



Foundations of Arti May 20, 2019 — 43. Monte-0	ficial Intelligence Carlo Tree Search: Introduction		
43.1 Introduction	1		
43.2 Monte-Carlo	o Methods		
43.3 Monte-Carlo	o Tree Search		
43.4 Summary			
M. Helmert (University of Basel)	Foundations of Artificial Intelligence	May 20, 2019	2 / 27



### Monte-Carlo Tree Search: Brief History

- Starting in the 1930s: first researchers experiment with Monte-Carlo methods
- ▶ 1998: Ginsberg's GIB player achieves strong performance playing Bridge ~> this chapter
- ▶ 2002: Auer et al. present UCB1 action selection for multi-armed bandits ~> Chapter 44
- 2006: Coulom coins the term Monte-Carlo Tree Search (MCTS) ~> this chapter
- 2006: Kocsis and Szepesvári combine UCB1 and MCTS into the most famous MCTS variant, UCT ~> Chapter 44

M. Helmert (University of Basel)

Foundations of Artificial Intelligence

May 20, 2019 5 / 27

Introduction

43. Monte-Carlo Tree Search: Introduction

Monte-Carlo Methods

# 43.2 Monte-Carlo Methods

### Monte-Carlo Tree Search: Applications

Examples for successful applications of MCTS in games:

- ▶ board games (e.g., Go ~→ Chapter 45)
- card games (e.g., Poker)
- ► AI for computer games (e.g., for Real-Time Strategy Games or Civilization)
- Story Generation (e.g., for dynamic dialogue generation in computer games)
- General Game Playing

Also many applications in other areas, e.g.,

- MDPs (planning with stochastic effects) or
- POMDPs (MDPs with partial observability)

M. Helmert (University of Basel)

Foundations of Artificial Intelligence

May 20, 2019 6 / 27

Introduction





May 20, 2019

11 / 27

Monte-Carlo Methods

43. Monte-Carlo Tree Search: Introduction

South to play, three tricks to win, trump suit ♣

Foundations of Artificial Intelligence

M. Helmert (University of Basel)

#### 43. Monte-Carlo Tree Search: Introduction

#### Monte-Carlo Methods: Example

Bridge Player GIB, based on Hindsight Optimization (HOP)

- perform samples as long as resources (deliberation time,
- **sample** hands for all players that are consistent with current knowledge about the game state
- ▶ for each legal move, compute if fully observable game that starts with executing that move is won or lost
- compute win percentage for each move over all samples
- play the card with the highest win percentage

Foundations of Artificial Intelligence

May 20, 2019 10 / 27





43. Monte-Carlo Tree Search: Introduction
Hindsight Optimization: Restrictions
HOP well-suited for partially observable games like most card games (Bridge, Skat, Klondike Solitaire)
must be possible to solve or approximate sampled game efficiently
often not optimal even if provided with infinite resources





May 20, 2019

Monte-Carlo Tree Search

17 / 27

## 43.3 Monte-Carlo Tree Search

M. Helmert (University of Basel)

Foundations of Artificial Intelligence

43. Monte-Carlo Tree Search: Introduction

Monte-Carlo Tree Search: Iterations

Each iteration consists of four phases:

- selection: traverse the tree by applying tree policy
  - **b** Stop when reaching terminal node (in this case, set  $n_{child}$  to that node and  $p_{\star}$  to its position and skip next two phases)...
  - $\blacktriangleright$  ... or when reaching a node  $n_{\text{parent}}$  for which not all successors are part of the tree.
- **•** expansion: add a missing successor  $n_{child}$  of  $n_{parent}$  to the tree
- **is simulation**: apply default policy from  $n_{child}$ until a terminal position  $p_{\star}$  is reached
- **backpropagation**: for all nodes *n* on path from root to  $n_{child}$ :
  - $\blacktriangleright$  increase N(n) by 1
  - $\blacktriangleright$  update current average  $\hat{u}(n)$  based on  $u(p_*)$

#### Monte-Carlo Tree Search: Idea

#### Monte-Carlo Tree Search (MCTS) ideas:

- perform iterations as long as resources (deliberation time, memory) allow:
- **build a partial game tree**, where nodes *n* are annotated with
  - utility estimate  $\hat{u}(n)$
  - $\blacktriangleright$  visit counter N(n)
- ▶ initially, the tree contains only the root node
- each iteration adds one node to the tree.

After constructing the tree, play the move that leads to the child of the root with highest utility estimate (as in minimax/alpha-beta).

M. Helmert (University of Basel)

Foundations of Artificial Intelligence

May 20, 2019 18 / 27













Monte-Carlo Tree Search

### Monte-Carlo Tree Search: Pseudo-Code

if is terminal (n posi	tion):	
utility := u(n.p)	osition)	
else:		
$p := n.get_unv$	isited_successor()	
if p is none:		
n' := apply	y_tree_policy(n)	
utility := v	$v_{isit_node}(n')$	
else:		
$p_\star:=appl$	$y_default_policy_until_end(p)$	
utility := u	$u(p_{\star})$	
n.add_chil	$d_node(p, utility)$	
undate visit count a	and estimate $(n, utility)$	
return <i>utility</i>		

43. Monte-Carlo Tree Search: Introduction
Summary
Monte-Carlo methods compute averages over a number of random samples.
Simple Monte-Carlo methods like Hindsight Optimization perform well in some games, but are suboptimal even with unbounded resources.
Monte-Carlo Tree Search (MCTS) algorithms iteratively build a search tree, adding one node in each iteration.
MCTS is parameterized by a tree policy and a default policy.

