

An Extensible and Modular Design and Implementation of Monte Carlo Tree Search for the JVM

Larkin Liu

University of Toronto

LARKIN.LIU@MAIL.UTORONTO.CA

Jun Tao Luo

Carnegie Mellon University

JTLUO@ANDREW.CMU.EDU

Abstract

Flexible implementations of Monte Carlo Tree Search (MCTS), combined with domain specific knowledge and hybridization with other search algorithms, can be a very powerful for the solution of problems in complex planning. We introduce *mctreesearch4j*, a standard MCTS implementation written as a standard JVM library following key design principles of object oriented programming. We define key class abstractions allowing the MCTS library to flexibly adapt to any well defined Markov Decision Process or turn-based adversarial game. Furthermore, our library is designed to be modular and extensible, utilizing class inheritance and generic typing to standardize custom algorithm definitions. We demonstrate that the design of the MCTS implementation provides ease of adaptation for unique heuristics and customization across varying Markov Decision Process (MDP) domains. In addition, the implementation is reasonably performant and accurate for standard MDP's. In addition, via the implementation of *mctreesearch4j*, the nuances of different types of MCTS algorithms are discussed.

Keywords: Monte Carlo Tree Search, Software Design, Markov Decision Process

1. Introduction

Sampling based strategies for the optimal planning of Markov Decision Processes (MDP) and adversarial games have been of great importance to the improvement to discovery of approximate solutions for complex high dimensionality domains where a full solution using complete methods are often infeasible. Algorithms such as Monte Carlo Tree Search (MCTS) (Kocsis and Szepesvári, 2006) have displayed a promising ability to generate competitive solutions for domains exhibiting high dimensionality search spaces.

In recent memory, MCTS used in combination with other game solving strategies have shown to be the world's state-of-the-art for high branching factor games such as Go (Silver et al., 2016). Other algorithms also frequently apply MCTS to solve for games with high branching factors such as Kriegspiel (Ciancarini and Favini, 2010) and Othello (Norvig, 1992). Our objective in this software is to create a extensible implementation of MCTS which is adaptable into multiple domains, applicable to either single player or multi player games.

©2021 Liu & Luo.

1.1 Discrete Markov Decision Process

The Markov Decision Process (MDP), described in (Puterman, 1994), is a well-developed framework for planning under uncertainty. The application of MDP's can be found in various optimization strategies to prescribe an optimal set of actions given a stochastic environment where perfect observations are available. An $MDP\langle S, A, \mathbb{P}, R \rangle$ designates a set of states S , where the agent traverses from S_t to S_{t+1} , for a horizon of T in T distinct time increments t (starting from $t = 1$). A is the action space, which is the set of actions the agent can take at any t and may or may not be a function of t itself. \mathbb{P} is the probability space that outlines the transition probability space for state-to-state transitions for any action taken at time t , $\mathbb{P} : \{S_t \times A \times S_{t+1}\} \rightarrow \{R \in \mathbb{R}\}$ subsequently returning a reward value, which belongs to \mathbb{R} .

In a fundamental MDP, the state at any time S_t is visible to the agent, and a corresponding reward is associated with each state. Therefore, if the state is observable, the value of each observable state is the reward associated with the state, and the availability of actions leading to rewards in the future. Evidently, they can form a recursive pattern to determine the value of the state, which can be solved via dynamic programming. Suppose we use $Q(S_t, a)$ to denote the immediate value when accounting for the immediate reward as well as the transitioned next state S_{t+1} . Where the Value of any state is denoted by $V(s)$.

$$Q(S_t, a) = R(S_t, a) + \gamma \sum_{S_{t+1} \in S} P(S_{t+1} | S_t, a) V(S_{t+1}) \quad (1)$$

$$V(S) = \operatorname{argmax}_{a \in A} Q(S, a) \quad (2)$$

Traditionally the exact solution to Eq. (2) can be found using dynamic programming algorithms such as Value Iteration or Policy Iteration (Bellman, 1957). Nevertheless, if the state space or branching factor is so large that traditional solutions such as dynamic programming cannot produce a solution within reasonable time or resource constraints, approximate methods via MCTS is a viable alternative for producing approximate optimal policies. The optimal policy denoted $\pi(S, a)$, is therefore the action, a that maximizes Eq. (2) at any particular state.

1.2 Monte Carlo Tree Search Algorithm

Monte Carlo Tree Search (MCTS) serves as an approach to solve the optimal policy for any well-defined MDP that does not always involve an exhaustive search through the entire search space. It has the potential to provide feasible solutions to high dimension and high branching MDP's, where dynamic programming techniques are not feasible provided constraints on computing resources. (Chang et al., 2010) demonstrates the specific use of the bandit algorithm, especially its ability to converge to the optimal policy of an MDP. Much research, such as (Coquelin and Munos, 2007) and (Kocsis and Szepesvári, 2006), demonstrates the effective usage of MCTS to obtain $\pi(S, a)$. The primary advantage of MCTS over dynamic programming solutions of the MDP is that, when implemented correctly, the number of iterations through the state space is drastically reduced. Nevertheless, the use

MONTE CARLO TREE SEARCH FOR THE JVM

of DP or MCTS is not mutually exclusive, as the combination of both algorithms can yield specific advantages and/or disadvantages (Feldman and Domshlak, 2014). The key notion of MCTS is the ability to guide the MC search down the tree efficiently, which borrows itself from the bandit strategy UCB1 (Auer et al., 2002).

$$UCT = E(S') + C\sqrt{\frac{2 \ln n}{n'}} \quad (3)$$

MCTS primarily makes use of a deterministic selection of actions and resulting outcomes to estimate the reward function of the system. MCTS is a tree search adaptation of the UCB1 Multi-Armed Bandit Strategy (Auer et al., 2002). When performing one tree traversal in MCTS, a series of actions is randomly played. This tree search is not entirely random as it is guided by the UCB1 illustrated in Eq. (3). The MCTS algorithm is distinctly divided into 4-phases, *Selection*, *Expansion*, *Simulation*, and *Backpropagation*, which are clearly illustrated in Fig. 1. In *Selection*, a policy deterministically selects which action to play, based on previously expanded states. This selection is typically guided by the UCT measure, from Eq. (3). In the *Expansion* phase, states that are unexplored are added to the tree in the fashion end of each node in the search tree. Subsequently, in the *Simulation* phase, a simulation is stochastically played out. Finally *Backpropagation* propagates the final reward of either a terminal state, or a node at an arbitrary depth limit, back to the root node. This 4-phase process is repeated until a maximum number of iterations or a convergence criteria is established.




Figure 1: Outline of MCTS - from (Browne et al., 2012).

2. Software Framework

mctreesearch4j is designed follow three key design principles. Firstly, we design for adaptability. Adaptability is defined as the ability for MDP domain to be easily integrated into the *mctreesearch4j* framework with ease via class abstractions. Our implementation seeks to simplify the adoption of MCTS solutions for a variety of domains. Secondly, we design a software that is fully compatible with the Java Virtual Machine (JVM), for the reasons to be discussed in Section 2.1. And lastly, we design to achieve a high degree of extensibility

within the framework. Extensibility is defined as the ability for key mechanisms to be reused, redefined, and enhanced, without sacrificing interoperability.

2.1 JVM Library

Open source implementations of MCTS exist, but have not gained widespread adoption. (Cowling et al., 2012) presents an implementation in Python, where multiple game engines can be plugged into the MCTS solver. (Lucas, 2018) presents a simple implementation written in Java, written with some categorization of the key MCTS mechanisms outlined in Fig. 1. However, both implementations do not provide easy access to heuristics, nor do they implement extensible and modular state action space abstractions. Chatzilygeroudis (2016) presents a low-level implementation of MCTS in C++. Yet because of its low-level implementation, it is not easy to interface a custom MDP definitions with the MCTS solver, hindering its adaptability across domains. Some MCTS implementations are available in Python, a widely adopted scripting language popular in Machine Learning, but suffer performance issues from utilizing an interpreter, causing computing overhead at runtime. Some workarounds do exist, such as dispatching to more performant libraries written in compiled code, but this is additional complexity we desire to avoid. Therefore, compiled languages such as Java or C are preferred for MCTS implementation primarily because of its more efficient resource consumption at runtime compared to native Python. Thus in the current landscape, there is a lack of open source MCTS implementations that provides the full list of features offered by *mctresearch4j* outlined in Section 2.

As of today, JVM language are widely adopted among many open source frameworks*. JVM languages are used in a variety of workloads including web, desktop and Android applications etc. The long history of the Java language also contributes to the JVM extensive knowledge base, language support and backwards compatibility. Furthermore, JVM is known for its wide adoption particularly in the creation of mobile games, many of which require online and/or offline planning solvable via MCTS. Therefore, we chose to develop an MCTS Framework for the JVM ecosystem.

The Kotlin programming language was chosen for the development of *mctresearch4j*. Kotlin features a more modern syntax that is succinct and expressive, containing powerful primitives, while simultaneously supporting a wide range of programming paradigms. This language is also popular among many software developers, particularly for the Android mobile operating system, a platform where *mctresearch4j* may find suitable adoption. The full compatibility of Kotlin with other JVM languages also allows the adoption by other languages within the JVM ecosystem. In addition Kotlin, along with many JVM languages, has strong support for type generics enabling us to build and reuse functions independent of types. It also offers better development experience with compile time type checks without explicit casts.

2.2 MDP Domain Abstraction

Attempts to rigorously define MDP interfaces to plug into custom MDP solvers have been attempted in the past. For example, RDDDL (Sanner, 2010) provides an interface for a formal

*<https://www.tiobe.com/tiobe-index/>

MONTE CARLO TREE SEARCH FOR THE JVM

Listing 1 Definition of the MDP abstract class.

```

1 abstract class MDP<StateType, ActionType> {
2     abstract fun transition(StateType, ActionType) : StateType
3     abstract fun reward(StateType, ActionType?, StateType) : Double
4     abstract fun initialState() : StateType
5     abstract fun isTerminal(StateType) : Boolean
6     abstract fun actions(StateType) : Collection<ActionType>
7 }

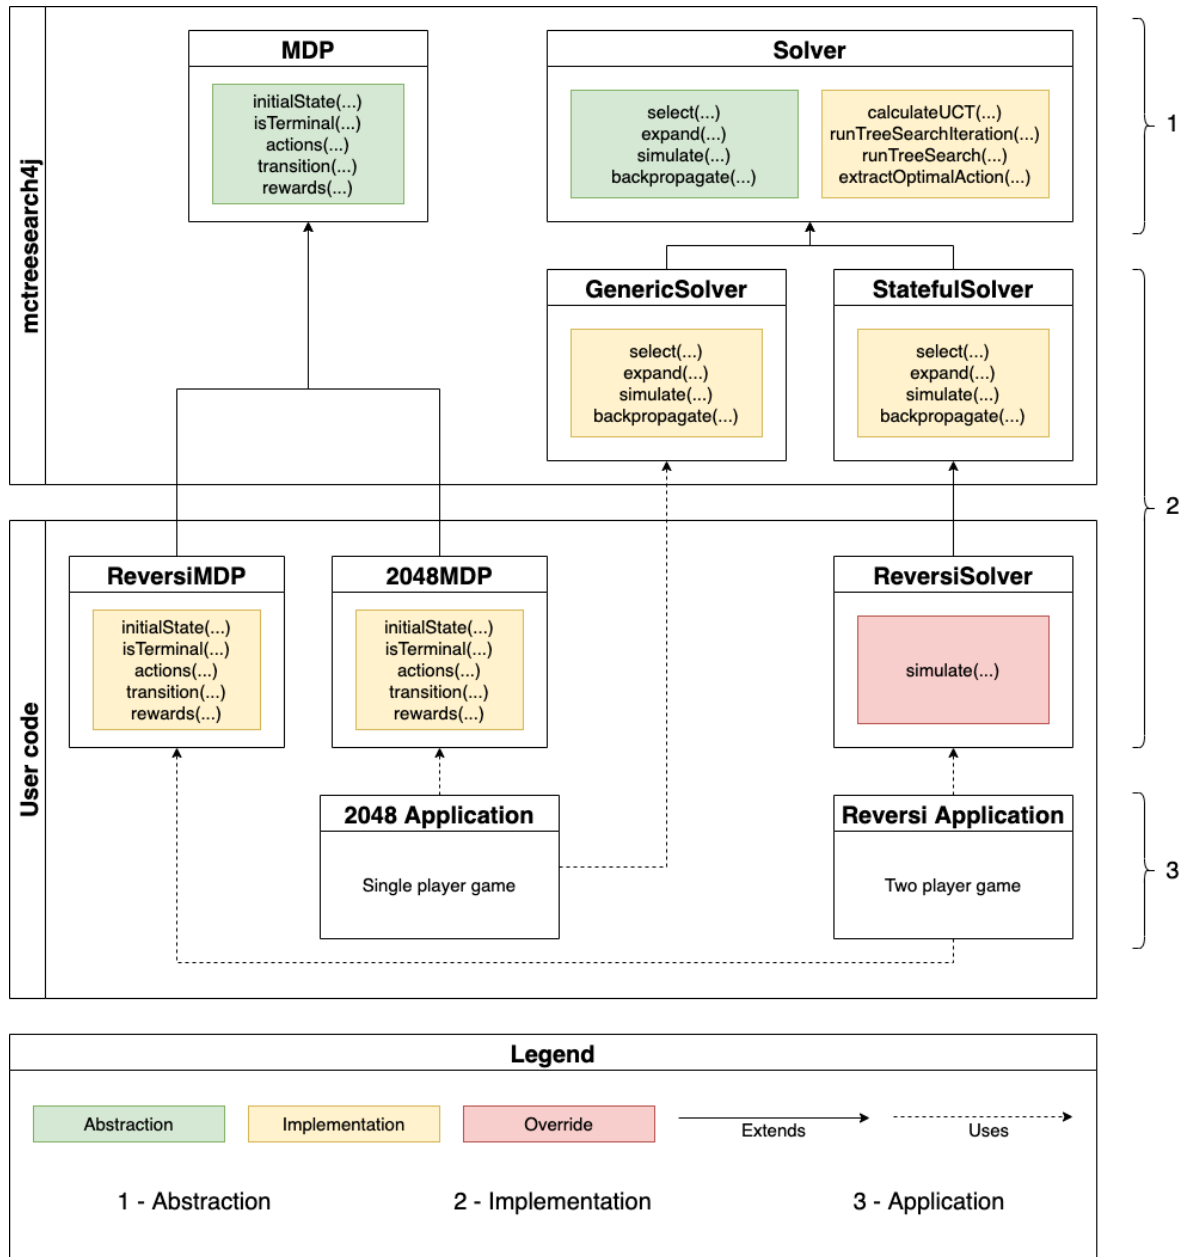
```

MDP definition, but it is rather complex for straightforward software implementations, and also adds additional overhead by introducing a new programming language specifically for the definition of MDP's. We simplify the formal MDP definition by abstracting the MDP into its base characteristics outlined in Listing 1.

The main abstraction that is used to define an MDP problem is the abstract class defined in Listing 1, using generic type variables. Each of the methods correspond to specific behaviour of a discrete MDP. In *mctreesearch4j* we use generic types **StateType** and **ActionType** to represent the MDP states and actions respectively. This abstract class has five members that must be implemented. These abstract class methods define the functionality of an MDP. The MDP abstraction will be used by core MCTS solvers to compute the optimal policy. The MDP interface can be written in any JVM language, we use Kotlin and Scala for this paper, with the Scala implementation from (Liu, 2021).

mctreesearch4j is designed to easily adapt to a variety of MDP definitions without complex mathematical constructs. The computational behaviour of *mctreesearch4j* is defined separately by user code outside of the core library, for example via a game controller, see Fig. 2. This allows programmers to define complex flexible systems, integrating seamlessly into turn-based game applications, including single and multi agent turn-based games such as 2048 and Reversi (Othello), respectively. In the user code layer of the framework, the programmer can define custom solvers, where unique heuristics can be implemented to add domain specific knowledge to improve the performance of the base MCTS solver. Thus the adaptability of the solver to a wide variety of MDP domains is one key design principle of *mctreesearch4j*, see Section 3 for an overview of adapted domains.

mctreesearch4j is able to use the MDP definitions to execute a tree search over the state action space. The MDP definitions used to interface with *mctreesearch4j* should optimize its save and load functionalities, as it can impact MCTS search performance if poorly constructed. Nevertheless, defining the MDP space programmatically gives the strong advantage of having the ability to define complex MDP's without the need for rigorous mathematical definitions. In our implementation, we aim to define a clear common declarative interface to provide abstractions for the solver and the MDP representation. The MDP abstraction ensures that the user defines the MDP scenario consistently such that it can be queried by the default solver provided in the library.

Figure 2: Inheritance and class structure of *mctreesearch4j*.

2.3 MCTS Solver Design

mctreesearch4j also implements the core Monte Carlo Tree Search algorithm in an extensible paradigm, maximizing the flexibility of the implementation by following Object Oriented Programming (OOP) principles. This is useful because it is possible to run MCTS with customized key mechanisms, and it is highly efficient if the programmer is able to reuse,

MONTE CARLO TREE SEARCH FOR THE JVM

Listing 2 Key Mechanisms of an MDP Solver

```

1 abstract class Solver<ActionType, NodeType>{
2     protected val verbose: Boolean,
3     protected val C: Double // Exploration Constant
4     ) where NodeType : Node<ActionType, NodeType> {
5     protected abstract var root: NodeType
6     abstract fun select(node: NodeType) : NodeType
7     abstract fun expand(node: NodeType) : NodeType
8     abstract fun simulate(node: NodeType) : Double
9     abstract fun backpropagate(node: NodeType, reward: Double)
10    open fun runTreeSearch(iterations: Int) { /* Predefined method */}
11 }

```

and/or redefine complete or partial segments of the key mechanisms. In Listing 2 we define such key mechanisms, all of which are inherited, and can be redefined or enhanced. This is important both for the efficiency of development, as well as the interoperability of the key MCTS mechanisms.

Listing 2 defines four abstract methods which are used to implement MCTS as illustrated in Fig. 1. The 4 phases, *Selection*, *Expansion*, *Simulation*, and *Backpropagation* defined in lines 6 to 9 serve as the key mechanisms to run *mctreesearch4j* and the class method `fun runTreeSearch()` is a predefined method from the base `class Solver()` responsible for executing the key MCTS solver mechanisms defined in Lines 6 to 9. These open functions, which can be overridden in subclasses, provide the flexibility to define the implementation of the solver that is appropriate for the MDP domain. Furthermore, we introduce the `NodeType` which is an abstract representation of the search space. It denotes how the state action space is represented in stages of the MCTS search, and must be explicitly defined by a subclass.

mctreesearch4j provides a default implementation known as `class GenericSolver`, and an alternate `class StatefulSolver`. The abstract `class Solver` serves as the base class for both versions, and defines a set of functionalities that all solver implementations must provide as well as a set of extensible utilities. Similar to the MDP abstraction, the solver uses a set of type parameters to provide strongly typed methods that unify and simplify the MCTS implementation. The programmer is free to select the data type or data structure that best defines how states and actions are represented in their MDP domain. Thus we can infer that, Generic and Stateful solvers have different representations of the `NodeType`.

The differentiation lies in their respective memory utilization of abstract node types to track states during MCTS iterations. The default `class GenericSolver` provides a leaner implementation, where actions are tracked and no explicit states are stored permanently. The states tracked with `class GenericSolver` are dynamic and the MDP transitions must be rerun when traversing the search tree during selection in order to infer the state. The `class StatefulSolver` keeps an explicit record of the visited states, and additional information, such as terminality and availability of downstream actions. The extra overhead of storing the state explicitly in the MCTS node, allows the algorithm to optimize its search

Listing 3 Generic Solver Selection.

```

1 fun select(node: ActionNode): ActionNode {
2     if (node.getChildren().isEmpty()) {return node}
3     var currentNode = node
4     simulateActions(node)
5     while (true) {
6         if (mdp.isTerminal(currentNode.state)) {return currentNode}
7         val currentChildren = currentNode.getChildren()
8         if (currentNode.isNotFullyExplored()) {return currentNode}
9         currentNode = currentChildren.maxByOrNull{a -> calculateUCT(a)}
10        simulateActions(currentNode)
11    }
12 }

```

using information from previously visited states. This is particularly useful for deterministic games, where a re-computation of the MDP transition is not necessary to determine the state of the agent after a particular taking a specific action. This differentiation results in different implementations of the *Selection* step, while maintaining identical implementations of *Expansion*, *Simulate* and *Backpropagation*. Both implementation iterates the algorithm in the most common pattern of *Selection*, *Expansion*, *Simulation*, *Backpropagation* until a certain number of iterations.

2.3.1 GENERIC SOLVER DESIGN

The Generic solver provides an MCTS solution where the states are not explicitly stored in the search tree, but instead inferred from the actions taken by the agent. Generic solvers are more memory efficient, and are preferred for scenarios of high state State-Action-State (SAS) branching factor. We define SAS branching factor as the average number of possible states that the agent can arrive at, given any action from any state. For stochastic domains, any action can lead to a multitude of states, so storing states explicitly in the search tree can be infeasible. In this case the Generic solver is preferred as it obviates the need to track states explicitly. A downside of the Generic solver is that additional computation is required to infer the state of the agent during the *Selection* phase of MCTS, as accomplished via the method on Line 4 of Listing 3. The Generic solver can also be extended to Partially Observable MDP's (POMDP's), where we abstract away the state space altogether, and solve over the estimation of state space as inferred by the action reward history, this is also known as the *Belief MDP* approach. In the Belief MDP, a POMDP is simplified into an MDP, by ignoring the potential discrepancy between the observation and true state, treating the current observation as the true state. This can create a tractable MDP solution at the cost of strict accuracy, and is in principle, fully programmable into *mctreesearch4j* via the Generic solver.

The most significant difference is the trade-off between memory and runtime. The Generic solver does not track the possible states for each node of the search tree and therefore reduces the memory consumption. The key benefit is the ability to track the state

MONTE CARLO TREE SEARCH FOR THE JVM

Listing 4 Stateful Selection.

```

1 fun select(node: StateNode): StateNode {
2     if (mdp.isTerminal(node.state)) {return node}
3     if (node.validActions.size > node.exploredActions.size) {return node}
4     var bestAction = exploredActions.getActionUCT(node, node.getChildren())
5     val newState = mdp.transition(node.state, bestAction)
6     val actionState = node.getChildren(bestAction)
7         .firstOrNull { s -> s.state == newState }
8     ?: return createNode(node, bestAction, newState)
9     return select(actionState)
10 }

```

space with actions alone. A disadvantage, is that the agent must replay the game engine to approximate the state at any time, which is where a computational bottleneck can arise.

2.3.2 STATEFUL SOLVER DESIGN

We provide an alternate solver design targeted for MDP domains with low SAS branching factors. In **StatefulSolver** the nodes of the solver explicitly track the state, the metadata contained within the state. An **NodeType** is used for Stateful search trees where the node explicitly contains both the state of the agent, and the inducing action that led to said state. The Stateful solver is computationally efficient when the SAS branching factor is feasible, as it is able to store and look up the current state of the agent and continue MCTS iterations from the recorded state. Stateful solvers are generally preferred for low state space MDP's, where the SAS branching factor, defined earlier in Section 2.3.1, is relatively manageable, such as in games where low SAS branching factors are exhibited. However, this optimization will incur additional memory usage.

Listing 4 illustrates the *Selection* and *Expansion* algorithm for Stateful MCTS. The Stateful solver keeps track of all possible states at each node illustrated in Lines 2 to 10. In the *Selection* phase, the Stateful solver iterates until a leaf node is reached.

2.3.3 EXPANSION, SIMULATION AND BACKPROPAGATION

The *Expansion*, *Simulation*, and *Backpropagation* steps for Generic and Stateful solvers remain virtually identical, notwithstanding specification of a different **NodeType** for each solver, also resulting minor nuances in instantiating said nodes in the Expansion phase. The base **NodeType**, which is used by both the Generic and Stateful Solvers, store key metadata such as the inducing action, expected reward, and number of visits. The **fun expand()** method randomly selects an unexplored node, and adds it to the search tree. This is done using the **fun createNode()** method, which differs between Generic and Stateful solvers.

In the Generic solver, no actual MDP state is required to create a new node of type **ActionNode**. The current state of the agent is computed based on the sequence of actions taken utilizing the specified **mdp** object, defined in Listing 1. The Stateful solver, on the other hand, explicitly requires the tracking of states when creating a **StateNode** via the **fun createNode()** method. The **fun simulate()** method runs a purely randomized ex-

ploration of the MDP state-action space from a specific state, until a terminal state or the simulation depth is reached. The `fun backpropagate()` method is used to update the reward values for the node where the simulation was run as well as all of its ancestors.

2.3.4 EXTENSIBILITY AND HEURISTICS

Though the default MCTS implementation described in Section 1.2 works well in many scenarios, there are situations where knowledge about specific problem domains can be applied to improve the MCTS performance. Improvements to MCTS, such as heuristics driven simulation, where domain knowledge can be exploited to improve solver performance. We demonstrate that a Reversi AI that uses heuristics derived from (Guenther and Klinik, 2004) is able to outperform the basic MCTS implementation. These heuristics are programmed via extensibility points in the *mctreesearch4j* solver implementation, where the key mechanisms can be altered or augmented. Our heuristic, illustrated in Fig. 3, introduces the `heuristicWeight` array, a 2D array storing domain specific ratings of every position on a Reversi board representing the desirability of that position on the board. The negative numbers represent a likely loss and positive numbers representing a likely win, again as represented in Fig. 3. This value is taken into consideration when traversing down the simulation tree. The `heuristicWeight` array adjusts the propensity to explore any position for both agents based on the heuristic's belief about the desirability of the position. The effectiveness of this heuristic is later presented in Section 3.3. To alter the MCTS simulation phase we override the `fun simulate()` method and create a new definition for it. The application of this `heuristicWeight` only requires minor alterations to the `fun simulate()` method, as illustrated from Lines 5 until 15 in Listing 6.




Figure 3: Reversi `heuristicWeight` lookup table.

The default solver behaviour continues running the MCTS iteration until the iteration count is reached. This may not be the best criteria and convergence metrics can be used to better determine when to stop the tree construction. For example, if no significant improvements in the best action's score is made, the algorithm can be terminated early. Additional functionality can also be added to the solver in the tree construction function. For example,

MONTE CARLO TREE SEARCH FOR THE JVM

Listing 5 Expansion, Simulation and Backpropagation pseudo-code.

```

1 fun expand(node: StateNode): StateNode {
2     if (node.isTerminal) {return node}
3     val unexploredActions = node.getUnexploredActions()
4     val actionTaken = unexploredActions.random()
5     return createNode(node, actionTaken)
6 }
7
8 fun simulate(node: NodeType): Double {
9     if (mdp.isTerminal(node.state)) {
10         return mdp.reward(node.parent?.state, node.inducingAction, node.state)
11     }
12     var depth = 0
13     var currentState = node.state
14     var discount = rewardDiscountFactor
15     while(true) {
16         val validActions = mdp.actions(currentState)
17         val randomAction = validActions.random()
18         val newState = mdp.transition(currentState, randomAction)
19         if (mdp.isTerminal(newState)) {
20             return mdp.reward(currentState, randomAction, newState) * discount
21         }
22         currentState = newState
23         depth++
24         discount *= rewardDiscountFactor
25         if (depth > simulationDepthLimit) {
26             return mdp.reward(currentState, randomAction, newState) * discount
27         }
28     }
29 }
30
31 fun backpropagate(node: NodeType, reward: Double) {
32     var currentNode = node
33     var currentReward = reward
34     while (true) {
35         currentStateNode.maxReward = max(currentReward, currentNode.maxReward)
36         currentStateNode.reward += currentReward
37         currentStateNode.n++
38         currentStateNode = currentStateNode.parent ?: break
39         currentReward *= rewardDiscountFactor
40     }
41 }

```

Listing 6 Heuristic implementation in Reversi.

```

1  override fun simulate(node: NodeType): Double {
2      /*... Original Simulation code ...*/
3      while(true) {
4          val validActions = mdp.actions(currentState)
5          var bestActionScore = Int.MIN_VALUE // Begin heuristic
6          var bestActions = mutableListOf<Point>()
7          for (action in validActions) {
8              val score = heuristicWeight[action.x][action.y]
9              if (score > bestActionScore) {
10                 bestActionScore = score
11                 bestActions = mutableListOf(action)
12             }
13             if (score == bestActionScore) {bestActions.add(action)}
14         }
15         val randomAction = bestActions.random() // End heuristic
16         val newState = mdp.transition(currentState, randomAction)
17         /*... Original Simulation code ...*/
18     }
19 }

```

if additional metrics need to be measured after each MCTS iteration, *mctreesearch4j* can be overridden to add this functionality. For example, the exploration term, C , from Eq. 3 can be exported via a class extension that tracks the exploration term.

3. Evaluation & Function

We evaluate our MCTS implementation on its ability to converge to an accurate solution. In our experiments, the MDP has one or many optimal solutions, and the solver should accurately converge to these solutions. We also measure the performance of the implementation by the average time to extract one optimal action taken from any state, we denote this simply as *runtime per decision* or Δ_d . Δ_d is an estimate of the time to make a decision averaged over 100 trials. This can vary from computer to computer, but serves as a rough estimate. We demonstrate that *mctreesearch4j* is adaptable across multiple domains, in addition to the fact that Δ_d can be kept at manageable levels for operation on a common laptop computer. The exact hardware specifications are outlined in Appendix A.1,

For adversarial search domains featuring two agents, the terminal state of victory or loss results in a score of +1 or -1 respectively, from the perspective of the agent when competing against the opponent. Later in Section 3.3 we demonstrate a adversarial setup where 2 MCTS solvers with novel characteristics, play against each other.

We examine the performance of the MCTS solver across five distinct MDP domains. These domains are GridWorld, a probabilistic single agent MDP illustrated in Fig. 4, the game of 2048 analyzed in Kohler et al. (2019), Push Your Luck (PYL), a single agent MDP

MONTE CARLO TREE SEARCH FOR THE JVM

defined in Geißer and Speck (2018), the game of Connect 4, a deterministic multi-agent game, and the game of Reversi, a deterministic adversarial multi-agent game.

3.1 Solver Runtime per Domain

Domain	No. Agents	Stochastic	Avg. Br. Fac.	Avg. Depth	Solver	Δ_d
GridWorld	1	yes	4	5	Gen.	0.1
PYL	1	yes	2	∞	Gen.	0.3
2048	1	yes	4	940	Gen.	0.5
Connect4	2	no	4	36	SF	0.06
Reversi	2	no	10	58	SF	0.12

Table 1: MDP domain characteristics, where Δ_d is in seconds.

Furthermore, we select either Generic or Stateful solvers based on the characteristics of the MDP domain, with the considerations outlined in Sections 2.3.2 and 2.3.1. The max tree depth of the MCTS was set to 1000, this number typically far exceeds the average depth before reaching a terminal state for most MDP domains, with the exception of 2048. The number of iterations set to 500 arbitrarily. The MDP domains, where *mctreesearch4j* was applied, contain one or two agents, and state action transitions can be stochastic or deterministic. We also note the average branching factor of the domain (Avg. Br. Fac.), and average depth (Avg. Depth)[†] of the search tree until a terminal state is reached. For some games, such as GridWorld, average depth is heavily determined by the state configuration, we refer to the configuration set out in Fig 4. We record the average runtime to compute the immediate next optimal action from an arbitrary state in the state space of the MDP domain in milliseconds[‡].

3.2 Solver Convergence

To demonstrate the accuracy of *mctreesearch4j*, we examine the performance of the solver with respect to convergence of rewards, convergence of exploration terms, and visiting of the optimal state subspace more than non-optimal state subspace with increasing iterations. Variables such as the exploration constant C , as denoted in Eq. (3), affect the overall exploration of *mctreesearch4j*. If C is too high, there can be over-exploration causing the UCT function to swing between optimal actions and non-optimal action selection. Over-exploration affects the smoothness of the exploration term convergence (James et al., 2017). If C is too low, then it is not able to sufficiently explore all states to find an optimal action.

We provide a simple example of convergence evaluation in the domain of Gridworld. The grid displayed in Fig. 4, represents the state space, and each cell demarcates the state which corresponds to a reward value (unfilled cell indicates a reward of 0). In our example any non-zero state represents a terminal state. The agent represented by \diamond , has 4 possible actions from its starting point. In the specific configuration shown by Fig. 4, the MCTS

[†]https://en.wikipedia.org/wiki/Game_complexity illustrates Avg. Br. Fac. and Avg. Depth.

[‡]The exact specification of the computer used to run the simulation is contained in Appendix A.1.

solver should ultimately select the actions \leftarrow or \downarrow , which are optimal, and *mctreesearch4j* should accurately converge to the appropriate values.




Figure 4: Illustration of the Gridworld MDP domain. An agent, is represented by \diamond , lies in a domain with terminal state reward +5, and terminal reward negative state -1.




Figure 5: Convergence of exploration terms

Our MCTS solver[§], with no additional heuristics, is capable of converging to the correct solution space. Fig. 5 displays a convergence of the exploration term $C\sqrt{2\log n/n'}$, which decreases as $n \rightarrow \infty$, and $n' \leq n$. Provided this, the correct optimal policy will yield smaller exploration terms for optimal versus non-optimal solutions because the number of visits to the non-optimal nodes, n' , will be generally smaller than the optimal nodes.

[§]The figures 5 to 7 represent solutions of the Stateful solver, nevertheless convergence plots using Generic solver yields similar behaviour.

MONTE CARLO TREE SEARCH FOR THE JVM




Figure 6: Convergence of rewards

Naturally, the reward value associated with each optimal state induced by the optimal policy action at $t + 1$ should yield higher backpropagated rewards. In Fig. 6 we see that the two optimal actions produce higher reward values than non-optimal solutions.




Figure 7: Convergence of visits

When the MCTS solver is accurately selecting the optimal solutions, it will continue to cause the agent to explore in the optimal subset of the state space, and reduce its exploration in the non-optimal subset of the state space as evidenced in Fig. 7. The cumulative number of visits corresponding to the optimal policy is proportionally increasing with respect to the number of MCTS iterations. Whereas for non-optimal solutions, the cumulative visits are significantly less because the solver will avoid visiting the non-optimal state subspace.

3.3 Adversarial Game Benchmarking

We demonstrate the possibility of heuristic adaptation to improve *mctreesearch4j* in an adversarial game. For this, we create a two-player game in the deterministic multiplayer domain of Reversi, where our AI will play as adversaries against each other. The Reversi heuristic is implemented to enable the search algorithm to compute more realistic estimates of reward, using specific domain knowledge (Guenther and Klinik, 2004). The intended effect of this heuristic is to enable a more efficient search of the state action space, as presented earlier in Section 2.3.4

The heuristic augmented *mctreesearch4j* is played against the original *mctreesearch4j* to prove that the adversarial MDP solver is capable of improving via domain knowledge. Table 2 demonstrates that when *mctreesearch4j* is augmented with heuristics it defeats the original *mctreesearch4j* solver up to 75% of the time, under ceteris paribus conditions. Nevertheless this win rate metric depends on several factors. Firstly, we see that the heuristic solver is only marginally better than the original solver when we limit the maximum simulation depth (Max. Sim. Depth), and the number of simulation iterations (No. Iter). The limitation on these two hyper-parameters limit the MCTS solver’s ability to explore the entire state action space to determine accurate rewards, thereby also limiting the effectiveness of the heuristic. The degree of contrast between *mctreesearch4j* with heuristic and without is greatly enhanced when the solver is provided enough simulation iterations and simulation depth to fully take advantage of the additional domain information.

No. Iter	Max. Sim. Depth	No. Wins	No. Loss	No. Ties	Win Rate
50	40	98	87	15	0.49
500	40	121	75	4	0.605
50	1000	111	75	14	0.555
500	1000	151	46	3	0.755

Table 2: Reversi results comparing MCTS solver with heuristics and without.

Via this experiment, we justify that our MCTS solver is capable of being modified to adapt new heuristics which improve the outcomes of the MDP policy for an adversarial game. The heuristic presented here actively biases the random simulation of MCTS to favour actions which induce better outcomes, based on domain knowledge. We also see that as we adjust hyper-parameters of the model, such as simulation depth and number of iterations, the heuristic augmented solver can become more effective.

4. Conclusion & Future Work

Due to the extensibility and modularity of *mctreesearch4j*, many improvements and experiments could be performed by extending and modifying the capabilities of the base software. *mctreesearch4j* is designed to enable researchers to experiment with various MCTS strategies while standardizing the functionality of the MCTS solver and ensuring reliability, where the experiments are reproducible and the solver compatible with common JVM runtimes [¶].

[¶]To be specific *mctreesearch4j* is fully compatible with Java Runtime Environment (JRE) 8 and 11.

4.1 Future Work

We encourage the effort to enhance *mctreesearch4j* further. For example, while the default implementations create a new node and adds it to the node chosen by the *Selection* phase, it is possible to rely on domain knowledge to avoid superfluous game states. For example, in Gridworld, where each unique position is functionally the same regardless of how it arrived to that position, there is no need to add a newly created node to the search tree. Instead, a cache of nodes representing known positions can be used to dramatically reduce the MCTS search space, yielding more accurate results. Furthermore, creating an mapping for more expressive and rigorous MDP defining meta-languages such as RDDL (Sanner, 2010) to bind to *mctreesearch4j*, could be an important feature in the near future. This enables complex and rigorous definitions of MDP's to be benchmarked using *mctreesearch4j*, opening many new opportunities for research.

Also, as described in Baier and Winands (2013), some games can be solved more efficiently if Minimax algorithms can be applied in the *Selection* phase. Instead of the default solver which uses UCT to select leaf nodes, a Minimax algorithm can be used instead to detect shallow wins and losses. This has been shown to work well for adversarial games to avoid traps set by the opposite player. For adversarial games, another common strategy is to improve simulation accuracy by incorporating a Minimax algorithm for the *Simulation* phase. This ensures the rewards computed are more realistic than two players who play randomly. This can significantly improve the MCTS performance. In addition, the *backpropagation* step can be augmented with Minimax concepts to mark known losing positions. This allows subsequent *Selection* phase to ignore parts of the search tree and speed up convergence (Baier and Winands, 2013). Finally, instead of running simulated games to compute a reward, it is possible to rely on an external source to predict the reward value of a given state. An example of this is the state-of-the-art Go AI, AlphaGo Silver et al. (2016), which uses Deep Learning trained on expert positions and via self-play to compute the expected reward for a given game position. Combining Deep Learning techniques with *mctreesearch4j* to provide external intelligence is an open avenue of research.

The phase order of MCTS defined in the base `class Solver()` can also be modified to explore different avenues. Although the most common MCTS phase order of *Selection* → *Expansion* → *Simulation* → *Backpropagation* works well in most scenarios, the MCTS phase order can be altered to *Selection* → *Simulation* → *Expansion* → *Backpropagation* as described in Winands et al. (2010). The advantages and/or disadvantages of altering the MCTS phase order is an under explored topic.

4.2 Summary

In closing, *mctreesearch4j* presents a framework which enables programmers to adapt an MCTS solver to a variety of MDP domains. This is important because software application was a main focus of *mctreesearch4j*. Furthermore, *mctreesearch4j* is fully compatible with JVM, and this design decision was made due to the excellent support of class structure and generic variable typing in Kotlin, and other JVM languages, as well as support for mobile applications. Yet most importantly, *mctreesearch4j* is modular and extensible, the key mechanism of MCTS are broken down, and the programmer is able inherit class charac-

teristics, redefine and/or re-implement certain sections of the algorithm while maintaining a high degree of MCTS standardization.

4.3 Software Availability

- <https://mvnrepository.com/artifact/ca.aqtech/mctreesearch4j> contains the compiled core library for *mctreesearch4j* available via the Maven Central Repository for the JVM ecosystem.
- <http://mctreesearch4j.aqtech.ca/> contains the source code for implementation of *mctreesearch4j* core library, as well as programs for various domains and game controllers.

Acknowledgements

We would like to acknowledge the kind support of our reviewers, and Maven Central Repository for repository hosting of our code base.

Appendix A.

A.1 Simulation Hardware Specifications

All of the benchmarks run in Section 3 were performed on an Apple Macbook Pro, containing a 2.6GHz dual-core Intel Core i5 processor with 3MB shared L3 cache with 8GB of 1600MHz DDR3L onboard memory. We expect that a computer with similar hardware capabilities will yield similar results.

References

- Peter Auer, Nicolò Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Mach. Learn.*, 47(2-3):235–256, May 2002. ISSN 0885-6125. doi: 10.1023/A:1013689704352. URL <https://doi.org/10.1023/A:1013689704352>.
- Hendrik Baier and Mark H. M. Winands. Monte-carlo tree search and minimax hybrids. In *2013 IEEE Conference on Computational Intelligence in Games (CIG)*, pages 1–8, 2013. doi: 10.1109/CIG.2013.6633630.
- Richard Bellman. A markovian decision process. *Journal of Mathematics and Mechanics*, 6(5):679–684, 1957. URL <http://www.jstor.org/stable/24900506>.
- Cameron Browne, Edward Powley, Daniel Whitehouse, Simon Lucas, Peter I. Cowling, Stephen Tavener, Diego Perez, Spyridon Samothrakis, Simon Colton, and et al. A survey of monte carlo tree search methods. *IEEE TRANSACTIONS ON COMPUTATIONAL INTELLIGENCE AND AI*, 2012.
- Hyeong Soo Chang, Jiaqiao Hu, Michael C. Fu, and Steven I. Marcus. Adaptive adversarial multi-armed bandit approach to two-person zero-sum markov games. *IEEE Trans. Autom. Control.*, 55(2):463–468, 2010. doi: 10.1109/TAC.2009.2036333. URL <https://doi.org/10.1109/TAC.2009.2036333>.

MONTE CARLO TREE SEARCH FOR THE JVM

- Konstantinos Chatzilygeroudis. Monte carlo tree search - c++. <https://github.com/resibots/mcts>, 2016.
- Paolo Ciancarini and Gian Piero Favini. Monte carlo tree search in kriegspiel. *Artificial Intelligence*, 174(11):670–684, 2010. ISSN 0004-3702. doi: <https://doi.org/10.1016/j.artint.2010.04.017>. URL <https://www.sciencedirect.com/science/article/pii/S0004370210000536>.
- Pierre-Arnaud Coquelin and Rémi Munos. Bandit algorithms for tree search. *CoRR*, abs/cs/0703062, 2007. URL <http://arxiv.org/abs/cs/0703062>.
- Peter Cowling, Edward Powley, and Daniel Whitehouse. Information set monte carlo tree search. *IEEE Transactions on Computational Intelligence and Ai in Games*, 4:120–143, 06 2012. doi: 10.1109/TCIAIG.2012.2200894.
- Zohar Feldman and Carmel Domshlak. Monte-carlo tree search: To MC or to dp? In Torsten Schaub, Gerhard Friedrich, and Barry O’Sullivan, editors, *ECAI 2014 - 21st European Conference on Artificial Intelligence, 18-22 August 2014, Prague, Czech Republic - Including Prestigious Applications of Intelligent Systems (PAIS 2014)*, volume 263 of *Frontiers in Artificial Intelligence and Applications*, pages 321–326. IOS Press, 2014. doi: 10.3233/978-1-61499-419-0-321. URL <https://doi.org/10.3233/978-1-61499-419-0-321>.
- Florian Geißer and David Speck. Prost-dd-utilizing symbolic classical planning in thts. 2018.
- Mathias Guenther and Jobas Klinik. A new experience: O-thell-us – an ai project. <https://courses.cs.washington.edu/courses/cse573/04au/Project/mini1/O-Thell-Us/0thellus.pdf>, 2004.
- Steven James, George Konidaris, and Benjamin Rosman. An analysis of monte carlo tree search, 2017. URL <https://aaai.org/ocs/index.php/AAAI/AAAI17/paper/view/14886>.
- Levente Kocsis and Csaba Szepesvári. Bandit based monte-carlo planning. In Johannes Fürnkranz, Tobias Scheffer, and Myra Spiliopoulou, editors, *Machine Learning: ECML 2006*, pages 282–293, Berlin, Heidelberg, 2006. Springer Berlin Heidelberg. ISBN 978-3-540-46056-5.
- Iris Kohler, Theresa Migler, and Foaad Khosmood. Composition of basic heuristics for the game 2048. In *Proceedings of the 14th International Conference on the Foundations of Digital Games*, FDG ’19, New York, NY, USA, 2019. Association for Computing Machinery. ISBN 9781450372176. doi: 10.1145/3337722.3341838. URL <https://doi.org/10.1145/3337722.3341838>.
- Larkin Liu. Connect 4 implementation in scala. <https://github.com/larkz/connect4-scala>, commit = 4549c00398e7987814d83a2bd0760bbaedeb879b, 2021.
- Simon Lucas. Montecarlots - java. <https://github.com/DrumerJoe21/MonteCarloTS>, 2018.

- Peter Norvig. *Paradigms of Artificial Intelligence Programming: Case Studies in Common Lisp*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1st edition, 1992. ISBN 1558601910.
- Martin L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, Inc., USA, 1st edition, 1994. ISBN 0471619779.
- Scott Sanner. Relational dynamic influence diagram language (rddl): Language description, 2010.
- David Silver, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George van den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, Sander Dieleman, Dominik Grewe, John Nham, Nal Kalchbrenner, Ilya Sutskever, Timothy Lillicrap, Madeleine Leach, Koray Kavukcuoglu, Thore Graepel, and Demis Hassabis. Mastering the game of Go with deep neural networks and tree search. *Nature*, 529(7587):484–489, January 2016. doi: 10.1038/nature16961.
- M. H. M. Winands, Y. Bjornsson, and J. Saito. Monte carlo tree search in lines of action. *IEEE Transactions on Computational Intelligence and AI in Games*, 2(4):239–250, 2010. doi: 10.1109/TCIAIG.2010.2061050.