| Introduction | Search Trees | Blind | Heuristic Search | MCTS | Summary | Quiz | References |
|--------------|--------------|-------|------------------|------|---------|------|------------|
| | | | | | | | |

Elements of DSAI AlphaGo Part 1: Introduction, Single-Agent Search Search & Learn: A Recipe for AI Action Decisions

Jörg Hoffmann



COMPUTER SCIENCE

Winter Term 2019/20

Elements of DSAI



Why AI Game Playing?

- Playing a game well is commonly viewed to require a form of "intelligence".
- Games capture a pure form of competition between opponents.
- Games are abstract and precisely defined, thus very easy to formalize and implement in a computer.

 \rightarrow Game playing is one of the oldest topics in AI.

 \rightarrow The dream of a machine that plays Chess is much older than AI!

Number of possible board positions: (let's do some combinatorics)



- Choose position for white king: 64
- $\bullet\,$ Choose position for white queen: 63
- 2 positions for bishops (Läufer): $\binom{62}{2}$
- Knights (Springer) $\binom{60}{2}$, rooks (Türme) $\binom{58}{2}$
- 8 positions for pawns (Bauern): $\binom{56}{8}$
- Similar for black pieces on 48 remaining squares.

 $\rightarrow 4.6*10^{42}.$ But that's just those positions where all pieces are present!

Estimated number of legal positions: ca. 10^{43} (estimated number of atoms in the universe: 10^{82}).

Introduction Search Trees Blind Heuristic Search MCTS Summary Quiz References Why is Game_Playing Hard?

Number of **possible** board positions:



- 3 possibilities per square: white, black, empty
- 19 * 19 = 361 squares

 $\rightarrow 3^{361} \approx 1.7 * 10^{172}$. (Reminder: 10^{82} atoms in the universe.)

Number of legal positions: ca. 10¹⁷¹ (actually it's 20816819938197998469947863334486277028652245388453054842563945682092741 961273801537852564845169851964390725991601562812854608988831442712971531 9317557736620397247064840935, see https://tromp.github.io/go/legal.html)

Hoffmann

Elements of DSAI

Introduction Search Trees Blind Heuristic Search OCODO Quiz References OCODO QUIZ REFERE

- Checkers (Dame): Since 1994, "Chinook" is world champion. In 2007, it was shown to be *unbeatable*: Checkers is solved, we know the optimal game strategy.
- Chess: In 1997, "Deep Blue" beat Garry Kasparov.



- 6 games, final score 3.5 : 2.5.
- Specialized Chess hardware, 30 nodes with 16 processors each.
- Nowadays, standard PC hardware plays at world champion level.
- Go: Best computer players ("Zen", "Mogo", "Crazystone") at the level of good amateurs, using sampling-based search methods.

Introduction Search Trees Blind Heuristic Search MCTS Summary Quiz References

Computer Chess: Famous Quotes

Claude Shannon (1949)

The chess machine is an ideal one to start with, since

- the problem is sharply defined both in allowed operations (the moves) and in the ultimate goal (checkmate),
- it is neither so simple as to be trivial nor too difficult for satisfactory solution,
- One of the second se
- the discrete structure of chess fits well into the digital nature of modern computers.

Alexander Kronrod (1965)

Chess is the drosophila of Artificial Intelligence.

 \rightarrow For those not in the know: "drosophila" = fruit fly, a species heavily used for research in genetics.

Hoffmann

Elements of DSAI

Computer Chess: Another Famous Quote

John McCarthy (1997)

In 1965, the Russian mathematician Alexander Kronrod said, "Chess is the Drosophila of artificial intelligence."

However, computer chess has developed much as genetics might have if the geneticists had concentrated their efforts on breeding racing Drosophilae. We would have some science, but mainly we would have very fast fruit flies.

Introduction Search Trees Blind Heuristic Search MCTS Summary Quiz References Very Fast Fruit Flies, 1 B.A. B.A.<

Computer Chess: (mostly similar in other games)

- Exhaustive search (Alpha-Beta, see later).
- Fast game-position evaluation functions, fine-tuned by human experts and training.
- Large database of known game positions for drawing analogies (from 2 million games, in 1997).
- Very large game opening databases.
- Very large game termination databases.
- Fast hardware.

 \rightarrow A mixture of (a) very fast search and human expertise. Learning plays only a minor role!

State of the Art, 3 A.A. (Anno AlphaGo)

AlphaGo, March 2016:

https://www.netflix.com/de-en/title/80190844

- AlphaGo beats Lee Sedol (winner of 18 world titles).
- Sampling-based search with guidance information from neural networks (NN), trained by expert data and self-play.

AlphaGo Zero, early 2017:

https://deepmind.com/blog/alphago-zero-learning-scratch/

- AlphaGo Zero beats AlphaGo using NN trained without expert data.
- Attention: Training time in "days" requires Google computing power! Intensely computation-expensive but massively parallelizable.

AlphaZero, late 2017:

 $\tt https://deepmind.com/blog/alphazero-shedding-new-light-grand-games-chess-shogi-and-go/$

AlphaZero beats world-class computer players in Go, chess, and shogi.



- Search Trees: What do we mean by "search" here?
 → We look at an example to illustrate the concept.
- Blind Search: How to explore the search tree?
 → Simple algorithms first (we'll be quick on this).
- Heuristic Search: How to explore the search tree *intelligently*?
 → The simple algorithms usually don't work. We need to inform the search about where it is promising to explore. Here's how.
- Monte-Carlo Tree Search (MCTS): Can we explore the search tree more sparsely?

 \rightarrow Sampling into the bargain. Recent successes, including AlphaGo.

Introduction 00000000 Search Trees

Heuristic Search

earch N 00000 0

MCTS

Summary Quiz

References

Positioning in the DSAI Phase Model

Blind



Introduction Search Trees Blind Heuristic Search MCTS Summary Quiz References

A Single-Agent Search Problem

\rightarrow Problem: Find a route to Moscow.



- Starting from an initial state ... (SB)
- ... apply actions ... (Using a road segment)
- ... to reach a goal state. (Moscow)
- Performance measure: Minimize summed-up action costs. (Road segment lengths)

Hoffmann

Elements of DSAI

Another Single-Agent Search Problem (The "15-Puzzle")

\rightarrow Problem: Move tiles to transform left state into right state.

| 9 | 2 | 12 | 6 |
|----|----|----|----|
| 5 | 7 | 14 | 13 |
| 3 | 4 | 1 | 11 |
| 15 | 10 | 8 | |

| 1 | 2 | 3 | 4 |
|----|----|----|----|
| 5 | 6 | 7 | 8 |
| 9 | 10 | 11 | 12 |
| 13 | 14 | 15 | |

- Starting from an initial state ... (Left)
- ... apply actions ... (Moving a tile)
- ... to reach a goal state. (Right)
- Performance measure: Minimize summed-up action costs. (Each move has cost 1, so we minimize the number of moves)

Hoffmann

Introduction Search Trees Blind Heuristic Search OCTS Summary Quiz References OCO Single-Agent Search Problems

Restricted environment: (yet practical applications, see next slide)

- Single-agent.
- Finite numbers of states and actions (in particular: discrete).
- Fully observable (agent knows everything).
- Deterministic (each action has only one outcome).
- Static (if the agent does nothing, the world doesn't change).

 \rightarrow All of these restrictions can be removed. For example, many of the basic ideas generalize easily to two-player zero-sum games.

Introduction Search Trees Blind Heuristic Search OCTS Summary Quiz References OCO Single-Agent Search Applications

Just to name a few:

- Route planning (e.g. Google Maps).
- Detecting bugs in software and hardware.
- Non-player-characters in computer games.
- Robot assembly sequencing. Planning of the assembly of complex objects. Actions = robot activities.
- Attack planning. Finding a hack into a secured network. Used for regular security testing. Actions = exploits.
- Query optimization in databases. Actions = rewriting operations.
- Sequence alignment in Bioinformatics. Actions = re-alignment operations.
- Natural language sentence generation. Actions = add another word to a partial sentence.



Start at the initial state. Then, step-by-step, expand a state by generating its successors $\ldots \rightarrow$ Search space.





An Example Search Space: The "8-Puzzle"



etc. ... (Blackboard)



Breadth-First Search: Expand nodes in the order they were generated.



Depth-First Search: Always expand the most recently generated node.



| Introduction 000000000 | Search Trees 000000 | Blind ⊙● | Heuristic Search | MCTS 00000 | Summary 00 | Quiz 000 | References |
|---------------------------|------------------------|-------------|------------------|---------------|---------------|-------------|------------|
| Algorith | m Prope | rties | | | | | |

Breadth-First Search:

- Does this guarantee to find a solution if there is one? Yes.
- Does this provide a guarantee on solution quality? Always finds a shortest solution.

Depth-First Search:

- Does this guarantee to find a solution if there is one? No, because search branches may be infinitely long (cycles).
- Does this provide a guarantee on solution quality? Trivially no.

 \rightarrow There are many more algorithms improving over these in a variety of ways. See, for example, [Russell and Norvig (2010)].

Hoffmann

Elements of DSAI

Introduction Search Trees Blind Heuristic Search MCTS Summary Quiz References

(Not) Playing Stupid

\rightarrow Problem: Find a route to Moscow.



- "Look at all locations 10km distant from SB, look at all locations 20km distant from SB," = Breadth-first search.
- "Keep following down arbitrary roads until you hit an ocean, then back up" = Depth-first search.
- "Focus on roads that go the right direction." = Heuristic search!



Heuristic Search: Basic Idea



 \rightarrow Heuristic function h estimates the cost of an optimal path from a state s to the goal; search prefers to expand states s with small h(s).

Hoffmann

Elements of DSAI

| Introduction 000000000 | Search Trees 000000 | Blind 00 | Heuristic Search | MCTS 00000 | Summary 00 | Quiz 000 | References |
|---------------------------|------------------------|-------------|------------------|---------------|---------------|-------------|------------|
| | | | | | | | |

Some Applications

GPS



Video Games



Robotics



Security Testing



Introduction, Single-Agent Search

Definition: Given a single-agent search problem with states S, a heuristic is a function $h: S \mapsto \mathbb{R}_0^+$.

Intended use:

• h estimates the cost of a cheapest path from s to a goal state.

Terminology:

- Ancient Greek $\varepsilon \upsilon \rho \iota \sigma \kappa \varepsilon \iota \nu$ (= "I find"); aka: $\varepsilon \upsilon \rho \eta \kappa \alpha$!
- Same word often used in CS for "rule of thumb", "imprecise solution method".

How to obtain h?

- Goal distance within a relaxed (i.e., simplified) problem.
- Or: learn from data, see later (games).

Hoffmann

Elements of DSAI

Introduction Search Trees Blind over the search of the sea

Heuristic Function from Relaxed Problem: Example



Problem: Find a route from Saarbrücken To Edinburgh.

Hoffmann

Elements of DSAI



Edinburgh



Relaxed Problem: Throw away the map.

Hoffmann

Elements of DSAI





Heuristic function h: Straight line distance.

Hoffmann

Elements of DSAI

| 1 6 | | DI.I. | A | | | | |
|--------------|--------------|-------|------------------|-------|---------|------|------------|
| 000000000 | 000000 | 00 | 00000000000 | 00000 | 00 | 000 | |
| Introduction | Search Trees | Blind | Heuristic Search | MCTS | Summary | Quiz | References |

h from Relaxed Problem: Another Example

| 9 | 2 | 12 | 6 | 1 | 2 | 3 | 4 |
|----|----|----|----|-------|----|----|----|
| 5 | 7 | 14 | 13 | 5 | 6 | 7 | 8 |
| 3 | 4 | 1 | 11 | 9 | 10 | 11 | 12 |
| 15 | 10 | 8 | | 13 | 14 | 15 | |

- Problem: Move tiles to transform left state into right state.
- Relaxed Problem: Allow to move each tile to *any cell* (in a single move), regardless of the situation.
- Heuristic function h: Number of misplaced tiles. (Here: 13)

| 1 6 | | DI.I. | A | | | | |
|--------------|--------------|-------|------------------|-------|---------|------|------------|
| 000000000 | 000000 | 00 | 00000000000 | 00000 | 00 | 000 | |
| Introduction | Search Trees | Blind | Heuristic Search | MCTS | Summary | Quiz | References |

h from Relaxed Problem: Another Example

| 9 | 2 | 12 | 6 | 1 | 2 | 3 | 4 |
|----|----|----|----|-------|----|----|----|
| 5 | 7 | 14 | 13 | 5 | 6 | 7 | 8 |
| 3 | 4 | 1 | 11 | 9 | 10 | 11 | 12 |
| 15 | 10 | 8 | | 13 | 14 | 15 | |

- Problem: Move tiles to transform left state into right state.
- Relaxed Problem: Allow to move each tile to *any neighbor cell*, regardless of the situation.
- Heuristic function h: Manhattan distance. (Here: 36)

Introduction Search Trees Blind Heuristic Search MCTS Summary Quiz References Heuristic Functions: Illustration Path Planning Vision Vision

Manhattan Distance, good case:



Introduction Search Trees Blind Heuristic Search MCTS Summary Quiz References Heuristic Functions: Illustration Path Planning Vision Vision

Manhattan Distance, bad case:



Introduction Search Trees Blind Heuristic Search MCTS Summary Quiz References

Search nodes ordered by ascending *h*:

```
function Greedy Best-First Search(problem) returns a solution, or failure

node \leftarrow a node n with n.State=problem.InitialState

frontier \leftarrow a priority queue ordered by ascending h, only element n

loop do

if Empty?(frontier) then return failure

n \leftarrow Pop(frontier)

if problem.GoalTest(n.State) then return Solution(n)

for each action a in problem.Actions(n.State) do

n' \leftarrow ChildNode(problem, n, a)

Insert(n', h(n'), frontier)
```

• Typically used: duplicates pruning. Insert n' only if n.State has not been seen before.

| C | | | | . р . | 1 . | | |
|--------------|--------------|-------|------------------|--------------|---------|------|------------|
| 000000000 | 000000 | 00 | 00000000000000 | 00000 | 00 | 000 | |
| Introduction | Search Trees | Blind | Heuristic Search | MCTS | Summary | Quiz | References |

Greedy Best-First Search: Route to Bucharest





Subscripts: *h*. Red nodes: removed by duplicate pruning.



Hoffmann

Elements of DSAI

| Introduction 000000000 | Search Trees | Blind 00 | Heuristic Search 00000000000000 | MCTS 00000 | Summary 00 | Quiz 000 | References |
|---------------------------|--------------|-------------|------------------------------------|---------------|---------------|-------------|------------|
| A^* | | | | | | | |

Ordered by ascending g + h: (g(n) = cost from initial state to n)

function A* (problem) returns a solution, or failure $node \leftarrow a node n$ with n.State=problem.InitialStatefrontier \leftarrow a priority queue ordered by ascending g + h, only element nloop do if Empty?(frontier) then return failure $n \leftarrow Pop(frontier)$ if problem.GoalTest(n.State) then return Solution(n) for each action a in problem.Actions(n.State) do $n' \leftarrow ChildNode(problem, n, a)$ Insert(n', g(n') + h(n'), frontier)

- Optimal if *h* is admissible, i.e. never over-estimates the true cost-to-goal.
- Duplicates pruning: possible but must take care to not affect optimality.

Introduction Search Trees Blind Heuristic Search coopooloooloo MCTS Summary coopooloooloo Quiz References Performance: A* in the 8-Puzzle

Without Duplicate Elimination; d =length of solution:

| | Numl | per of search nodes | generated |
|----|------------------|---------------------|------------------------|
| | Iterative | А | * with |
| d | Deepening Search | misplaced tiles h | Manhattan distance h |
| 2 | 10 | 6 | 6 |
| 4 | 112 | 13 | 12 |
| 6 | 680 | 20 | 18 |
| 8 | 6384 | 39 | 25 |
| 10 | 47127 | 93 | 39 |
| 12 | 3644035 | 227 | 73 |
| 14 | - | 539 | 113 |
| 16 | - | 1301 | 211 |
| 18 | - | 3056 | 363 |
| 20 | - | 7276 | 676 |
| 22 | - | 18094 | 1219 |
| 24 | - | 39135 | 1641 |

 Introduction
 Search Trees
 Blind
 Heuristic Search
 MCTS
 Summary
 Quiz
 References

 Empirical Performance:
 A* in Path Planning



Live illustration:

http://qiao.github.io/PathFinding.js/visual/

Hoffmann

Elements of DSAI

Introduction Search Trees Blind oo Heuristic Search October Summary Quiz References

Setting: Online, fixed time budget. Decide which action to take.

Algorithm: Current state s, choose action a

```
while time not up do
select a transition s \xrightarrow{a} s'
rollout from s' until terminal state t
increment \#expansions(a), update average reward(a) with reward(t)
return an a for s with maximal average reward(a)
```

Notes:

- Sparse search, rollouts go deep fast.
- Guidance information can be injected into "select" and "rollout".
- Rollouts require "terminal state" to be reached in finite number of transitions. Intrinsic in games; not as clear in general.
- "Reward" at terminal state: intrinsic in games; here, based on path success (goal reached?) and, in the successful cases, plan cost.

Introduction Search Trees Blind Heuristic Search OCON Quiz References

Monte-Carlo Sampling: Illustration



Introduction Search Trees Blind Heuristic Search OCO Quiz References

Monte-Carlo Sampling: Illustration



Introduction Search Trees Blind Heuristic Search OCON Quiz References

Monte-Carlo Sampling: Illustration



 Introduction
 Search
 Trees
 Blind
 Heuristic Search
 MCTS
 Summary
 Quiz
 References

 000000000
 00000000000
 00000000000
 00000000000
 00000000000
 00000000000
 00000000000

Monte-Carlo Sampling: Illustration



Introduction Search Trees Blind Heuristic Search OCON Quiz References

Monte-Carlo Sampling: Illustration



Expansions: 0, 0 avg. reward: 0, 0

Introduction Search Trees Blind Heuristic Search MCTS Summary Quiz References

Monte-Carlo Tree Search (MCTS)

Algorithm: Same, but maintain a search tree \boldsymbol{T}

```
while time not up do
select actions within T up to a state s' and s' \xrightarrow{a'} s'' s.t. s'' \notin T
with bias to maximize reward
rollout from s'' until terminal state t
add s'' to T
update, from a' up to root, #expansions and average rewards
return an a for s with maximal average reward(a)
When executing a, keep the part of T below a
```

Notes:

- Search sparsity, information from deep samples, requirement terminal states and reward: as in Monte-Carlo Sampling.
- Action decisions (select) within tree. Maximize reward: exploitation; only bias, not exclusive: exploration, convergence to optimal choices.
- Different strategies/guidance within tree vs. rollouts.
- Tree remembers previous work, reduces redundant work later.

Hoffmann

Elements of DSAI

| 000000000 | 000000 | 00 | 00000000000 | 00000 | 00 | 000 | | | |
|--------------------|--------|----|-------------|-------|----|-----|--|--|--|
| MCTS: Illustration | | | | | | | | | |







| Introduction 000000000 | Search Trees | Blind 00 | Heuristic Search 000000000000 | MCTS 000€0 | Summary 00 | Quiz 000 | References | | |
|---------------------------|--------------|-------------|----------------------------------|---------------|---------------|-------------|------------|--|--|
| MCTS: Illustration | | | | | | | | | |







| NACTO | 111 | | | | | | |
|--------------|--------------|-------|------------------|-------|---------|------|------------|
| | | | | 00000 | | | |
| Introduction | Search Trees | Blind | Heuristic Search | MCTS | Summary | Quiz | References |

MCTS: Illustration





MCTS: Illustration



| 000000000 | 000000 | 00 | 000000000000 | 00000 | 00 | 000 | | | | |
|--------------|--------------|-------|------------------|-------|---------|------|------------|--|--|--|
| Introduction | Search Trees | Blind | Heuristic Search | MCTS | Summary | Quiz | References | | | |

MCTS: Illustration





Basic strategy: Exploitation vs. exploration

- Exploitation: Actions with high average \rightarrow Focus on "what works".
- Exploration: Actions not tried a lot yet \rightarrow Try to find better solutions.
- Balance the two, e.g. UCT [Kocsis and Szepesvári (2006)].

 \rightarrow Strong link to reinforcement learning: exploitation = "use the action currently believed to be best"; exploration = "use some other action".

Additional information: Incorporate search guidance!

- Heuristic function: State quality estimate to guide "select" and/or "rollout".
- Approximate policy (action-choice function): Same, as action-selection bias.
- Based on problem relaxations (→ no training phase needed), or on learning (→ reinforcement learning across search iterations).

| Introduction | Search Trees | Blind | Heuristic Search | MCTS | Summary | Quiz | References | |
|--------------|--------------|-------|------------------|-------|---------|------|------------|--|
| 000000000 | 000000 | 00 | 000000000000 | 00000 | ●0 | 000 | | |
| Summary | | | | | | | | |

- Game playing is a traditional benchmark for AI. The AlphaGo/Zero line of works has dramatically advanced the state of the art, but there are many challenges still.
- The success of AlphaGo is based on a combination of search with learned search guidance.
- Single-agent search problems require to find a path of actions leading from an initial state to a goal state.
- Heuristic functions *h* map states to an estimate of goal distance. Greedy best-first search and A^{*} use *h* for ordering search nodes.
- Monte-Carlo tree search (MCTS) is a sparse form of search, going deep through rollouts. One needs to balance exploitation vs. exploration, similarly as in reinforcement learning. Heuristic functions and approximate policies can be integrated/learned.

| Introduction | Search Trees | Blind | Heuristic Search | MCTS | Summary | Quiz | References |
|--------------|--------------|-------|------------------|-------|---------|------|------------|
| 000000000 | 000000 | 00 | 000000000000 | 00000 | ○● | 000 | |
| Reading | | | | | | | |

 Chapter 3: Solving Problems by Searching, Sections 3.1 – 3.4 [Russell and Norvig (2010)].

Content: Covers search trees and blind search in more detail.

• Chapter 3: Solving Problems by Searching, Sections 3.5 and 3.6 [Russell and Norvig (2010)].

Content: Covers heuristic search in more detail.

https://en.wikipedia.org/wiki/Monte_Carlo_tree_search
 Content: Quite a nice read giving lots of background. (MCTS is not

covered by Russel/Norvig yet)

• *Trial-based Heuristic Tree Search for Finite Horizon MDPs* [Russell and Norvig (2010)].

Content: Nice overview paper systematizing the space of MCTS-style algorithms in (probabilistic) AI Planning.

Hoffmann

| Introduction 000000000 | Search Trees 000000 | Blind 00 | Heuristic Search 000000000000 | MCTS 00000 | Summary 00 | Quiz ●00 | References | | |
|---|------------------------|-------------|----------------------------------|---------------|---------------|-------------|------------|--|--|
| Quiz: V | ∕ideo Gan | nes | | | | | | | |
| Quiz | | | | | | | | | |
| When was the first game-playing computer built? | | | | | | | | | |
| (A): 19 | 41 | | (B) | : 1950 | | | | | |

(C): 1958 (D): 1965

 \rightarrow In 1941, a small box beat humans at Nim (take away objects from heaps, player taking the last object looses).

Quiz

Does the video game industry attempt to make the computer opponents as intelligent as possible?

(A): Yes (B): No

 \rightarrow In some cases, yes (I suppose). In general, no. For example, in Ego-Shooter games, if your computer opponents did the best they can, you'd be shot immediately and always (similar: Rambo movies et al).

Hoffmann

Elements of DSAI

Introduction 000000000 Search Trees

Heuristic Search

arch M

S Sun

ary Quiz o●o References

Quiz: Gorillas and Koalas



Blind

Quiz

How to solve this?

1. 1K,1G; 2. 1K; 3. 2K; 4. 1K; 5. 1K,1G.



Quiz

And this?

Video: http://www.youtube. com/watch?v=W9NEWxabGmg

Quiz: Heuristic Error in the Blocksworld



- n blocks, 1 hand.
- A single action either takes a block with the hand or puts a block we're holding onto some other block/the table.
- The goal is a set of statements "on(x,y)".

Quiz

Consider h := number of goal statements that are not currently true. Is the error (difference from true goal distance) bounded by a constant?

(A): Yes.

(B): No.

 \rightarrow No. There are examples where the error grows linearly in n. Example: Block b_1 is currently beneath a stack of b_n, \ldots, b_2 and the goal is on (b_1, b_2) . Then h(s) = 1 but true goal distance is n as we have to move b_n, \ldots, b_2 away first.

Hoffmann

Elements of DSAI

| Introduction 000000000 | Search Trees 000000 | Blind 00 | Heuristic Search | MCTS 00000 | Summary 00 | Quiz 000 | References | |
|---------------------------|------------------------|-------------|------------------|---------------|---------------|-------------|------------|--|
| References I | | | | | | | | |

- Levente Kocsis and Csaba Szepesvári. Bandit based Monte-Carlo planning. In Proceedings of the 17th European Conference on Machine Learning (ECML 2006), volume 4212 of Lecture Notes in Computer Science, pages 282–293, 2006.
- Stuart Russell and Peter Norvig. Artificial Intelligence: A Modern Approach (Third Edition). 2010.