

# Twist: a trivially simple draughts engine + extensible microframework for game-playing agents

Tobia Tesan

## Abstract

In this document we give an overview of adversarial search techniques and introduce the game of draughts and its role in AI research in section 1.

We detail the requirements and implementation of a general framework capable of hosting game-playing agents in section 2 and the implementation of a minimally simplified draughtsboard in subsection 2.1.

We quickly review the theoretical foundations and discuss the implementation of an Alpha-Beta agent equipped with quiescence search and the killer heuristic in subsection 3.1 and of a UCT-based agent in subsection 3.2.1.

Finally, in section 4 we draw conclusions and propose future extensions.

## Contents

<b>1</b>	<b>Introduction</b>	<b>3</b>
<b>2</b>	<b>A micro-framework for two-player games</b>	<b>4</b>
2.1	Implementation of the game of Draughts . . . . .	6
<b>3</b>	<b>Agents</b>	<b>7</b>
3.1	Minimax . . . . .	7
3.1.1	Quiescence Search Minimax . . . . .	9
3.1.2	Killer Heuristic . . . . .	11
3.2	Monte Carlo Tree Search . . . . .	11
3.2.1	Bandit problems and UCT . . . . .	12
<b>4</b>	<b>Conclusion and future work</b>	<b>15</b>
<b>A</b>	<b>Long listings</b>	<b>18</b>
A.1	UCT Agent . . . . .	18
A.2	BasicAlphaBeta . . . . .	20
A.3	SimpleKillerAlphaBeta . . . . .	21

## List of Tables

## List of Figures

1	Draughtsboard with PDN numbering . . . . .	3
2	A nonquiescent position . . . . .	10
3	Five iterations of MTCS . . . . .	13

## List of Listings

1	Game.scala . . . . .	5
2	Player.scala . . . . .	6
3	AbstractAlphaBeta . . . . .	8
4	NaiveDraughtsEvaluation . . . . .	9
5	QuiescenceAlphaBeta . . . . .	10

1		2		3		4	
	5		6		7		8
9		10		11		12	
	13		14		15		16
17		18		19		20	
	21		22		23		24
25		26		27		28	
	29		30		31		32

Figure 1: Draughtsboard with PDN numbering. Red starts on  $1 \div 12$ , White starts on  $21 \div 32$ ;  $1 \div 4$  and  $29 \div 32$  are called “King’s Row”; for a man to reach King’s Row on the opposite side of the board results in a promotion to king.

## 1 Introduction

Draughts is a well-known family of two-player games, the most played of which is English draughts (or American checkers).

The rules issued by the World Checkers Draughts Federation [Fed12], which claims for itself the title of “official world governing body for the game of Checkers”, define draughts as “a board game of skill played between two players who, following a fixed set of rules, attempt to win the game by either removing all of their opponent’s playing pieces from the draughts board, or by rendering their opponent’s pieces immobile.”

The draughtsboard for English draughts with its 32 squares is reproduced in Figure 1, with squares numbered according to PDN (Portable Draughts Notation); from here on, with “draughts” we will be referring to the game of English draughts as defined by the WCDF rules, which will be assumed to be known. PDN numbering will also be liberally used in code fragments.

With a branching factor estimated as 2.5 by Lu [Lu93] and 6.4 by Guerra [Gue11] draughts is a somewhat less complex game than the “drosophila of artificial intelligence”, chess, with its branching factor estimated to be around 40; this hasn’t prevented several high-profile attempts at devising draughts-playing engines that include early efforts by Samuel [Sam67] and, most famously, Chinook, written by the team of Jonathan Schaeffer, which proved a worthy opponent for world champion Marion Tinsley in 1996. [Sch+07]

Draughts is a weakly solved game, in other words perfect play by both sides leads to a draw [Sch+07]; moreover, in the classical framework of game theory Draughts is a two-player, deterministic, perfect-information game: for this class of games variations on the minimax algorithm have been standard since Shannon’s seminal 1950 work [Sha50] but are currently being challenged by Monte Carlo methods, particularly since Sylvain Gelly et al’s successful implementation [GW06] of Kocsis and Szepesvari’s UCT algorithm [KS06] in the Go engine MoGo.

## 2 A micro-framework for two-player games

Despite the popularity of computer draughts, when we set out to extend an existing program with MTCS or advanced evaluation functions and set up a framework to carry out benchmarks we were faced with a lack of appropriate programs in source code form in the public domain or under a sufficiently permissive license.

More precisely, nearly all programs we found fell into one of the following categories:

- “Industrial-strength” programs, generally hard to extend and modify, and either
  - architecturally complex
  - highly optimized at the expense of simplicity
  - low-level (e.g. written in C with extensive manual memory handling)
- Programs too tightly coupled with a specific agent or class of agent (typically Minimax)
- Programs too tightly coupled with their UI and/or with the human vs IA mode of play, hard to extend in order to allow for IA vs. IA play and statistics collection.

**Requirements** We eventually resolved to write our own, with the aim that it be

- Architecturally simple
- Strongly decoupled, in particular regarding agents and games, and versatile
- Written in a high-level language

**Choice of programming language** Scala [Ode] was chosen as a programming language; among its benefits are its high portability thanks to the JVM (and, indeed, its ability to go beyond what the JVM affords – for example the ability of its dialect Scala.js to run in a browser [Doe13]), a reasonable type system that can make debugging and writing correct programs easier and its ability to mix stateless, expressive functional programming with imperative programming with side effects, which is ideal for writing high-level functional code while retaining the ability enter algorithms found in literature in ALGOL-like languages verbatim.

**Architecture** The code is available at [Twist].

The architecture is almost entirely specified by files `Game.scala` and `Player.scala`, shown in Listing 1 and Listing 2.

In `Game.scala` positions are partitioned into `TerminalPositions` and their complement, `LivePositions`<sup>1</sup>.

---

<sup>1</sup>`LivePosition` has nothing to do with the notion of “dead position”, related to that of quiescence, introduced by Turing in 1950.

```

0 trait Game[G <: Game[G]] {
1   def startingPosition(): LivePosition[G]
2 }
3
4 trait Move[G <: Game[G]] {}
5
6 sealed abstract class Side {
7   def opposite(): Side
8 }
9
10 case object Min extends Side {
11   def opposite() = Max
12 }
13
14 case object Max extends Side {
15   def opposite() = Min
16 }
17
18 trait TerminalPosition[G <: Game[G]] {
19
20   /**
21    * @return 0 if draw, -1 if Min wins, +1 if Max wins
22    */
23   def utility: Integer
24 }
25
26 /**
27  * A position that is not terminal
28  */
29 trait LivePosition[G <: Game[G]] {
30
31   /**
32    * @return the side to move, /if/ there are available moves
33    */
34   def sideToMove: Side
35
36   /**
37    * @return a non-empty Move -> Position map
38    *
39    * Unless the search space is trivial it's advisable to lazily evaluate positions
40    */
41   def successor(): Map[Move[G], Either[LivePosition[G], TerminalPosition[G]]]
42 }

```

Listing 1: Game.scala

```

0 trait DebugStats[+A <: AI[_]] {
1   def getNodes: Int
2   def toString: String
3 }
4
5 case class MoveWithStats[G <: Game[G], +M <: Move[G], +S <: DebugStats[AI[G]](
6   val move: M,
7   val stats: S
8 )
9
10 sealed trait Player[G <: Game[G]] {
11   def apply(p: LivePosition[G]): Move[G]
12 }
13 trait Human[G <: Game[G]] extends Player[G]
14 trait AI[G <: Game[G]] extends Player[G] {
15   def debug(p: LivePosition[G]): MoveWithStats[G, Move[G], DebugStats[AI[G]]
16   def apply(p: LivePosition[G]): Move[G] = debug(p).move
17 }

```

Listing 2: `Player.scala`; in order to implement an agent it is sufficient to extend `AI` with an appropriate implementation of `debug`, which returns a move and a `DebugStats` object containing statistics about the search that yielded said move (e.g. how many nodes expanded, how many cuts...)

The method `successor`<sup>2</sup> acts as the authoritative generator of legal moves and returns a map of moves to successor states; lazy evaluation, afforded by Scala, is necessary to make this simple interface viable for games with non-trivial state space.

## 2.1 Implementation of the game of Draughts

An implementation of the game of Draughts is to be found in `Draughts.scala`, accompanied by a test suite living under `test/`.

The implementation of the game will not be discussed at length nor reproduced here, but for the details of practical relevance that follow.

Firstly, our board implements faithfully the WCDF rules with the exception of the definition of a draw given in §1.32.

For simplicity **a draw is reached after a fixed number of plys have been played without a winner**. The default number of plys is set to 100; Guerra gives the average game length as 60 ply [Gue11].

Moreover, we work under the implicit assumption that utility is in the range  $[-1, 1]$ , where 1 and -1 are the utility of a win for either side, as suggested in [Sha50].

---

<sup>2</sup>Modeled after the definitions presented in the second edition of [AIMA], rather than those in the third ed. that include an explicit `result` function

## 3 Agents

### 3.1 Minimax

The Minimax algorithm, first introduced by Shannon in [Sha50] and illustrated at length in [AIMA], performs a complete depth-first exploration of the game tree in order to derive a *minimax value* for each child node and play the optimal move.

Its naive implementation has time complexity  $O(b^m)$  and space complexity  $O(bm)$  [AIMA] and is thus impractical; its most popular refinement is  $\alpha/\beta$  pruning, likely independently invented in the 1950s by various authors.

Alpha-Beta pruning, also illustrated in great detail in [AIMA], propagates heuristic bounds on the value of a position while traversing the game tree, corresponding to the minimum (resp. maximum) value that the agent about to move (resp. agent's opponent) can achieve.

Subtrees that are outside this range are cut off, resulting, in the best case, in a  $O(b^{m/2})$  complexity.

In our family of Minimax implementations we will use the equivalent Negamax formulation, described in [KM75], for ease of implementation, in which the minimax value  $F$  of a position  $p$  is defined as

$$F(p) = \begin{cases} f(p) & d = 0 \\ \max(-F(p_1), \dots, -F(p_d)) & d > 0 \end{cases}$$

**Implementation** The skeleton of our Alpha-Beta implementation, found in the class `AbstractAlphaBeta` inside `Minimax.scala`, is shown in Listing 3.

The most basic concretization of `AbstractAlphaBeta` – the classical Alpha-Beta algorithm – is implemented in class `BasicAlphaBeta`, reproduced in appendix A.2.

Notice how `AbstractAlphaBeta[G]` expects that an `Evaluation` object, appropriate for the game `G` we want to apply our agent to, be passed as parameter.

A basic evaluation function is provided for the game of draughts in `Draughts.scala` and shown in figure Listing 4, along with further evaluation functions to be discussed in ??.

A cutoff is implemented by throwing an exception `Cut`, which will forcibly abandon the evaluation of current subtree.

**Move ordering** It is a well known fact that the efficiency of Alpha-Beta in terms of expanded nodes over naive Minimax is fully realized only through an appropriate ordering of moves.

We have provided a simple ordering function in `BasicDraughtsMoveOrdering` (not reproduced) that privileges

1. among capture moves, the ones that maximize the amount of captured pieces
2. among ordinary, non-capturing moves, those that reach a farther place on the board or travel a longer distance.

```

0 abstract class AbstractAlphaBeta[G <: Game[G]](e: MinimaxEvaluation[G],
1                                     o: AlphaBetaOrdering[G],
2                                     depth: Int,
3                                     maximize: Boolean = false)
4   extends AI[G] {
5   def onTerminal(t: TerminalPosition[G],
6                 plyLeft: Int,
7                 nega: Int): (Double, AlphaBetaStats[G])
8
9   def onStatic(l: LivePosition[G],
10              plyLeft: Int,
11              nega: Int): (Double, AlphaBetaStats[G])
12
13  def otherwise(l: LivePosition[G],
14              plyLeft: Int,
15              alpha: Double,
16              beta: Double,
17              nega: Int): (Double, AlphaBetaStats[G])
18
19  case class Cut(val v: Double, val stats: AlphaBetaStats[G]) extends Exception
20
21  def iter(p: Either[LivePosition[G], TerminalPosition[G]],
22          plyLeft: Int,
23          alpha: Double = Double.NegativeInfinity,
24          beta: Double = Double.PositiveInfinity,
25          nega: Int = 1): (Double, AlphaBetaStats[G]) =
26    p match {
27      case Right(t: TerminalPosition[G]) =>
28        onTerminal(t, plyLeft, nega)
29      case Left(l: LivePosition[G]) =>
30        if (plyLeft <= 0)
31          onStatic(l, plyLeft, nega)
32        else
33          otherwise(l, plyLeft, alpha, beta, nega)
34    }
35
36  def debug(
37    p: LivePosition[G]: MoveWithStats[G, Move[G], AlphaBetaStats[G]] = {
38    val evaluatedMoves = p.successor.map(mp => mp._1 -> iter(mp._2, depth))
39    val cumulativeStats = evaluatedMoves
40      .map(_._2._2)
41      .fold(new AlphaBetaStats[G](0, 0, 0, 0))(_ + _)
42    val rankedMoves = evaluatedMoves.toList
43      .map((t: ((Move[G]), (Double, AlphaBetaStats[G]))) => (t._1, t._2._1)) // Discard stats
44      .sortWith(_._2 < _._2)
45    if (maximize)
46      MoveWithStats[G, Move[G], AlphaBetaStats[G]](rankedMoves.last._1,
47        cumulativeStats)
48    else
49      MoveWithStats[G, Move[G], AlphaBetaStats[G]](rankedMoves.head._1,
50        cumulativeStats)
51  }
52 }

```

Listing 3: AbstractAlphaBeta



The assumption that motivates the second point is that the ordering is particularly critical in the early phases of the game, when the search space is larger and deeper<sup>3</sup>, and in opening and middle game players typically will try to advance, in order to ultimately reach King’s Row and obtain a promotion.

Subsection 3.1.2 discusses the killer heuristic as an improvement.

```

0 object NaiveDraughtsEvaluation extends MinimaxEvaluation[Draughts] {
1   def apply(p: LivePosition[Draughts]): Double =
2     p match {
3       case (p: LiveDraughtsPosition) =>
4         (p.board.filter(_ == Some(Man(Max))).size * 1 +
5          p.board.filter(_ == Some(Man(Min))).size * -1 +
6          p.board.filter(_ == Some(King(Max))).size * 2 +
7          p.board.filter(_ == Some(King(Min))).size * -2).toDouble / 24
8     }
9   }
10 }

```

Listing 4: NaiveDraughtsEvaluation

### 3.1.1 Quiescence Search Minimax

A key improvement in minimax-like algorithms is quiescence search, used to contrast the “horizon effect”, which is the tendency of naive Minimax to be blind to developments that human players would call “obvious” in consequence of a given move when they occur one or more moves after the cutoff depth.

A typical example is found in Figure 2: assuming White is to move, if Red were to take a material-based static evaluation for the position the resulting estimate might be exceedingly optimistic.

The horizon effect is acknowledged as early as [GEC67], where it is countered with the “secondary search” approach.

The evaluation function is therefore trustworthy only for positions that are *quiescent*.

The idea of quiescence search is to equip the agent with some heuristic function to determine quiescence and, if necessary, carry out a further search for nonquiescent positions that extends beyond cutoff depth.

**Implementation** We have equipped the Alpha-Beta algorithm with quiescence search by integrating a further instance of `BasicAlphaBeta`, as seen in Listing 5.

Our augmented agent launches a simple minimax search of a given depth `extraPlys` if at cutoff depth the position is non-quiescent.

We have chosen to use a naive implementation that inevitably takes the static evaluation after `extraPlys` further plys, even if the position reached then is nonquiescent, as this is the simplest way to guarantee termination irrespective of the choice of `q`.

---

<sup>3</sup>This is trivially true under the assumption that a draw is forced after the  $n$ -th ply without a winner, as discussed in subsection 2.1




Figure 2: A nonquiescent position. White moves.

```

0 /**
1  * Simple AlphaBeta enhanced with quiescence search
2  *
3  * @param extraPlys fixed depth of extra search if "ordinary" search depth is exhausted
4  * on a non-quiescent position
5  */
6 class QuiescenceAlphaBeta[G <: Game[G]](e: MinimaxEvaluation[G],
7     o: AlphaBetaOrdering[G],
8     depth: Int,
9     q: QuiescenceCheck[G],
10    extraPlys: Int,
11    maximize: Boolean = false)
12 extends BasicAlphaBeta[G](e, o, depth, maximize) {
13 override def iter(l: Either[LivePosition[G], TerminalPosition[G]],
14     plyLeft: Int,
15     alpha: Double = Double.NegativeInfinity,
16     beta: Double = Double.PositiveInfinity,
17     nega: Int = 1): (Double, AlphaBetaStats[G]) =
18 1 match {
19 case Right(t: TerminalPosition[G]) => onTerminal(t, plyLeft, nega)
20 case Left(l: LivePosition[G]) =>
21   if (plyLeft <= 0)
22     if (q(l))
23       onStatic(l, plyLeft, nega)
24     else
25       // Position is non-quiescent, we do a further local, small
26       // minimax search
27       (new BasicAlphaBeta[G](e, o, extraPlys, maximize))
28         .iter(Left(l), extraPlys, alpha, beta, nega)
29   else
30     otherwise(l, plyLeft, alpha, beta, nega)
31 }
32 }

```

Listing 5: QuiescenceAlphaBeta

Notice how the quiescence search contributes to the collected statistics, which will allow us to estimate if and how it is competitive with a conventional minimax search search of depth `ply + extraPly`.

### 3.1.2 Killer Heuristic

Literature describes several variations on the “killer heuristic”; Gillogly, who used it in the TECH program, gives the common underlying intuition as follows: “a move which generates a prune in one set of moves may also generate a prune in the adjacent set (first cousin positions) so that this «killer» move should be tried first” [Gil72].

Therefore, programs using some form of killer heuristic will save moves that are refutations on a killer list or buffer, and examine the moves at each node as they are generated to see whether one of them matches a move on the killer list [Gil72].

According to [AN77], implementations can differ, among other things, on whether a separate list is kept for each ply.

We have therefore provided two separate implementations, `SimpleKillerAlphaBeta` (reproduced in appendix A.3) and `KillerAlphaBeta` (not reproduced): the former uses a single killer list, whereas the latter uses a separate killer list for each ply.

We have chosen for simplicity to make the killer list not persistent between different searches; this has been experimentally shown to suffice in yielding an appreciable gain in ??.

## 3.2 Monte Carlo Tree Search

Monte Carlo Tree Search, shortened in MTCS, is a kind of simulation-based search [GS11]; it belongs to a family of search algorithms that evaluate nodes by repeatedly carrying out *simulations* – i.e. descending a single path in a game tree according to a randomized *simulation policy* – in order to estimate their utility.

Abramson first introduced the expected outcome model in [Abr90] and shown that the game-theoretic value of a game-tree node is approximated by the expected value of the game’s outcome given random play from that node on.

An important difference with Minimax-based approaches is that the expected outcome approach considers the relative merit of game-tree nodes rather than board positions, and is thus, in its simplest formulation, independent of domain knowledge.

“Flat Monte Carlo” is the simplest simulation search algorithm and proceeds by uniform selection [Bro+12]; its drawback is that it makes an ineffective use of computational resources.

Intuitively, uniform selection implies that an equal number of simulations are spent on “interesting” and “uninteresting” moves alike. In the extreme case, if a position allows for exactly three moves, one of which directly leads to a loss and the other two are root to a large and “interesting” subtree that may or may not lead to wins, losses and draws, at the limit  $1/3$  of simulations will be spent on the former move, when they could be spent on gathering more information for moves two and three.

We would like to include features of a Shannon Type B strategy (one that focuses on “plausible moves”) in flat Monte Carlo, which in the limit is a Shannon type A strategy,

by focusing our efforts on promising moves – however, we wouldn’t want to run the risk of discarding a good move after a few unlucky simulations.

The Monte Carlo Tree Search algorithm thus incorporates the following steps [Bro+12] in every iteration, until a pre-allocated budget is exhausted:

- Starting from the root node, the tree is descended according to the *simulation policy*, until a non-terminal state with unexpanded children is reached and a node is then chosen to be expanded.
- A simulation is carried out according to the *default policy*.
- The result of the simulation is backpropagated through the parent nodes.

Figure 3, excerpted from [GS11], illustrates the progression of the algorithm on a game tree.

### 3.2.1 Bandit problems and UCT

As mentioned in the previous paragraph, we would like to focus on moves that appear to be “interesting”, but we wouldn’t want to run the risk of discarding a good move based on inconclusive data.

In this respect, our search problem can be framed as a “bandit problem”, an element in a class of problems akin to that of repeatedly choosing to play one out of  $k$  slot machines<sup>4</sup> with distinct, unknown distributions of rewards in order to maximise the cumulative reward.

This is the nature of the *exploitation-exploration dilemma* studied in bandit problems [Bro+12], i.e. the need to balance exploitation the “arm” believed to be optimal at a given time with the need to gather more information in order to have adequate confidence on the underlying distribution of arms that appear suboptimal.

Kocsis and Szepesvári [KS06] have shown that in Monte-Carlo tree search it is possible to treat each state of the search tree as a multi-armed bandit, in which each action corresponds to an arm of the bandit and have proposed UCT, an extremely popular variant of MTCS.

In UCT, the tree policy selects actions by using the UCB1 algorithm, which maximises an upper confidence bound on the value of actions; UCT is proven to converge on the minimax action value function.

When applying the tree policy to a node  $s$ , a child node  $a$  is selected to maximise

$$\underbrace{Q(s, m)}_{\text{MC value}} + c \underbrace{\sqrt{\frac{\log N(s)}{N(s, m)}}}_{\text{exploration bonus}}$$

where  $Q(s, m)$  is the “Monte-Carlo value” of taking move  $m$  from the position corresponding to node  $s$  – or the mean outcome of previous all simulations,  $N(s, m)$  is the number of times move  $m$  has been taken from  $s$  previously and  $N(s)$  is the number of times  $s$  has been expanded; the value is augmented by an exploration bonus that is highest for rarely visited state-action pairs [GS11].

---

<sup>4</sup>known as “one-armed bandits”

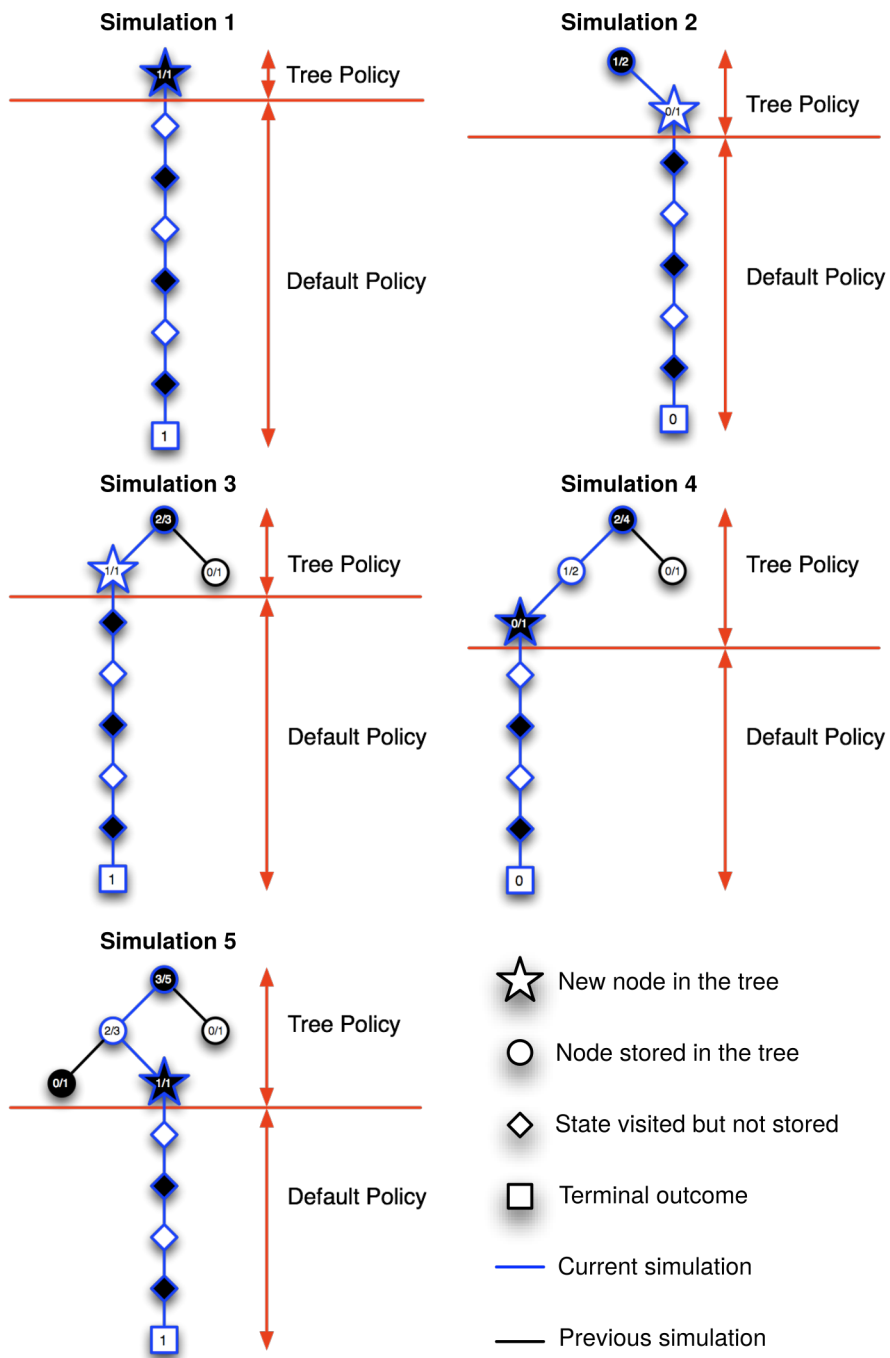


Figure 3: Five iterations of MTCS (from [GS11])

**Implementation** In `MTCS.scala`, reproduced in appendix A.1, we give a recursive implementation that follows closely the pseudocode for UCT reported in [GS11].

The budget is measured in total nodes expanded instead of wall clock time, in order to make it independent of optimizations in the implementation (or lack thereof).

Although domain knowledge can and often is incorporated in the tree policy – in fact, according to [Bro+12], the full benefit of MCTS is “typically not realised” until the algorithm is thus adapted – this basic implementation of MTCS is domain-agnostic.

## 4 Conclusion and future work

We have presented an architecture that can host different sorts of two-player, deterministic games and different sorts of agents that play such games.

We have presented an implementation of the game of draughts and two agents, one based on classical minimax and one based on MTCS.

A number of possible extensions and optimizations can be thought of; we give a necessarily non-exhaustive list.

The simplest possible extension probably is the addition of a new game using the framework specified in `Game.scala`.

Furthermore, the framework itself could be extended as to allow for games with more than two players, stochastic games such as Backgammon or games of imperfect information, for example through a subclassing of `Position` and a modification of class `Match` in order to serve an appropriately “obfuscated” version of each position to different players.

The UCT implementation could be extended with domain-specific optimizations and additional heuristics such as the history heuristic could be implemented in the Alpha-Beta agent.

Given enough CPU time, it would be interesting to run extensive benchmarks to characterize precisely the performance of agents and fit a function to predict the optimal value of  $c$  for UCT in relation to the allocated budget.

A further interesting nontrivial extension might be using a genetic, evolutionary or swarm-based approach to derive optimal evaluation functions.

While hardly an essential improvement, the usability of the framework also could be improved (thus facilitating the more fundamental additions discussed in the previous paragraphs) by providing an extended framework for benchmarks, extending class `Match` in order to realize a client-server model or, finally, adding a graphical user interface for play and visualization of statistics, perhaps leveraging `Scala.js` for in-browser execution.

## References

- [Abr90] Bruce Abramson. “Expected-outcome: A general model of static evaluation”. In: *IEEE transactions on pattern analysis and machine intelligence* 12.2 (1990), pp. 182–193.
- [AIMA] Stuart Russell, Peter Norvig, and Artificial Intelligence. “Artificial Intelligence: A Modern Approach”. In: *Artificial Intelligence. Prentice-Hall, Englewood Cliffs* (1995).
- [AN77] Selim G Akl and Monroe M Newborn. “The principal continuation and the killer heuristic”. In: *Proceedings of the 1977 annual conference*. ACM. 1977, pp. 466–473.
- [Bro+12] Cameron B Browne et al. “A survey of monte carlo tree search methods”. In: *IEEE Transactions on Computational Intelligence and AI in games* 4.1 (2012), pp. 1–43.
- [Doe13] Sébastien Doeraene. *Scala. js: Type-directed interoperability with dynamically typed languages*. Tech. rep. 2013.
- [Fed12] World Draughts Checkers Federation. *Rules of Checkers*. 2012.
- [GEC67] Richard D Greenblatt, Donald E Eastlake III, and Stephen D Crocker. “The Greenblatt chess program”. In: *Proceedings of the November 14-16, 1967, fall joint computer conference*. ACM. 1967, pp. 801–810.
- [Gil72] James J Gillogly. “The technology chess program”. In: *Artificial Intelligence* 3 (1972), pp. 145–163.
- [GS11] Sylvain Gelly and David Silver. “Monte-Carlo tree search and rapid action value estimation in computer Go”. In: *Artificial Intelligence* 175.11 (2011), pp. 1856–1875.
- [Gue11] Joao Carlos Correia Guerra. “Classical Checkers”. PhD thesis. Master’s thesis, Instituto Superior Técnico, 2011.
- [GW06] Sylvain Gelly and Yizao Wang. “Exploration exploitation in go: UCT for Monte-Carlo go”. In: *NIPS: Neural Information Processing Systems Conference On-line trading of Exploration and Exploitation Workshop*. 2006.
- [KM75] Donald E Knuth and Ronald W Moore. “An analysis of alpha-beta pruning”. In: *Artificial intelligence* 6.4 (1975), pp. 293–326.
- [KS06] Levente Kocsis and Csaba Szepesvári. “Bandit based Monte-Carlo planning”. In: *ECML*. Vol. 6. Springer. 2006, pp. 282–293.
- [Lu93] Chien-Ping Paul Lu. “Parallel search of narrow game trees”. MA thesis. University of Alberta, 1993.
- [Ode] Martin Odersky. *The Scala language specification*.
- [Sam67] Arthur L Samuel. “Some studies in machine learning using the game of checkers. II—recent progress”. In: *IBM Journal of research and development* 11.6 (1967), pp. 601–617.



- [Sch+07] Jonathan Schaeffer et al. “Checkers is solved”. In: *science* 317.5844 (2007), pp. 1518–1522.
- [Sha50] Claude E Shannon. “XXII. Programming a computer for playing chess”. In: *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 41.314 (1950), pp. 256–275.
- [Twist] Tobia Tesan. *Twist*. <https://github.com/tobiatesan/twist>. 2017.

# A Long listings

## A.1 UCT Agent

```
0 /**
1  * Implements an (inefficient) UCT agent based on Gelly 2011
2  *
3  * @param budget the number of total nodes (including default policy) to explore before ↔
4  *   stopping
5  * @param c the exploration constant; the larger it is, the more the algorithm favors exploration
6  *   over exploitation.
7  */
8 class UCTAgent[G <: Game[G]](val budget: Int,
9                               val c: Double,
10                              val maximize: Boolean = false,
11                              val r: Random)
12
13   extends AI[G] {
14     type Utility = Int
15     type NodeCount = Int
16
17     def SimDefault[G <: Game[G]](
18       p: Either[LivePosition[G], TerminalPosition[G]]: (Utility, NodeCount) =
19       p match {
20         case Right(t: TerminalPosition[G]) => (t.utility, 0)
21         case Left(l: LivePosition[G]) => {
22           val map = l.successor()
23           val randomMove = r.shuffle(map.keys.toList).head
24           val (util, nodes) = SimDefault((map.get(randomMove).get))
25           (util, nodes + 1)
26         }
27       }
28
29     def SelectMove[G <: Game[G]](p: LivePosition[G],
30                                  t: UCTNode[G],
31                                  c: Double): Move[G] = {
32       def N_s: Double = t.den.toDouble
33       def N_s_a(a: Move[G]): Double =
34         t.children.get(a).map(_.den).getOrElse(0).toDouble
35       def Q_s_a(a: Move[G]): Double =
36         t.children.get(a).map((a) => a.num / a.den).getOrElse(0).toDouble
37       def argmax[A](a: Seq[A], f: A => Double) = a.sortBy[Double](f).last
38       def argmin[A](a: Seq[A], f: A => Double) = a.sortBy[Double](f).head
39       if (t.maxNode)
40         argmax(p.successor.keys.toSeq,
41               (a: Move[G]) => Q_s_a(a) + c * Math.sqrt(Math.log(N_s) / N_s_a(a)))
42       else
43         argmin(p.successor.keys.toSeq,
44               (a: Move[G]) => Q_s_a(a) - c * Math.sqrt(Math.log(N_s) / N_s_a(a)))
45     }
46
47     /**
48     * @return Updated tree, spent budget and value of last simulation to be propagated
49     */
50     def Simulate(p: LivePosition[G],
```

```

49     node: UCTNode[G]): (UCTNode[G], NodeCount, Utility) = {
50     val move = SelectMove(p, node, c)
51     p.successor.get(move).get match {
52     case (Right(t)) =>
53         // Terminal game position, return utility
54         (UCTNode(p,
55             node.num + t.utility,
56             node.den + 1,
57             node.maxNode,
58             node.children),
59             1,
60             t.utility)
61     case (Left(l)) =>
62         if (node.children contains move) {
63             // Already in tree and not terminal, continue with tree policy
64             ((rec: (UCTNode[G], NodeCount, Utility)) => {
65                 (UCTNode(p,
66                     node.num + rec._3, // Backup
67                     node.den + 1,
68                     node.maxNode,
69                     node.children updated (move, rec._1)),
70                     rec._2 + 1,
71                     rec._3)
72                 })(Simulate(1, node.children.get(move).get))
73             } else {
74                 // Not in tree, not terminal: expand
75                 ((sim: (Utility, NodeCount)) => {
76                     (UCTNode[G](p,
77                         node.num + sim._1,
78                         node.den + 1,
79                         node.maxNode,
80                         (node.children updated (move,
81                             UCTNode(1, sim._1, 1, !(node.maxNode), Map.empty))),
82                         sim._2 + 1,
83                         sim._1)
84                     })(SimDefault(Left(1)))
85                 }
86             }
87         }
88
89     def UCTSearch(p: LivePosition[G],
90         budget: NodeCount): (Move[G], NodeCount, UCTNode[G]) = {
91         var tree = new UCTNode[G](p, 0, 0, maximize, Map.empty)
92         var budget_* = budget
93         while (budget_* > 0) {
94             val (newtree, spent, _) = Simulate(p, tree)
95             tree = newtree
96             budget_* = budget_* - spent
97         }
98         (SelectMove(p, tree, 0), budget - budget_*, tree)
99     }
100
101     def debug(
102         p: LivePosition[G]): MoveWithStats[G, Move[G], DebugStats[UCTAgent[G]]] =
103         ((x: (Move[G], Int, UCTNode[G])) =>

```

```

104 MoveWithStats[G, Move[G], DebugStats[UCTAgent[G]]](
105     x._1,
106     new UCTStats(x._2, x._3))(UCTSearch(p, budget))
107 }

```

## A.2 BasicAlphaBeta

```

0 class BasicAlphaBeta[G <: Game[G]](e: MinimaxEvaluation[G],
1     o: AlphaBetaOrdering[G],
2     depth: Int,
3     maximize: Boolean = false)
4 extends AbstractAlphaBeta[G](e, o, depth, maximize) {
5 def onTerminal(t: TerminalPosition[G],
6     plyLeft: Int,
7     nega: Int): (Double, AlphaBetaStats[G]) = {
8     (nega * t.utility.toDouble, new AlphaBetaStats(0, 1, 0, 0))
9 }
10
11 def onStatic(l: LivePosition[G],
12     plyLeft: Int,
13     nega: Int): (Double, AlphaBetaStats[G]) = {
14     (nega * e(l), new AlphaBetaStats(0, 0, 1, 0))
15 }
16
17 def otherwise(p: LivePosition[G],
18     plyLeft: Int,
19     alpha: Double,
20     beta: Double,
21     nega: Int): (Double, AlphaBetaStats[G]) = {
22 var alpha_* = alpha
23 var runningStats = new AlphaBetaStats[G](1, 0, 0, 0)
24 try {
25     p.successor.toSeq
26     .sortWith((x: (Move[G], _), y: (Move[G], _)) => o.lt(x._1, y._1))
27     .foreach { m =>
28         {
29             val (negv, stats) =
30                 this.iter(m._2, plyLeft - 1, -beta, -alpha_*, -nega)
31             val v = -negv
32             alpha_* = max(alpha_*, v)
33             runningStats = runningStats + stats
34             if (alpha_* >= beta)
35                 throw new Cut(alpha_*, runningStats)
36         }
37     }
38     (alpha_*, runningStats)
39 } catch {
40 case (p: Cut) => (p.v, p.stats + new AlphaBetaStats[G](0, 0, 0, 1))
41 }
42 }
43 }

```

### A.3 SimpleKillerAlphaBeta

```
0 /**
1  * Simple extension of BasicAlphaBeta with killer heuristic
2  * implemented with a _single_ list for all plys.
3  *
4  * According to Akl77, "programs differ in the number of killer moves
5  * saved, the number of matches looked for, and on whether a separate
6  * list is kept for each ply", so there is at least precedent.
7  *
8  * @param killerSize *total* size of the killer list
9  */
10 class SimpleKillerAlphaBeta[G <: Game[G]](e: MinimaxEvaluation[G],
11      o: AlphaBetaOrdering[G],
12      depth: Int,
13      killerSize: Int = 10,
14      maximize: Boolean = false)
15   extends BasicAlphaBeta[G](e, o, depth, maximize) {
16
17   var killerList: FiniteQueue[Move[G]] = new FiniteQueue(Queue(), killerSize)
18
19   object killerOrdering extends Ordering[Move[G]] {
20     /*
21     * if x is on the killer list and y is not x < y (comes first)
22     * if y " " then y < x
23     * defer to usual ordering otherwise
24     */
25     def compare(x: Move[G], y: Move[G]): Int = {
26       if (killerList.contains(x) && !killerList.contains(y))
27         -1
28       else if (!killerList.contains(x) && killerList.contains(y))
29         +1
30       else
31         o.compare(x, y)
32     }
33   }
34
35   override def debug(
36     p: LivePosition[G]): MoveWithStats[G, Move[G], AlphaBetaStats[G]] = {
37     // Wipe killer list at each new search
38     killerList = new FiniteQueue(Queue(), killerSize)
39     super.debug(p)
40   }
41
42   override def otherwise(l: LivePosition[G],
43     plyLeft: Int,
44     alpha: Double,
45     beta: Double,
46     nega: Int = 1): (Double, AlphaBetaStats[G]) = {
47     var alpha_* = alpha
48     var runningStats = new AlphaBetaStats[G](1, 0, 0, 0)
49
50     try {
51       l.successor()
52       .toSeq
```

```

53 .sortWith((x: (Move[G], _), y: (Move[G], _)) =>
54     killerOrdering.lt(x._1, y._1))
55 .foreach { m =>
56     {
57         val (negv, stats: AlphaBetaStats[G]) =
58             this.iter(m._2, plyLeft - 1, -beta, -alpha_*, -nega)
59         val v = -negv
60         alpha_* = max(alpha_*, v)
61         runningStats += (stats)
62         if (alpha_* >= beta) {
63             killerList = killerList.enqueue(m._1)
64             throw new Cut(alpha_*, runningStats)
65         }
66     }
67 }
68 (alpha_*, runningStats)
69 } catch {
70     case (c: Cut) => {
71         (c.v, c.stats + new AlphaBetaStats[G](0, 0, 0, 1))
72     }
73 }
74 }
75 }

```