations.

$nbTest(s) = nbTest(s) + 1$

Let $lb(s) = r/n - deviation_{\text{Hoeffding}}\left(\delta_0/(\#S \times 2^{nbTest(s)}), n\right)$.

Let $ub(s) = r/n + deviation_{\text{Hoeffding}}\left(\delta_0/(\#S \times 2^{nbTest(s)}), n\right)$.

What is important in these formula is that the sum of the $\delta_0/(\#S \times 2^{nbTests(s)})$, for $s \in S$ and $nbTest(s) \in \{1, 2, 3 \dots\}$, is at most δ_0 . By union bound, this implies that the overall risk is at most δ_0 . The proof of the consistency and of the termination assumptions are therefore immediate consequences of Hoeffding's bounds (we could use Bernstein's bounds if we believed that small standard deviations matter). A (better) variant, based on $\sum_{n \geq 1} 1/n^2 = \pi^2/6$ is

Function *computeBounds*(s) (variant 2)

Static internal variable: $nbTest(s)$, initialized at 0.

Let n be the number of times s has been simulated.

Let r be the average reward over those s simulations.

$nbTest(s) = nbTest(s) + 1$

Let $lb(s) = r/n - deviation_{\text{Hoeffding}}\left(\delta_0/(\#S \times \left(\frac{\pi^2 nbTest(s)^2}{6}\right)), n\right)$.

Let $ub(s) = r/n + deviation_{\text{Hoeffding}}\left(\delta_0/(\#S \times \left(\frac{\pi^2 nbTest(s)^2}{6}\right)), n\right)$.

We show precisely the consistency of *computeBounds* below.

Lemma 1 (Consistency of *computeBounds*.) *For all S finite, for all algorithms calling *computeBounds* and simulating modifications in arbitrary order, with probability at least $1 - \delta_0$, for all s and after each simulation, $lb(s) \leq e(s) \leq ub(s)$.*

The proof is removed due to length constraints. $\square$

Our algorithm, BGP (Bandit-based Genetic Programming), based on the *computeBounds* function above, is as follows:

BGP algorithm.

$S = S_0$ = some initial set of modifications.

while $S \neq \emptyset$ **do**

Select $s \in S$ // the selection rule is not specified here
// (the result is independent of it)

Let n be the number of simulations of modification s .

Simulate s n more times (*i.e.* now s has been simulated $2n$ times).

//this ensures $nbTests(s) = O(\log(n(s)))$

computeBounds(s)

if $lb(s) > 0.501$ **then**

```

    Accept  $s$ ; exit the program.
  else if  $ub(s) < 0.504$  then
     $S \leftarrow S \setminus \{s\}$  //  $s$  is discarded.
  end if
end while

```

We do not specify the selection rule. The result below is independent of the particular rule.

Theorem 1 (Consistency of BGP). *When using variant 1 or variant 2 of computeBounds, or any other version ensuring consistency (Eq. 2) and termination (Eq. 3), BGP is consistent in the sense that:*

1. *if at least one modification s has efficiency $> .504$, then with probability at least $1 - \delta_0$ a modification with efficiency $> .501$ will be selected (and the algorithm terminates).*
2. *if no modification has efficiency $> .504$, then with probability at least $1 - \delta_0$ the algorithm will*
 - (a) *either select a modification with efficiency $> .501$ (and terminate);*
 - (b) *or select no modification and terminate.*

Remark: The constants 0.501 and 0.504 are arbitrary provided that the latter is greater or equal to the former. The proof is removed due to length constraints. $\square$

We have only considered $|S| < \infty$. The extension to $S = \{s_1, s_2, s_3, \dots\}$ countable is straightforward but removed due to length constraints.

5 Experiments

Life is a Game of Go in which rules have been made unnecessarily complex, according to an old proverb. As a matter of fact, Go has very simple rules, is very difficult for computers, is central in education in many Asian countries (part of school activities in some countries) and has NP-completeness properties for some families of situations [12], and PSPACE-hardness for others [21], and EXPTIME-completeness for some versions [23]. It has also been chosen as a testbed for artificial intelligence by many researchers. The main tools, for the game of Go, are currently MCTS/UCT (Monte-Carlo Tree Search, Upper Confidence Trees); these tools are also central in many difficult games and in high-dimensional planning. An example of nice Go game, won by MoGo as white in 2008 in the GPW Cup, is given in Fig. 1 (left). Since these approaches have been defined [7, 10, 17], several improvements have appeared like First-Play Urgency [25], Rave-values [5, 14] (see <ftp://ftp.cgl.ucsf.edu/pub/pett/go/ladder/mcgo.ps> for B. Bruegman’s unpublished paper), patterns and progressive widening [11, 8], better than UCB-like (Upper Confidence Bounds) exploration terms [20], large-scale parallelization [13, 9, 6, 16], automatic building of huge opening books [2]. Thanks to all these improvements, our implementation MoGo already won even games against a professional player in 9x9 (Amsterdam, 2007; Paris, 2008; Taiwan 2009), and recently won with handicap 6 against a professional player

Fig. 1. Left: A decisive move (number 28) played by MoGo as white, in the GPW Cup 2008. Right: An example from Senseis of good large pattern in spite of a very bad small pattern. The move 2 is a good move.

(Tainan, 2009), and with handicap 7 against a top professional player, Zhou Junxun, winner of the LG-Cup 2007 (Tainan, 2009). Besides impressive results for the game of Go, MCTS/UCT have been applied to non-linear optimization [4], optimal sailing [17], active learning [24]. The formula used in the bandit is incredibly complicated, and it is now very hard to improve the current best formula [20].

Here we will consider only mutations consisting in adding patterns in our program MoGo. Therefore, accepting a mutation is equivalent to accepting a pattern. We experiment random patterns for biasing UCT. The reader interested in the details of this is referred to [20]. Our patterns contain jokers, black stones, empty locations, white stones, locations out of the goban, and are used as masks over all the board: this means that for a given location, we consider patterns like “there is a black stone at coordinate $+2,+1$, a stone (of any color) at coordinate $+3,0$, and the location at coordinate $-1,-1$ is empty”. This is a very particular form of genetic programming. We consider here the automatic generation of patterns for biasing the simulations in 9x9 and 19x19 Go. Please note that: (1) When we speak of good or bad shapes here, it is in the sense of “shapes that should be more simulated by a UCT-like algorithm”, or “shapes that should be less simulated by a UCT-like algorithm”. This is not necessarily equivalent to “good” or “bad” shapes for human players (yet, there are correlations). (2) In 19x19 Go, MoGoCVS is based on tenths of thousands of patterns as in [8]. Therefore, we do not start from scratch. A possible goal would be to have similar results, with less patterns, so that the algorithm is faster (the big database of patterns provides good biases but it is very slow). (3) In 9x9 Go, there are no big library of shapes available; yet, human expertise has been encoded in MoGo, and we are far from starting from scratch. Engineers have spent hundreds of hours manually optimizing patterns. The goals are both (i) finding shapes that should be more simulated (ii) finding shapes that should be less simulated.

Section 5.1 presents our experiments for finding good shapes in 9x9 Go. Section 5.2 presents our experiments for finding bad shapes in 9x9 Go. Section

5.3 presents our unsuccessful experiments for finding both good and bad shapes in 19x19, from MoGoCVS and its database of patterns as in [8]. Section 5.4 presents results on MoGoCVS with patterns removed, in order to improve the version of MoGoCVS without the big database of pattern.

5.1 Finding good shapes for simulations in 9x9 Go

Here the baseline is MoGo CVS. All programs are run on one core, with 10 000 simulations per move. All experiments are performed on Grid5000. The selection rule, not specified in BGP, is the upper bound as in UCB[19, 3]: we simulate s such that $ub(s)$ is maximal. We here test modifications which give a positive bias to some patterns, *i.e.* we look for shapes that should be simulated more often.

For each iteration, we randomly generate some individuals, and test them with the BGP algorithm. For the three first iterations, 10 patterns were randomly generated; the two first times, one of these 10 patterns was validated; the third time, no pattern was validated. Therefore, we have three version of MoGo: MoGoCVS, MoGoCVS+P1, and MoGoCVS+P1+P2, where P1 is the pattern validated at the first iteration and P2 is the pattern validated at the second iteration. We then tested the relative efficiency of these MoGos as follows:

Tested code	Opponent	Success rate
MoGoCVS + P1	MoGoCVS	$50.78\% \pm 0.10\%$
MoGoCVS + P1 + P2	MoGoCVS + P1	$51.2\% \pm 0.20\%$
MoGoCVS + P1 + P2	MoGoCVS	$51.9\% \pm 0.16\%$

We also checked that this modification is also efficient for 100 000 simulations per move, with success rate $52.1 \pm 0.6\%$ for MoGoCVS+P1+P2 against MoGoCVS. There was no pattern validated during the third iteration, which was quite expensive (one week on a cluster). We therefore switched to another variant; we tested the case $|S_0| = 1$, *i.e.* we test one individual at a time. We launched 153 iterations with this new version. There were therefore 153 tested patterns, and none of them was validated.

5.2 Finding bad shapes for simulations in 9x9 Go

We now switched to the research of negative shapes, *i.e.* patterns with a negative influence of the probability, for a move, to be simulated. We kept $|S_0| = 1$, *i.e.* only one pattern tested at each iteration. There were 173 iterations, and two patterns P3 and P4 were validated. We verified the quality of these negative patterns as follows, with mogoCVS the version obtained in the section above:

Tested code	Opponent	Success rate
MoGoCVS + P1 + P2 + P3	MoGoCVS + P1 + P2	$50.9\% \pm 0.2\%$
MoGoCVS + P1 + P2 + P3	MoGoCVS	$52.6\% \pm 0.16\%$
MoGoCVS + P1 + P2 + P3 + P4	MoGoCVS + P1 + P2 + P3	$50.6\% \pm 0.13\%$
MoGoCVS + P1 + P2 + P3 + P4	MoGoCVS	$53.5\% \pm 0.16\%$

This leads to an overall success of 53.5% against MoGoCVS, obtained by BGP.

5.3 Improving 19x19 Go with database of patterns

In 19x19 Go, all tests are performed with 3500 simulations per move. Here also, we tested the case $|S_0| = 1$, *i.e.* we test one individual at a time. We tested only positive biases. The algorithm was launched for 62 iterations. Unfortunately, none of these 62 iterations was accepted. Therefore, we concluded that improving these highly optimized version was too difficult. We switched to another goal: having the same efficiency with faster simulations and less memory (the big database of patterns strongly slows the simulations and takes a lot of simulations), as discussed below.

5.4 Improving 19x19 Go without database of patterns

We therefore removed all the database of patterns; the simulations of MoGo are much faster in this case, but the resulting program is nonetheless weaker because simulations are far less efficient (see *e.g.* [20]). Fig. 1 (right) presents a known (from Senseis <http://senseis.xmp.net/?GoodEmptyTriangle#toc1>) difficult case for patterns: move 2 is a good move in spite of the fact that locally (move 2 and locations at the east, north, and north east) form a known very bad pattern (termed empty triangle), termed empty triangle, and is nonetheless a good move due to the surroundings.

We keep $|S_0| = 1$, 127 iterations. There were six patterns validated, validated at iterations 16, 22, 31, 57, 100 and 127. We could validate these patterns Q1, Q2, Q3, Q4, Q5, Q6 as follows. MoGoCVS+AE means MoGoCVS equipped with the big database of patterns extracted from games between humans.

Tested code	Opponent	Success rate
MoGoCVS + Q1	MoGoCVS	50.9% $\pm$ 0.13%
MoGoCVS + Q1 + Q2	MoGoCVS + Q1	51.2% $\pm$ 0.28%
MoGoCVS + Q1 + Q2 + Q3	MoGoCVS + Q1 + Q2	56.7% $\pm$ 1.50%
MoGoCVS + Q1 + ... + Q4	MoGoCVS + Q1 + Q2 + Q3	52.1% $\pm$ 0.39%
MoGoCVS + Q1 + ... + Q5	MoGoCVS + Q1 + ... + Q4	51.1% $\pm$ 0.20%
MoGoCVS + Q1 + ... + Q6	MoGoCVS + Q1 + ... + Q5	54.1% $\pm$ 0.78%
MoGoCVS + Q1 + Q2	MoGoCVS	53.4% $\pm$ 0.50%
MoGoCVS + Q1 + Q2 + Q3	MoGoCVS	57.3% $\pm$ 0.49%
MoGoCVS + Q1 + ... + Q4	MoGoCVS	59.4% $\pm$ 0.49%
MoGoCVS + Q1 + ... + Q5	MoGoCVS	58.6% $\pm$ 0.49%
MoGoCVS + Q1 + ... + Q6	MoGoCVS	61.7% $\pm$ 0.49%
MoGoCVS	MoGoCVS + AE	26.6% $\pm$ 0.20%
MoGoCVS + Q1	MoGoCVS + AE	27.5% $\pm$ 0.49%
MoGoCVS + Q1 + Q2	MoGoCVS + AE	28.0% $\pm$ 0.51%
MoGoCVS + Q1 + Q2 + Q3	MoGoCVS + AE	30.9% $\pm$ 0.46%
MoGoCVS + Q1 + ... + Q4	MoGoCVS + AE	32.1% $\pm$ 0.43%
MoGoCVS + Q1 + ... + Q5	MoGoCVS + AE	30.9% $\pm$ 0.46%
MoGoCVS + Q1 + ... + Q6	MoGoCVS + AE	32.8% $\pm$ 0.47%

An important property of BGP is that all validated patterns are confirmed by these independent experiments. We see however that in 19x19, we could reach

roughly 30% of success rate against the big database built on human games (therefore our BGP version uses far less memory than the other version); we will keep this experiment running, so that maybe we can go beyond 50 %. Nonetheless, we point out that we already have 60 % against the version without the database, and the performance is still increasing (improvements were found at iterations 16,22,57,100,122,127, with regular improvements - we have no plateau yet) - therefore we successfully improved the version without patterns, which is lighter (90% of the size of MoGoCVS is in the database).

6 Conclusions

We proposed an original tool for genetic programming. This tool is quite conservative: it is based on a set of admissible modifications, and has strong theoretical guarantees. Interestingly, the application of this theory to GP was successful, with in particular the nice property that all patterns selected during the GP run could be validated in independent experiments. We point out that when humans test modifications of MoGo, they usually test their algorithms based on simple confidence intervals, without taking into account the fact that, as they test multiple variants, one of these variants might succeed just by chance - it happened quite often that modifications accepted in the CVS were later removed, causing big delays and many non-regression tests. This is in particular true for this kind of applications, because the big noise in the results, the big computational costs of the experiments, imply that people can't use p-values like 10^{-10} - with BGP, the confidence intervals can be computed at a reasonable confidence level, and the algorithm takes care by itself of the risk due to the multiple simultaneous hypothesis testing. In 9x9 Go, BGP outperformed human development, and the current CVS of MoGo is the version developed by BGP. In 19x19 Go, we have an improvement over the default version of MoGo, but not against the version enabling the use of big databases - we nonetheless keep running the experiments as the success rate is still increasing and we had a big improvement for light versions.

References

1. J.-Y. Audibert, R. Munos, and C. Szepesvari. Use of variance estimation in the multi-armed bandit problem. In *NIPS 2006 Workshop on On-line Trading of Exploration and Exploitation*, 2006.
2. P. Audouard, G. Chaslot, J.-B. Hoock, J. Perez, A. Rimmel, and O. Teytaud. Grid coevolution for adaptive simulations; application to the building of opening books in the game of go. In *Proceedings of EvoGames*, 2009.
3. P. Auer, N. Cesa-Bianchi, and P. Fischer. Finite time analysis of the multiarmed bandit problem. *Machine Learning*, 47(2/3):235–256, 2002.
4. A. Auger and O. Teytaud. Continuous lunches are free plus the design of optimal optimization algorithms. *Algorithmica*, page 2009.
5. B. Bruegmann. Monte carlo go, 1993.

6. T. Cazenave and N. Jouandeau. On the parallelization of UCT. In *Proceedings of CGW07*, pages 93–101, 2007.
7. G. Chaslot, J.-T. Saito, B. Bouzy, J. W. H. M. Uiterwijk, and H. J. van den Herik. Monte-Carlo Strategies for Computer Go. In P.-Y. Schobbens, W. Vanhoof, and G. Schwanen, editors, *Proceedings of the 18th BeNeLux Conference on Artificial Intelligence, Namur, Belgium*, pages 83–91, 2006.
8. G. Chaslot, M. Winands, J. Uiterwijk, H. van den Herik, and B. Bouzy. Progressive strategies for monte-carlo tree search. In P. Wang et al., editors, *Proceedings of the 10th Joint Conference on Information Sciences (JCIS 2007)*, pages 655–661. World Scientific Publishing Co. Pte. Ltd., 2007.
9. G. Chaslot, M. Winands, and H. van den Herik. Parallel Monte-Carlo Tree Search. In *Proceedings of the Conference on Computers and Games 2008 (CG 2008)*, 2008.
10. 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.
11. R. Coulom. Computing elo ratings of move patterns in the game of go. In *Computer Games Workshop, Amsterdam, The Netherlands*, 2007.
12. M. Crasmaru. On the complexity of Tsume-Go. 1558:222–231, 1999.
13. S. Gelly, J. B. Hoock, A. Rimmel, O. Teytaud, and Y. Kalemkarian. The parallelization of monte-carlo planning. In *Proceedings of the International Conference on Informatics in Control, Automation and Robotics (ICINCO 2008)*, pages 198–203, 2008. To appear.
14. S. Gelly and D. Silver. Combining online and offline knowledge in UCT. In *ICML '07: Proceedings of the 24th international conference on Machine learning*, pages 273–280, New York, NY, USA, 2007. ACM Press.
15. J. H. Holland. Genetic algorithms and the optimal allocation of trials. *SIAM J. Comput.*, 2(2):88–105, 1973.
16. H. Kato and I. Takeuchi. Parallel monte-carlo tree search with simulation servers. In *13th Game Programming Workshop (GPW-08)*, November 2008.
17. L. Kocsis and C. Szepesvári. Bandit-based monte-carlo planning. In *ECML'06*, pages 282–293, 2006.
18. J. R. Koza. *Genetic Programming: On the Programming of Computers by means of Natural Evolution*. MIT Press, Massachusetts, 1992.
19. T. Lai and H. Robbins. Asymptotically efficient adaptive allocation rules. *Advances in Applied Mathematics*, 6:4–22, 1985.
20. C.-S. Lee, M.-H. Wang, G. Chaslot, J.-B. Hoock, A. Rimmel, O. Teytaud, S.-R. Tsai, S.-C. Hsu, and T.-P. Hong. The Computational Intelligence of MoGo Revealed in Taiwan's Computer Go Tournaments. *IEEE Transactions on Computational Intelligence and AI in games*, 2009.
21. D. Lichtenstein and M. Sipser. Go is polynomial-space hard. *J. ACM*, 27(2):393–401, 1980.
22. V. Mnih, C. Szepesvári, and J.-Y. Audibert. Empirical Bernstein stopping. In *ICML '08: Proceedings of the 25th international conference on Machine learning*, pages 672–679, New York, NY, USA, 2008. ACM.
23. J. M. Robson. The complexity of go. In *IFIP Congress*, pages 413–417, 1983.
24. P. Rolet, M. Sebag, and O. Teytaud. Optimal active learning through billiards and upper confidence trees in continuous domains. In *Proceedings of the ECML conference*, 2009.
25. Y. Wang and S. Gelly. Modifications of UCT and sequence-like simulations for Monte-Carlo Go. In *IEEE Symposium on Computational Intelligence and Games, Honolulu, Hawaii*, pages 175–182, 2007.