
This is an electronic reprint of the original article.
This reprint may differ from the original in pagination and typographic detail.

Kujanpää, Kalle; Kannala, Juho; Babadi, Amin; Ilin, Alexander; Zhao, Yi; Pajarinen, Joni
Continuous Monte Carlo Graph Search

Published in:
AAMAS '24: Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems

DOI:
[10.5555/3635637.3662960](https://doi.org/10.5555/3635637.3662960)

Published: 01/01/2024

Document Version
Publisher's PDF, also known as Version of record

Published under the following license:
CC BY

Please cite the original version:
Kujanpää, K., Kannala, J., Babadi, A., Ilin, A., Zhao, Y., & Pajarinen, J. (2024). Continuous Monte Carlo Graph Search. In *AAMAS '24: Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems* (pp. 1047-1056). ACM. <https://doi.org/10.5555/3635637.3662960>



Continuous Monte Carlo Graph Search

Kalle Kujanpää^{*†‡}

Aalto University
Espoo, Finland
kalle.kujanpaa@aalto.fi

Amin Babadi^{*‡}

Bugbear Entertainment,
Aalto University
Helsinki, Finland
amin.babadi@bugbear.fi

Yi Zhao[§]

Aalto University
Espoo, Finland
yi.zhao@aalto.fi

Juho Kannala^{*†‡}

Aalto University
Espoo, Finland
juho.kannala@aalto.fi

Alexander Ilin^{*†‡}

Aalto University, System 2 AI
Espoo, Finland
alexander.ilin@aalto.fi

Joni Pajarinen^{*‡§}

Aalto University
Espoo, Finland
joni.pajarinen@aalto.fi

ABSTRACT

Online planning is crucial for high performance in many complex sequential decision-making tasks. Monte Carlo Tree Search (MCTS) employs a principled mechanism for trading off exploration for exploitation for efficient online planning, and it outperforms comparison methods in many discrete decision-making domains such as Go, Chess, and Shogi. Subsequently, extensions of MCTS to continuous domains have been developed. However, the inherent high branching factor and the resulting explosion of the search tree size are limiting the existing methods. To address this problem, we propose Continuous Monte Carlo Graph Search (CMCGS), an extension of MCTS to online planning in environments with continuous state and action spaces. CMCGS takes advantage of the insight that, during planning, sharing the same action policy between several states can yield high performance. To implement this idea, at each time step, CMCGS clusters similar states into a limited number of stochastic action bandit nodes, which produce a layered directed graph instead of an MCTS search tree. Experimental evaluation shows that CMCGS outperforms comparable planning methods in several complex continuous DeepMind Control Suite benchmarks and 2D navigation and exploration tasks with limited sample budgets. Furthermore, CMCGS can be scaled up through parallelization, and it outperforms the Cross-Entropy Method (CEM) in continuous control with learned dynamics models.

KEYWORDS

Continuous Control; Planning; Reinforcement Learning; Model-Based Reinforcement Learning; MCTS; Online Planning

ACM Reference Format:

Kalle Kujanpää, Amin Babadi, Yi Zhao, Juho Kannala, Alexander Ilin, and Joni Pajarinen. 2024. Continuous Monte Carlo Graph Search. In *Proc. of the 23rd*

^{*}Equal Contribution

[†]Department of Computer Science

[‡]Finnish Center for Artificial Intelligence (FCAI)

[§]Department of Electrical Engineering and Automation



This work is licensed under a Creative Commons Attribution International 4.0 License.

Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), N. Alechina, V. Dignum, M. Dastani, J.S. Sichman (eds.), May 6 – 10, 2024, Auckland, New Zealand. © 2024 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), Auckland, New Zealand, May 6 – 10, 2024, IFAAMAS, 10 pages.

1 INTRODUCTION

Monte Carlo Tree Search (MCTS) is a well-known online planning algorithm for solving decision-making problems in discrete action spaces [15, 16]. MCTS achieves super-human performance in various domains such as Atari, Go, Chess, and Shogi when a learned transition model is available [59]. Therefore, recent research has tried to extend MCTS to environments with continuous state and action spaces [26, 37, 39, 41, 55].

However, current MCTS approaches for continuous states and actions are limited [11, 14, 37, 55]. MCTS approaches that build the search tree by discretizing the action space or otherwise limiting the growth of the tree, such as progressive widening approaches [11, 14, 16], do not scale up well to high-dimensional action spaces and complex problems due to the search space increasing exponentially with the planning horizon and action dimensionality [9, 74]. To alleviate this problem, learning-based approaches [37, 55] reduce the required planning horizon when a sufficiently large number of training samples is available but do not solve the underlying problem of increasing search space size.

To solve the search space explosion problem in MCTS-based continuous planning, this paper presents Continuous Monte Carlo Graph Search (CMCGS), an extension of MCTS to the continuous control problem. Similarly to MCTS, CMCGS employs an iterative mechanism for building a search graph that yields the next action to execute. At each iteration, a series of operators is used to grow the search graph and update the information stored in the graph nodes. CMCGS uses state clustering and Gaussian action bandits to deal with the challenges posed by high-dimensional continuous states and actions and limit the search space explosion.

Our experiments show that CMCGS outperforms strong MCTS and other planning baselines in many challenging benchmarks given limited interaction with the environment, including high-dimensional environments from the DeepMind Control Suite (DMC) [68]. We also show that CMCGS can be efficiently parallelized to scale up and demonstrate its generality and robustness with learned dynamics models in high-dimensional visual tasks, where CMCGS outperforms the Cross-Entropy Method (CEM) [57, 71]. Moreover, CMCGS shows strong performance by being superior to CEM in environments that require complex exploration, and our algorithm




Figure 1: Core steps in one iteration of Continuous Monte Carlo Graph Search (CMCGS). a) Starting from the root node, the graph is navigated via action sampling and node selection until a sink node is reached. b) If there is enough experience collected in the final layer of the graph and the maximum depth has not been reached, a new layer containing a new node N is initialized. c) A trajectory of random actions is simulated from the graph’s sink node to approximate the value of the state. d) The computed accumulated reward is backed up through the selected nodes, updating their replay memories, policies, and state distributions. e) If a new cluster of experience data is found in a previous layer of the graph, all nodes in that layer are updated based on the new clustering information (in this example, the node N is split into two new nodes $N1$ and $N2$).

also performs well in difficult, high-dimensional environments with informative rewards, where MCTS-based methods fail.

2 RELATED WORK

Monte Carlo Tree Search (MCTS), a decision-time planning method in discrete action spaces [9, 15, 40], is a core component of the success of the computer Go game [22, 61, 62]. MCTS also shows advantages in general game playing [3, 20, 23, 24]. Several recent works combine MCTS with a learned dynamic model lifting the requirement for access to a simulator [59, 73]. Grouping similar or equivalent states reached via different trajectories in MCTS has been proposed [12], but recent methods mostly focus on the discrete state setting [17, 42], which differs from our continuous setup.

Online planning in continuous action spaces. The number of actions in continuous action spaces is infinite, which significantly increases the size of the MCTS search tree and makes applying out-of-the-box MCTS infeasible. Instead, CEM is commonly used in the continuous action domain as an online planning method [13, 25, 56, 57, 71]. Many methods combine CEM with a value function or policy in different ways to improve its performance [8, 30, 31, 44, 50]. The main limitation of CEM compared to MCTS (and this work) is that CEM models the whole action trajectory using a single sampling distribution. This could be translated into the context of MCTS as using only one node at each layer of the search tree, and it can limit the exploration capabilities of CEM in environments where there are several ways of controlling the agent, for example, going around an obstacle using two different ways. Covariance Matrix Adaptation Evolution Strategy (CMA-ES) is an evolutionary

gradient-free algorithm for black-box continuous optimization [28, 29]. Although it has not seen much use in continuous control due to its computational complexity and memory requirements, applying it as an alternative to CEM has been proposed [18, 32, 49].

MCTS in continuous action spaces Adopting MCTS in continuous action spaces requires mitigating the requirement of enumerating all actions as in the discrete case. Several works use progressive widening to increase the number of child actions of a node as a function of its visitation frequency [11, 14, 16, 46]. The action space can be split into bins by factorizing across action dimensions [27, 65]. A line of work extends MCTS to continuous action spaces via Hierarchical Optimistic Optimization (HOO) [10, 45, 48, 54]. HOO hierarchically partitions the action space by building a binary tree to incrementally split the action space into smaller subspaces. Voronoi Optimistic Optimization (VOO) builds on similar ideas as HOO. It uses Voronoi partitioning to sample actions efficiently and is guaranteed to converge [39]. VG-UCT relies on action value gradients to perform local policy improvement [41]. Furthermore, several works grow the search tree based on sampling. Kernel regression can be used to generalize the value estimation and, thus, limit the number of actions sampled per node [74]. However, our experiments indicate that despite the presented strategies, these methods still struggle with the rapidly growing search space, especially with high-dimensional states and actions. We propose tackling this challenge by representing the search space as a graph instead of a tree, clustering similar states into search nodes, and using Gaussian action bandits for the nodes. In practice, CMCGS significantly improves on the continuous MCTS methods, especially with limited search budgets.

State abstractions in MCTS Our approach to clustering similar states into stochastic action bandit nodes is closely related to state abstraction and aggregation methods. State abstraction helps mitigate the impact of the combinatorial explosion in planning [1]. In state aggregation, many states are combined into abstract ones [43, 69] and it has been proposed for handling the high branching factor in stochastic environments [35]. Hierarchical MCTS through action abstraction is closely related [6]. Progressive Abstraction Refinement in Sparse Sampling (PARSS) builds and refines state abstractions as the search progresses [36]. However, these methods, including PARSS, generally assume a discrete action space, incompatible with our setup. State abstractions can also be used in a continuous setting when combined with progressive widening [63]. Our results indicate that relying on progressive widening makes the proposed method uncompetitive with CMCGS. Elastic MCTS uses state abstractions for strategy game playing, but lacks a progressive widening or discretization strategy to handle the high-dimensional continuous action spaces prevalent in our experiments [72].

There are also learning-based approaches to online planning that represent the policy using a neural network and sample from it when expanding the tree [2, 37]. By leveraging the learned policy, these methods achieve promising performance with limited search budgets. However, we focus only on the core search component and leave comparisons with learning-based methods to future work.

3 PRELIMINARIES

Planning and search is arguably the most classic approach for optimal control [19, 47, 67]. The idea is to start with a (rough) estimation of the action trajectory and gradually improve it through model-based simulation of the environment. In the case of environments with complex dynamics, this is usually done online by optimizing only a small part of the action trajectory at each step. Then, the first action of the optimized trajectory is executed. After that, the whole process is repeated, starting from the newly visited state. This process is also known as closed-loop control, or model-predictive control (MPC). MPC is an effective approach for optimal control in complex real-time environments [5, 21, 58].

Monte Carlo Tree Search (MCTS) [15] is one of the most popular MPC algorithms in environments such as video games [34, 52] that require principled exploration. MCTS starts by initializing a single-node search tree using the current state s_t as the root and grows the tree in an iterative manner, where one node is added to the tree at each iteration. The iterative process of MCTS is comprised of the following four key steps [9]:

- (1) **Selection (Tree Policy):** In this step, the algorithm selects one of the tree nodes to be expanded next. This is done by starting from the root and navigating to a child node until a node with at least one unexpanded child is visited using a selection criterion, whose goal is to balance the exploration-exploitation trade-off.
- (2) **Expansion:** In this step, MCTS expands the search tree by adding a (randomly chosen) unvisited child to the selected node from the previous step.
- (3) **Simulation (Default Policy):** After expanding the selected node, the newly added node is evaluated using a trajectory of random actions and computing the return.

- (4) **Backup:** In the final step, the computed return is backed up through the navigated nodes, updating the statistics stored in each node.

For a more detailed description of the MCTS algorithm, the reader is referred to Browne et al. [9].

4 CONTINUOUS MONTE CARLO GRAPH SEARCH

We propose an algorithm for continuous control called Continuous Monte Carlo Graph Search (CMCGS). We assume a discrete-time, discounted, and fully-observable Markov Decision Process (MDP) $M = \langle \mathcal{S}, \mathcal{A}, P, R, \gamma \rangle$ with continuous states and actions. \mathcal{S} is the state space, \mathcal{A} the action space, P the transition probability function, R the reward function, and $\gamma \in [0, 1]$ the discount factor. The objective of CMCGS is to find the action $a \in \mathcal{A}$ that maximizes the return in M by performing closed-loop control with model-based simulation of the environment. CMCGS is an iterative online planning algorithm in which each iteration contains the same four steps as MCTS with an additional width expansion step (see Fig. 1). Instead of building a search tree, CMCGS performs a search on a layered graph in which each layer corresponds to one step of the underlying MDP. We assume that we have access to a dynamics model $p(s_{t+1}, r_{t+1} | s_t, a_t)$ that can either be an exact model of the environment or a learned approximation.

The core idea of CMCGS is to use state clustering, Gaussian action bandits, and a layered directed search graph to tackle the challenges posed by search state explosion and high-dimensional continuous states and actions. Each node q in the t -th layer of the CMCGS graph corresponds to a *cluster of states visited at timestep t* . The node is characterized by its state distribution $p_q(s_t)$ and policy $\pi_q(a_t)$, which describe the states assigned to this cluster and the preferred actions taken from this cluster during planning. Both distributions $p_q(s_t)$ and $\pi_q(a_t)$ are modeled as Gaussian distributions with diagonal covariance matrices. The policy $\pi_q(a_t)$ is initialized as a Gaussian $\mathcal{N}(\mu = \frac{1}{2}(a_{\min} + a_{\max}), \sigma = \frac{1}{2}(a_{\max} - a_{\min}))$. During planning, the distributions are fit to the data in the replay buffer. The replay buffer contains tuples $e = (s_t, a_t, s_{t+1}, R, q)$, where q is the node to which state s_t belongs and R is the observed trajectory return. We denote by \mathcal{D}_q all the tuples in the replay buffer that correspond to node q and by Q_t the set of nodes in layer t .

The search graph is initialized with d_{init} layers and one node per layer. Initializing the graph with multiple layers improves the early search efficiency. The CMCGS algorithm then iterates the following steps (see Fig. 1 for a graphical illustration).

- (a) **Selection:** The selection mechanism is repeatedly applied to navigate through the search graph starting from the root node until a sink node (a node without children) or a terminal state of the MDP is reached (Fig. 1a). When navigating through the graph, we use an epsilon-greedy-inspired policy: with probability ϵ , we sample an action from the node policy, the Gaussian distribution $\pi_q(a)$, and with probability $1 - \epsilon$, we uniformly sample one of the top actions in \mathcal{D}_q (greedy action) and modify it with a small Gaussian noise for local improvement¹:

¹This approach enables the agent to leverage past experience efficiently while exploring new options, leading to better performance on complex tasks.

$\epsilon_{\text{top}} \sim \mathcal{N}(0, \mathcal{E}_{\text{top}} \cdot (\mathbf{a}_{\text{max}} - \mathbf{a}_{\text{min}}))$. The current state \mathbf{s}_t is updated to \mathbf{s}_{t+1} according to the dynamics model $p(\mathbf{s}_{t+1}, r_{t+1} | \mathbf{s}_t, \mathbf{a}_t)$, and the node in the next layer is selected by maximizing the probability of the updated state according to the node state distributions:

$$q \leftarrow \arg \max_{q \in \mathcal{Q}_{t+1}} p_q(\mathbf{s}_{t+1})$$

- (b) **Depth expansion:** If the last layer of the graph has collected enough experience (the number of samples is greater than threshold m) and the maximum graph depth has not been reached, we add a new layer to the graph (Fig. 1b). The new layer contains a single node initially. Without the threshold m , the depth of the search graph would grow too fast.
- (c) **Simulation:** Similar to MCTS, we simulate a random trajectory using random actions starting from the sink node and compute the trajectory return R (Fig. 1c). The length of the rollout is controlled by a hyperparameter N_r , but the rollout is interrupted if a terminal state is encountered.
- (d) **Backup:** We traverse the search graph backward and update the node distributions $p_q(\mathbf{s})$ and $\pi_q(\mathbf{a})$ for all the nodes visited in the selection procedure (Fig. 1d). The replay buffers \mathcal{D}_q of the visited nodes are updated. For each affected node q , the Gaussian state distribution $p_q(\mathbf{s})$ is fitted to the states \mathbf{s} in the node replay buffer \mathcal{D}_q such that its mean is equal to the state mean and standard deviation to the state standard deviation. For the action distributions $\pi_q(\mathbf{a})$, we use a threshold-based strategy to encourage sufficient exploration: the node policy $\pi_q(\mathbf{a})$ is updated only if the replay buffer of the node q contains enough experience ($|\mathcal{D}_q| > m/2$). The dimension means of $\pi_q(\mathbf{a})$ are fitted to the highest-scoring (elite) experiences in \mathcal{D}_q , similarly to CEM. To update the dimension variances of $\pi_q(\mathbf{a})$, we adopt the Bayes' rule by defining a conjugate prior (inverse gamma) on the variances, compute the posteriors using the elite samples, and select the posterior means as the variances. This prevents the variance from decaying too fast. Formally, we assume that the dimension variances σ^2 of $\pi_q(\mathbf{a})$ follow the inverse gamma distribution with prior hyperparameters α and β . Given n elite samples $(\mathbf{a}_{t,1}, \dots, \mathbf{a}_{t,n})$ and the new dimension means μ of $\pi_q(\mathbf{a})$, the posterior hyperparameters of the inverse gamma distribution are:

$$\alpha_{\text{posterior}} = \alpha_{\text{prior}} + \frac{n}{2}, \quad \beta_{\text{posterior}} = \beta_{\text{prior}} + \frac{\sum_{i=1}^n (\mathbf{a}_{t,i} - \mu)^2}{2}$$

and the new dimension variances σ^2 of $\pi_q(\mathbf{a})$ are equal to the means of the inverse gamma posterior distributions:

$$\sigma^2 = \frac{\beta_{\text{posterior}}}{\alpha_{\text{posterior}} - 1}. \quad (1)$$

- (e) **Width expansion:** To prevent the rapid search graph growth and maintain efficiency, CMCGS computes the maximum desired number C_t of nodes in layer t using a heuristic rule

$$C_t = \min(n_{\text{max}}, \lfloor n_t/m \rfloor)$$

where n_t is the number of transitions collected at timestep t and stored in the replay buffer \mathcal{D} , m the threshold hyperparameter, and n_{max} the hard limit on the number of nodes. If $|\mathcal{Q}_t|$, the number of nodes in layer t , is smaller than C_t , we attempt to cluster the states visited at timestep t into $|\mathcal{Q}_t| + 1$ clusters.

We use Agglomerative Clustering with Ward linkage in our experiments due to its fast running time with the amount of data needed by CMCGS [51, 70]. Using Euclidean distance as a metric is feasible both in low-dimensional state spaces and learned latent spaces. The clustering is approved if each cluster contains at least $m/2$ states to prevent degenerate clusters that are unlikely to be selected for states encountered in the future. Then, a new node is added to the layer (Fig. 1e), all experience tuples in the replay buffer are re-assigned to the corresponding nodes and the state and action distributions $p_q(\mathbf{s})$ and $\pi_q(\mathbf{a})$ of the affected nodes are updated. Otherwise, no new node is created, and clustering is attempted again when the layer has collected $m/2$ new samples to limit the number of calls to the clustering algorithm and, therefore, speed up the algorithm.

When the simulation budget has been exhausted, the agent's next action is determined by the experience stored at the root node. If the search has access to a deterministic environment simulator, the first action of the trajectory with the highest return is chosen. If searching with a learned dynamics model or the environment is stochastic, the average of the top actions in the replay memory of the root node is returned to prevent the exploitation of model inaccuracies or individual samples. Each of the five steps of CMCGS can be performed efficiently using batch operations, making it possible to collect multiple trajectories in parallel. The pseudocode of the CMCGS algorithm is shown in Algorithm 1.

5 EXPERIMENTS

We run four sets of experiments to evaluate the performance, generality, and robustness of CMCGS. First, we present a motivating example where CEM fails due to its inability to explore properly in an environment where a multimodal representation of the action distribution is useful. Second, we evaluate CMCGS in continuous control tasks with limited sample budgets to demonstrate that it outperforms relevant baselines in a variety of tasks with different levels of exploration required. Third, we demonstrate that CMCGS can be parallelized and can successfully utilize learned models for large-scale planning in complex tasks by showing that it achieves significantly better planning performance than CEM in image-based continuous control environments. Finally, we perform extensive ablation studies to better understand the design choices of CMCGS.

5.1 Toy Environment

We design a toy environment where the agent samples N actions a_1, \dots, a_N sequentially and receives a basic reward of $r = 0.5$ if the absolute value of each action is greater than 1. An additional reward of 0.5 is given if the sign of all actions is the same. The reward is formalized as

$$r = \begin{cases} 0.5, & \text{if } \forall i \in \{1, \dots, N\} : |a_i| > 1 \\ 1, & \text{if } \forall i \in \{1, \dots, N\} : |a_i| > 1 \wedge \\ & \forall i \in \{2, \dots, N\} : \text{sgn}(a_i) = \text{sgn}(a_1) \end{cases}$$

All action bandits are initialized to $\mathcal{N}(0, 1)$ to eliminate the effect of the initial action distribution on the performance, the action space is equal to \mathbb{R} , and the state equal to the y-coordinate of the agent (see Fig. 2). Finding the basic reward of $r = 0.5$ is easy, but performing principled exploration to find the additional component

Algorithm 1 Continuous Monte Carlo Graph Search

```

1: function CMCGS( $s_0$ )
2:    $q_{\text{root}} \leftarrow \text{Node}(s_0, \text{parent}=\text{None})$ 
3:   Initialize search graph with  $d_{\text{init}}$  layers and  $q_{\text{root}}$  as the root.
4:   while within computational budget do
5:      $\tau, R, s_d \leftarrow \text{GRAPHPOLICY}(s_0, q_{\text{root}})$ 
6:     if  $s_d$  is not terminal then
7:        $R \leftarrow R + \text{RANDOMROLLOUT}(s_d)$ 
8:      $\text{BACKUP}(\tau, R)$ 
9:   return  $\text{GETBESTACTION}(q_{\text{root}})$ 
10: function GRAPHPOLICY( $s, q$ )
11:    $\tau \leftarrow \emptyset$ 
12:    $R \leftarrow 0$ 
13:    $d \leftarrow 1$ 
14:    $\phi \leftarrow \begin{cases} 1, & \text{if } \mathcal{U}(0, 1) < \epsilon \\ 0, & \text{otherwise} \end{cases}$ 
15:   while  $s$  is not terminal and  $d \leq N_{\text{layers}}$  do
16:      $a \sim \phi \pi_q(a) + (1 - \phi) \pi_{q, \text{greedy}}(a)$ 
17:     Apply action  $a$ , observe new state  $s'$ , reward  $r$ 
18:      $\tau \leftarrow \tau \cup \{(s, a, s', q)\}$ 
19:      $R \leftarrow R + r$ 
20:     if  $s'$  is not terminal and  $d = N_{\text{layers}}$  then
21:       Try adding a new layer with one node to the graph.
22:     if  $s'$  is terminal or  $d = N_{\text{layers}}$  then
23:       break
24:      $q \leftarrow \arg \max_{q \in Q_{t+1}} p_q(s')$ 
25:      $s \leftarrow s'$ 
26:      $d \leftarrow d + 1$ 
27:   return  $\tau, R, s$ 
28: function BACKUP( $\tau, R$ )
29:   for each  $\langle s_t, a_t, s_{t+1}, q_t \rangle \in \tau$  do
30:      $\text{BUFFERSTORE}(\langle s_t, a_t, R, s_{t+1}, q_t \rangle)$ 
31:     Compute  $C_t$ , the desired number of clusters in layer  $t$ 
32:     if  $|Q_t| < C_t$  then
33:       Cluster states  $\{s_t\}$  in layer  $t$  into  $|Q_t| + 1$  clusters
34:       if clustering was successful then
35:         Replace nodes in layer  $t$  with the new nodes
36:         for  $q \in Q_t$  do
37:           Estimate state distribution  $p_q(s)$ 
38:           Estimate policy  $\pi_q(a)$ 
39:     if no new cluster found in layer  $t$  then
40:       Update  $p_q(s)$  and  $\pi_q(a)$  of the visited node  $q_t$ 

```

is a challenge. Furthermore, being able to represent complex multimodal action distributions is a valuable asset in this environment. We let $N = 5$ and compare our method to random shooting and CEM. The results are plotted in Table 1 and the agents' behavior is illustrated in Fig. 2. The results show that CEM struggles to explore the environment properly, and CMCGS outperforms random shooting due to directed exploration. CEM sometimes fails to receive the basic reward of 0.5 due to the inability to represent the multimodal action distribution natural to this problem.

Table 1: The success rates of achieving the small and large rewards and the average rewards achieved by different algorithms in the toy environment plus-minus two standard errors.

Method	$r \geq 0.5$	$r = 1.0$	Average Reward
CMCGS	1.00	0.99	0.995 \pm 0.002
Random shooting	1.00	0.89	0.943 \pm 0.010
CEM	0.74	0.57	0.655 \pm 0.028




Figure 2: An illustration of exploration in the toy environment. For clarity, the actions have been clipped to be in the range $[-1, 1]$. The agent, the chosen action, the explored trajectories, and the different-sized rewards are represented by the blue dot, red arrow, grey dashed lines, and green dots, respectively. The search nodes of CMCGS and the corresponding state and action means are illustrated with black dots and arrows. The left image illustrates how CMCGS explores with state-dependent policies. In the other pictures, we see how CEM fails in the environment. CEM can either fail to discover the large reward and choose a suboptimal action (center) or completely fail to handle the multimodality required by the environment (right).

5.2 Continuous Control with Limited Interaction

Second, we compare CMCGS to the strong planning algorithms (1) Monte Carlo Tree Search with Progressive Widening [11, 14, 16, MCTS-PW], (2) Voronoi Optimistic Optimization Applied to Trees [39, VOOT], (3) MCTS with Iteratively Refined State Abstractions (MCTS-ABS) [63, MCTS-ABS], (4) Cross-Entropy Method [56, CEM], and (5) Covariance Matrix Adaptation Evolution Strategy [28, CMA-ES] to demonstrate the power of CMCGS in continuous control tasks when the number of environment interactions is a bottleneck. This can be the case in, for instance, real-life robotics or when a computationally very intensive and accurate simulator is used. In these tasks, we evaluate the planning performance of the algorithms on real dynamics. We also include standard random shooting (RS) as a baseline.

The environments we use in these experiments include:

- (1) **2D Navigation:** We developed an environment for the control of a particle that should reach the goal without colliding with the obstacles. We used two variations of this environment, one with circular and another with rectangular obstacles. In these environments, a new action is applied whenever the particle reaches a vertical gray line. The main challenge

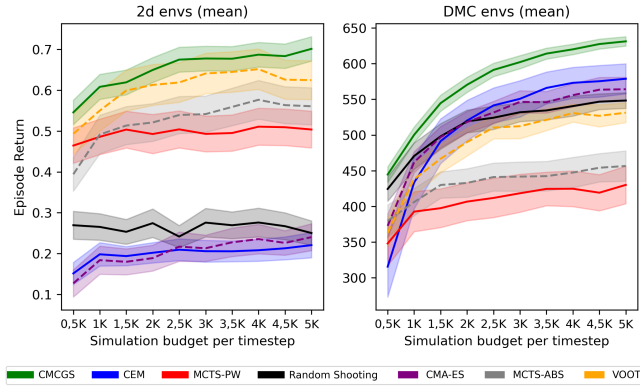


Figure 3: The mean episode returns averaged over the custom 2D environments that test exploration (left) and the environments from DeepMind Control Suite (right).

in these environments is that the agent has to explore to find the green reward and avoid colliding with the obstacles.

- (2) **2D Reacher:** Since *2D Navigation* environments only have 1 Degree of Freedom (DOF), we also developed a variation of the classic reacher task for a 2D multi-link arm. In this environment, only a sparse binary reward is used for reaching the goal that block the path. We use 15- and 30-linked arms to evaluate the scalability of CMCGS to high-dimensional action spaces.
- (3) **DeepMind Control Suite:** We also used several popular environments provided by DeepMind Control Suite [66] in our experiments. These environments pose a wide range of challenges for continuous control with low- and high-dimensional states and actions. In this part of the evaluation, we use proprioceptive observations.

We evaluate CMCGS and the chosen baselines using different simulation budgets per control timestep. For each budget, we used 400 random seeds for 2D navigation tasks, 200 for 2D reacher tasks, and 100 for DMC tasks. All experiments used the same set of hyperparameters that were carefully tuned to maximize the performance on the chosen environments.

We plot the means of the episode returns in Fig. 3 and the complete results in Fig. 4. The plots show the mean plus-minus two standard errors. CMCGS demonstrates the best overall performance compared to the baselines. In the four 2D environments that test exploration, CMCGS is the best, and it is also the best overall in the DMC environments where exploration is less challenging. CMCGS is highly robust to the different sizes of simulation budgets and achieves the best overall mean performance in both the 2D and DMC environments at every simulation budget. CEM uses a single bandit for each time step, MCTS-PW, VOOT, and MCTS-ABS create a search tree that grows with each time step, and CMA-ES creates a distribution over all time steps with temporal correlations but assumes a uni-modal distribution and does not divide planning into different nodes for different parts of the state distributions [53, 56]. Therefore, CEM and CMA-ES perform poorly in tasks that

require strong exploration, and VOOT, MCTS-PW, and MCTS-ABS perform poorly in the DMC tasks that require taking advantage of informative rewards and dealing with high-dimensional action spaces. CMCGS uses several nodes, similar to MCTS, for each layer but strongly limits the number of nodes at each time step. CMCGS explores successfully but can also take advantage of informative rewards to find the best actions with the node policies $\pi_q(a)$.

CMA-ES outperforms CMCGS in the *cheetah* and *walker* environments, most likely due to the covariance matrix that takes correlations over all time steps into account. However, the benefit of the full covariance matrix is nonexistent in other environments. CMA-ES is sometimes inferior to CEM despite CEM assuming a much simpler sampling distribution without any temporal correlations, which was also observed in Duan et al. [18]. The MCTS-based algorithms VOOT, MCTS-PW, and MCTS-ABS perform well in low-dimensional and sparse reward environments such as the 2d navigation and reacher environments and *ball_in_cup_catch*, but they struggle with high-dimensional action spaces, especially in the *walker* environments. We observe that the progressive widening strategy is insufficient for making the MCTS-based methods handle long planning horizons with the search space growing exponentially. On our 2D environments and DMC walker tasks, the running time of CMCGS has the same magnitude as the fastest method, random shooting, and is superior to that of CMA-ES and MCTS-ABS. The clustering in CMCGS is not highly time-consuming because of the limitations on how often clustering is attempted and the relatively small size of the replay memory of CMCGS.

5.3 Large-Scale CMCGS with Learned Models

We demonstrate the generality of CMCGS by evaluating it with learned dynamics models in high-dimensional continuous control environments and comparing it to CEM. For evaluation, we use pixel-based environments from the DeepMind Control Suite [66] and PlaNet [25] as the underlying algorithm for learning the dynamics model. We use an open-sourced implementation of PlaNet and pre-trained models by Arulkumaran [4] and replace CEM with CMCGS as the planner during evaluation. PlaNet has an encoder that learns a latent representation from pixels, and CMCGS performs planning and clustering in this learned latent space.

We copy the hyperparameters for CEM from the original paper [25] and either perform equally many simulator interactions with CMCGS or spend equal time. We let CMCGS collect 400 trajectories in parallel to make the running time comparable to CEM. The mean episode returns in Table 2 show that CMCGS outperforms CEM in six of the seven environments for the same number of environment steps and in five of the seven environments in total and is significantly better overall. We find that CMCGS does not exploit the model inaccuracies, which makes it a robust and general planning algorithm for high-dimensional continuous control tasks.

5.4 Ablations

We perform several ablation studies to analyze the impact of our design choices on the performance of CMCGS. To make the returns from our custom environments and DMC environments comparable, we scale the former by adding one and multiplying by 500. The impacts of various design choices related to the action selection

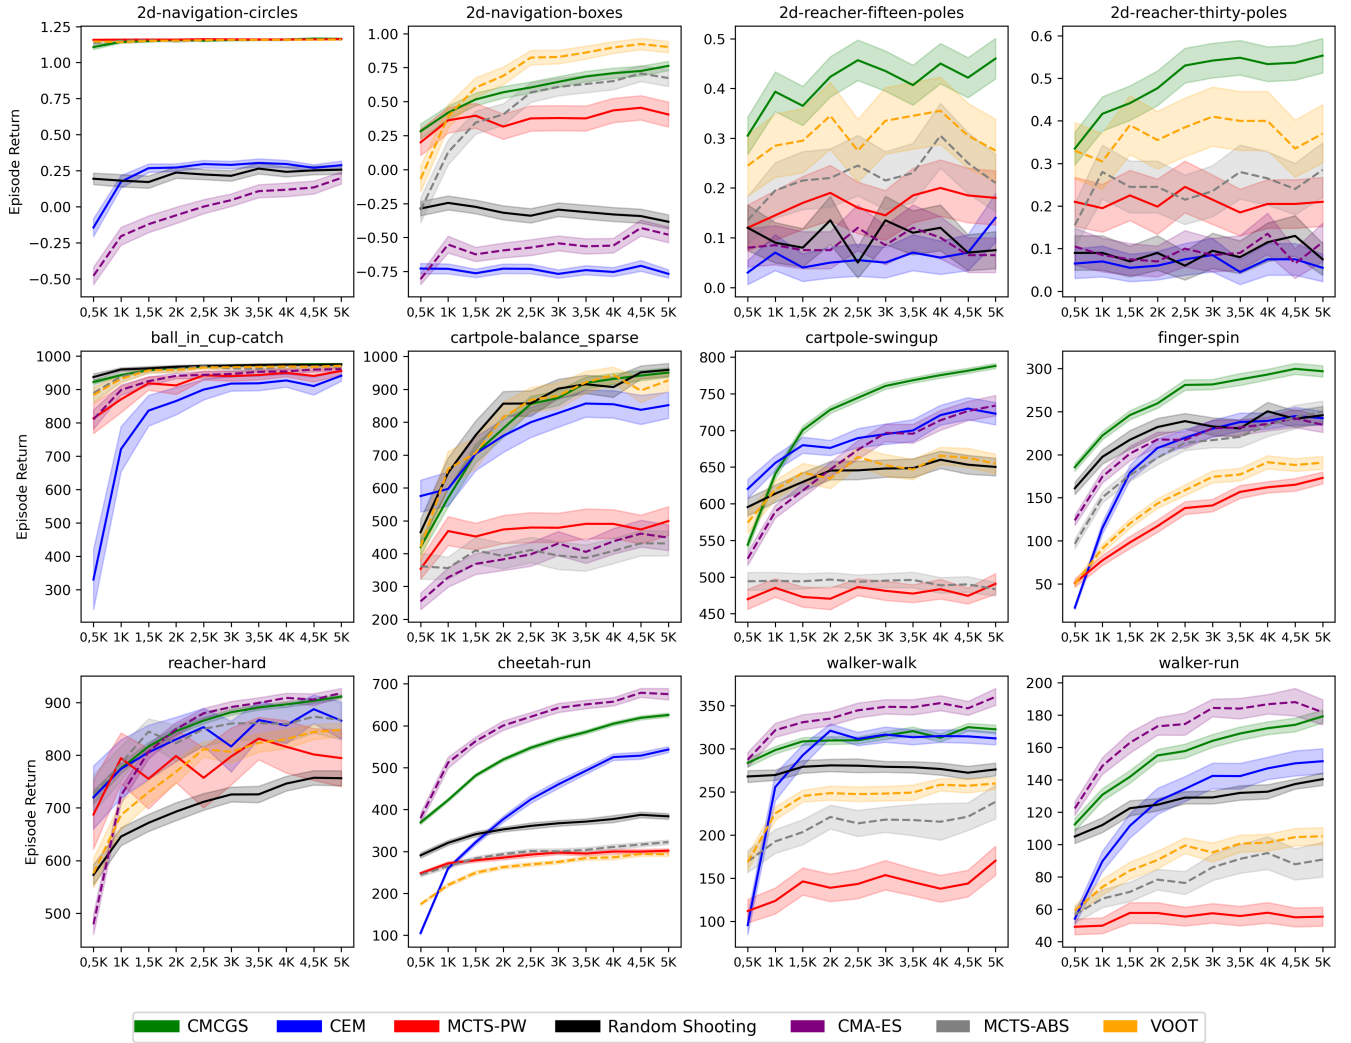


Figure 4: Reward plots for different simulation budgets per timestep. 2d-X environments (first row) have been developed by the authors to test exploration. The rest are from DeepMind Control Suite with proprioceptive observations. The proposed CMCGS shows the best performance overall.

are shown in Table 3. Replacing the initial Gaussian distribution for $\pi_q(\mathbf{a})$ with actions sampled from a uniform distribution until the first update degrades the performance. Furthermore, adding uniform exploration with probability p_{unif} throughout the planning process hurts the performance. The Bayesian variance updates (see Eq. 1) are crucial for high performance and outperform CEM-like variance updates, where the new variance for $\pi_q(\mathbf{a})$ is equal to the variance of the elite actions. We also evaluate discarding the CEM-inspired elite actions entirely and simply letting the means of $\pi_q(\mathbf{a})$ equal the reward-weighted means of the actions in the replay buffer, but this approach fails to improve the performance of our method. We also experiment with splitting the action space into Voronoi cells [39] and sampling the greedy action from the cell corresponding to the best action found in \mathcal{D}_q , but it is not competitive with the current greedy action selection mechanism, where we

sample one of the top actions and add noise to it. When $\mathcal{E}_{\text{top}} = 0$, we do not add any noise to the greedy actions sampled from the replay buffer \mathcal{D}_q . When $N_{\text{top}} = 1$, we take the best action deterministically from \mathcal{D}_q . Finally, we evaluate values of ϵ different from the chosen one, $\epsilon = 0.7$. When $\epsilon = 1$, we always sample from $\pi_q(\mathbf{a})$ and never use the top actions found in \mathcal{D}_q . None of the ablations presented here improve on the chosen CMCGS design, but many of them are close to the original result-wise, which highlights the robustness of CMCGS. In Table 4, we compare clustering algorithms. KMeans and GMM are not significantly inferior to hierarchical clustering in terms of performance, but they are slower to run.

6 LIMITATIONS AND FUTURE WORK

In this section, we explain the main limitations of our work and how they could be addressed in future work.

Table 2: The average rewards attained in the DMC environments with image observations using learned PlaNet dynamics models plus-minus two standard errors. There are two variants of CEM. In the first variant, the number of simulator steps is equal to the number of simulator steps used by CMCGS. In the second variant of CEM, the running times of the algorithms have been matched on our workstation.

Environment	CEM (steps)	CEM (time)	CMCGS
Ball-in-Cup Catch	713.4 \pm 31.6	728.4 \pm 32.8	844.6 \pm 20.6
Cartpole Balance	996.7 \pm 0.1	991.0 \pm 0.3	993.4 \pm 0.2
Cartpole Swingup	763.1 \pm 4.5	778.0 \pm 12.1	857.0 \pm 0.8
Cheetah Run	556.5 \pm 8.5	615.8 \pm 12.7	628.3 \pm 16.6
Finger Spin	651.1 \pm 9.5	857.7 \pm 2.4	833.9 \pm 3.2
Reacher Easy	828.4 \pm 46.1	832.6 \pm 45.7	887.1 \pm 36.1
Walker Walk	863.1 \pm 21.3	901.9 \pm 9.1	950.6 \pm 4.6
Mean	767.5 \pm 23.2	815.0 \pm 22.6	856.3 \pm 17.1

Table 3: The average rewards attained in 2D and DMC environments for simulation budget 2,500 plus-minus two standard errors for different action bandit design choices.

Ablation	Result
CMCGS	679.6 \pm 22.5
Initial action distribution uniform	666.1 \pm 21.1
Uniform exploration with $p_{\text{unif}} = 0.05\%$	671.9 \pm 21.4
Uniform exploration with $p_{\text{unif}} = 0.1\%$	668.5 \pm 21.8
No Bayesian variance updates	633.1 \pm 26.6
Reward-weighted mean	673.3 \pm 21.2
VOO-sampling	658.7 \pm 23.6
Greedy actions without noise, $\mathcal{E}_{\text{top}} = 0$	670.2 \pm 22.4
No sampling of greedy actions, $\mathcal{N}_{\text{top}} = 1$	647.5 \pm 23.6
No greedy actions, $\epsilon = 1.0$	649.1 \pm 22.5
Higher proportion of greedy actions, $\epsilon = 0.5$	668.7 \pm 23.2

Table 4: The average rewards attained in 2D and DMC environments for simulation budget 2,500 plus-minus two standard errors for different clustering algorithms.

Clustering algorithm	Result
Hierarchical (Ward linkage)	679.6 \pm 22.5
KMeans	674.2 \pm 31.2
GMM	670.8 \pm 22.1

Bootstrapping: In our current implementation, CMCGS builds the search graph from scratch at each timestep. Model-predictive control can benefit from using the best solution from previous timestep(s) to bias the optimization process towards more promising trajectories if the model is accurate [7, 33]. In CMCGS, this could be done by biasing the mean of the action policies of the search nodes towards the old best actions. A more trivial approach would be to simply inject the old best trajectory into the replay memory of the search nodes.

Learning: MCTS can benefit from learned value functions, policy priors, or rollout policies that have been trained using imitation learning from pre-existing datasets [38, 60, 64]. We believe that this could improve the performance of CMCGS.

Arbitrary graph structures: Currently, CMCGS uses a layered directed acyclic graph to represent the state space, which does not allow the search graph to re-use previous experience in environments where the agent can navigate back and forth between different states (such as searching through a maze). This limitation could be lifted by allowing the algorithm to build arbitrary graph structures (such as complete graphs) to better utilize the structure of the state space during planning.

Theoretical properties: We leave the analysis of the theoretical properties of CMCGS and its convergence to future work. Note that our ablation study in Table 3 shows that adding uniform exploration and Voronoi optimistic sampling during the planning process has a negative impact on the performance of CMCGS in the 2D and DMC environments that we evaluated it on.

7 CONCLUSION

In this paper, we propose Continuous Monte Carlo Graph Search (CMCGS), an extension of the popular MCTS algorithm for solving decision-making problems with high-dimensional continuous state and action spaces. CMCGS builds up on the observation that different regions of the state space ask for different action bandits for estimating the reward and representing the action distribution. Based on this observation, CMCGS builds a layered search graph where, at each layer, the visited states are clustered into several stochastic action bandit nodes, which allows CMCGS to solve complex high-dimensional continuous control problems efficiently. CMCGS outperforms MCTS-PW, VOOT, MCTS-ABS, CEM, and CMA-ES in several complex continuous control environments such as DeepMind Control Suite benchmarks. Experiments in sparse-reward custom environments indicate that CMCGS has significantly better exploration capabilities than the evaluated baselines, especially given a high-dimensional action space.

CMCGS can be efficiently parallelized, which makes the algorithm practically relevant also for large-scale planning. Our experiments with the learned PlaNet dynamics models show that CMCGS can be scaled up to be competitive with CEM as a planning component of general model-based reinforcement learning algorithms that learn from pixels, even when the number of environment interactions is not a limiting factor. We believe that the proposed CMCGS algorithm could be a building block for a new family of Monte Carlo methods applied to decision-making problems with continuous action spaces.

ACKNOWLEDGMENTS

We acknowledge the computational resources provided by the Aalto Science-IT project and CSC, Finnish IT Center for Science. The work was funded by Research Council of Finland (aka Academy of Finland) within the Flagship Programme, Finnish Center for Artificial Intelligence (FCAI). J. Pajarinen and Y. Zhao were partly supported by Research Council of Finland (345521).

ETHICS STATEMENT

Our work does not appear to have any immediate adverse effects on society. Our research did not involve any human subjects or participants. We only use pre-trained neural network models that have not been trained on sensitive or confidential information. Neither did we use any such information ourselves. Our model does not directly pertain to making real-world decisions involving individuals. However, our work is focused on improving continuous control, which could be applicable to, for instance, robotics. Therefore, it is, in principle, possible that our methods could be employed in unforeseen harmful ways, such as military or law enforcement robotics.

REFERENCES

- [1] David Abel, David Hershkowitz, and Michael Littman. 2016. Near optimal behavior via approximate state abstraction. In *International Conference on Machine Learning*. PMLR, 2915–2923.
- [2] Zaheen Farraz Ahmad, Levi Leis, and Michael Bowling. 2020. Marginal Utility for Planning in Continuous or Large Discrete Action Spaces. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6–12, 2020, virtual*, Hugo Larochelle, Marc'Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin (Eds.). <https://proceedings.neurips.cc/paper/2020/hash/14da15db887a4b50efe5c1bc66537089-Abstract.html>
- [3] Thomas Anthony, Zheng Tian, and David Barber. 2017. Thinking Fast and Slow with Deep Learning and Tree Search. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4–9, 2017, Long Beach, CA, USA*, Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett (Eds.). 5360–5370. <https://proceedings.neurips.cc/paper/2017/hash/d8e1344e27a5b08cdf5d027d9b8d6de-Abstract.html>
- [4] Kai Arulkumaran. 2019. Pytorch Implementation of PlaNet: A Deep Planning Network for Reinforcement Learning. <https://github.com/Kaixhin/PlaNet>.
- [5] Amin Babadi, Kourosh Naderi, and Perttu Härmäläinen. 2018. Intelligent middle-level game control. In *Proceedings of IEEE Conference on Computational Intelligence and Games (CIG)*. IEEE, 25–32.
- [6] Aijun Bai, Siddharth Srivastava, and Stuart Russell. 2016. Markovian State and Action Abstractions for MDPs via Hierarchical MCTS. In *IJCAI*. 3029–3039.
- [7] Mohak Bhardwaj, Sanjiban Choudhury, and Byron Boots. 2021. Blending MPC & Value Function Approximation for Efficient Reinforcement Learning. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3–7, 2021*. OpenReview.net. https://openreview.net/forum?id=RqCC_00Bg7V
- [8] Mohak Bhardwaj, Ankur Handa, Dieter Fox, and Byron Boots. 2020. Information theoretic model predictive q-learning. In *Learning for Dynamics and Control*. PMLR, 840–850.
- [9] Cameron B Browne, Edward Powley, Daniel Whitehouse, Simon M Lucas, Peter I Cowling, Philipp Rohlfshagen, Stephen Tavener, Diego Perez, Spyridon Samothrakis, and Simon Colton. 2012. A survey of monte carlo tree search methods. *IEEE Transactions on Computational Intelligence and AI in games* 4, 1 (2012), 1–43.
- [10] Sébastien Bubeck, Rémi Munos, Gilles Stoltz, and Csaba Szepesvári. 2008. Online Optimization in X-Armed Bandits. In *Advances in Neural Information Processing Systems 21, Proceedings of the Twenty-Second Annual Conference on Neural Information Processing Systems, Vancouver, British Columbia, Canada, December 8–11, 2008*, Daphne Koller, Dale Schuurmans, Yoshua Bengio, and Léon Bottou (Eds.). Curran Associates, Inc., 201–208. <https://proceedings.neurips.cc/paper/2008/hash/f387624df552cea2f369918c5e1e12bc-Abstract.html>
- [11] Guillaume M JB Chaslot, Mark HM Winands, H Jaap van den Herik, Jos WHM Uiterwijk, and Bruno Bouzy. 2008. Progressive strategies for Monte-Carlo tree search. *New Mathematics and Natural Computation* 4, 03 (2008), 343–357.
- [12] Benjamin E Childs, James H Brodeur, and Levente Kocsis. 2008. Transpositions and move groups in Monte Carlo tree search. In *2008 IEEE Symposium on Computational Intelligence and Games*. IEEE, 389–395.
- [13] Kurtland Chua, Roberto Calandra, Rowan McAllister, and Sergey Levine. 2018. Deep Reinforcement Learning in a Handful of Trials using Probabilistic Dynamics Models. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3–8, 2018, Montréal, Canada*, Samy Bengio, Hanna M. Wallach, Hugo Larochelle, Kristen Grauman, Nicolò Cesa-Bianchi, and Roman Garnett (Eds.). 4759–4770. <https://proceedings.neurips.cc/paper/2018/hash/3de568f8597b94bda53149c7d7f5958c-Abstract.html>
- [14] Adrien Couëtoux, Jean-Baptiste Hoock, Nataliya Sokolovska, Olivier Teytaud, and Nicolas Bonnard. 2011. Continuous upper confidence trees. In *International Conference on Learning and Intelligent Optimization*. Springer, 433–445.
- [15] Rémi Coulom. 2006. Efficient selectivity and backup operators in Monte-Carlo tree search. In *International conference on computers and games*. Springer, 72–83.
- [16] Rémi Coulom. 2007. Computing “elo ratings” of move patterns in the game of go. *ICGA journal* 30, 4 (2007), 198–208.
- [17] Johannes Czech, Patrick Korus, and Kristian Kersting. 2021. Improving AlphaZero Using Monte-Carlo Graph Search. In *Proceedings of the International Conference on Automated Planning and Scheduling*, Vol. 31. 103–111.
- [18] Yan Duan, Xi Chen, Rein Houthoofd, John Schulman, and Pieter Abbeel. 2016. Benchmarking Deep Reinforcement Learning for Continuous Control. In *Proceedings of the 33rd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19–24, 2016 (JMLR Workshop and Conference Proceedings, Vol. 48)*, Maria-Florina Balcan and Kilian Q. Weinberger (Eds.). JMLR.org, 1329–1338. <http://proceedings.mlr.press/v48/duan16.html>
- [19] Richard E Fikes and Nils J Nilsson. 1971. STRIPS: A new approach to the application of theorem proving to problem solving. *Artificial intelligence* 2, 3–4 (1971), 189–208.
- [20] Hilmar Finnsson and Yngvi Björnsson. 2008. Simulation-based approach to general game playing. In *Aaai*, Vol. 8. 259–264.
- [21] Raluca D Gaina, Simon M Lucas, and Diego Pérez-Liébana. 2017. Population seeding techniques for rolling horizon evolution in general video game playing. In *Evolutionary Computation (CEC), 2017 IEEE Congress on*. IEEE, 1956–1963.
- [22] Sylvain Gelly and David Silver. 2011. Monte-Carlo tree search and rapid action value estimation in computer Go. *Artificial Intelligence* 175, 11 (2011), 1856–1875.
- [23] Jean-Bastien Grill, Florent Althé, Yunhao Tang, Thomas Hubert, Michal Valko, Ioannis Antonoglou, and Rémi Munos. 2020. Monte-Carlo Tree Search as Regularized Policy Optimization. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13–18 July 2020, Virtual Event (Proceedings of Machine Learning Research, Vol. 119)*. PMLR, 3769–3778. <http://proceedings.mlr.press/v119/grill20a.html>
- [24] Xiaoxiao Guo, Satinder P. Singh, Honglak Lee, Richard L. Lewis, and Xiaoshi Wang. 2014. Deep Learning for Real-Time Atari Game Play Using Offline Monte-Carlo Tree Search Planning. In *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8–13 2014, Montreal, Quebec, Canada*, Zoubin Ghahramani, Max Welling, Corinna Cortes, Neil D. Lawrence, and Kilian Q. Weinberger (Eds.). 3338–3346. <https://proceedings.neurips.cc/paper/2014/hash/8bb88f80d334b1869781beb89f7b73be-Abstract.html>
- [25] Danijar Hafner, Timothy P. Lillicrap, Ian Fischer, Ruben Villegas, David Ha, Honglak Lee, and James Davidson. 2019. Learning Latent Dynamics for Planning from Pixels. In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9–15 June 2019, Long Beach, California, USA (Proceedings of Machine Learning Research, Vol. 97)*, Kamalika Chaudhuri and Ruslan Salakhutdinov (Eds.). PMLR, 2555–2565. <http://proceedings.mlr.press/v97/hafner19a.html>
- [26] Perttu Härmäläinen, Sebastian Eriksson, Esa Tanskanen, Ville Kyrki, and Jaakko Lehtinen. 2014. Online motion synthesis using sequential monte carlo. *ACM Transactions on Graphics (TOG)* 33, 4 (2014), 51.
- [27] Jessica B. Hamrick, Abram L. Friesen, Feryal M. P. Behbahani, Arthur Guez, Fabio Viola, Sims Witherspoon, Thomas Anthony, Lars Holger Buesing, Petar Velickovic, and Theophane Weber. 2021. On the role of planning in model-based deep reinforcement learning. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3–7, 2021*. OpenReview.net. <https://openreview.net/forum?id=IrM64DGB21>
- [28] Nikolaus Hansen, Sibylle D Müller, and Petros Koumoutsakos. 2003. Reducing the time complexity of the derandomized evolution strategy with covariance matrix adaptation (CMA-ES). *Evolutionary computation* 11, 1 (2003), 1–18.
- [29] Nikolaus Hansen and Andreas Ostermeier. 1996. Adapting arbitrary normal mutation distributions in evolution strategies: The covariance matrix adaptation. In *Proceedings of IEEE international conference on evolutionary computation*. IEEE, 312–317.
- [30] Nicklas Hansen, Hao Su, and Xiaolong Wang. 2022. Temporal Difference Learning for Model Predictive Control. In *International Conference on Machine Learning, ICML 2022, 17–23 July 2022, Baltimore, Maryland, USA (Proceedings of Machine Learning Research, Vol. 162)*, Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvári, Gang Niu, and Sivan Sabato (Eds.). PMLR, 8387–8406. <https://proceedings.mlr.press/v162/hansen22a.html>
- [31] Nathan Hatch and Byron Boots. 2021. The Value of Planning for Infinite-Horizon Model Predictive Control. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 7372–7378.
- [32] Verena Heidrich-Meisner and Christian Igel. 2009. Uncertainty handling CMA-ES for reinforcement learning. In *Proceedings of the 11th Annual conference on Genetic and evolutionary computation*. 1211–1218.
- [33] Martin Herceg, CN Jones, and M Morari. 2015. Dominant speed factors of active set methods for fast MPC. *Optimal Control Applications and Methods* 36, 5 (2015), 608–627.
- [34] Christoffer Holmgård, Michael Cerny Green, Antonios Liapis, and Julian Togelius. 2018. Automated playtesting with procedural personas through MCTS with

- evolved heuristics. *IEEE Transactions on Games* 11, 4 (2018), 352–362.
- [35] Jesse Hostetler, Alan Fern, and Tom Dietterich. 2014. State aggregation in Monte Carlo tree search. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 28.
- [36] Jesse Hostetler, Alan Fern, and Thomas Dietterich. 2017. Sample-based tree search with fixed and adaptive state abstractions. *Journal of Artificial Intelligence Research* 60 (2017), 717–777.
- [37] Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Mohammadamin Barekatain, Simon Schmitt, and David Silver. 2021. Learning and Planning in Complex Action Spaces. In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18–24 July 2021, Virtual Event (Proceedings of Machine Learning Research, Vol. 139)*, Marina Meila and Tong Zhang (Eds.). PMLR, 4476–4486. <http://proceedings.mlr.press/v139/hubert21a.html>
- [38] Bilal Kartal, Pablo Hernandez-Leal, and Matthew E Taylor. 2019. Action guidance with MCTS for deep reinforcement learning. In *Proceedings of the AAAI conference on artificial intelligence and interactive digital entertainment*, Vol. 15. 153–159.
- [39] Beomjoon Kim, Kyungjae Lee, Sungbin Lim, Leslie Pack Kaelbling, and Tomás Lozano-Pérez. 2020. Monte Carlo Tree Search in Continuous Spaces Using Voronoi Optimistic Optimization with Regret Bounds. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7–12, 2020*. AAAI Press, 9916–9924. <https://aaai.org/ojs/index.php/AAAI/article/view/6546>
- [40] Levente Kocsis and Csaba Szepesvári. 2006. Bandit based monte-carlo planning. In *European conference on machine learning*. Springer, 282–293.
- [41] Jongmin Lee, Wonseok Jeon, Geon-Hyeon Kim, and Kee-Eung Kim. 2020. Monte-Carlo Tree Search in Continuous Action Spaces with Value Gradients. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7–12, 2020*. AAAI Press, 4561–4568. <https://aaai.org/ojs/index.php/AAAI/article/view/5885>
- [42] Edouard Leurent and Odalric-Ambrym Maillard. 2020. Monte-carlo graph search: the value of merging similar states. In *Asian Conference on Machine Learning*. PMLR, 577–592.
- [43] Lihong Li, Thomas J Walsh, and Michael L Littman. 2006. Towards a unified theory of state abstraction for MDPs. In *AI&M*.
- [44] Kendall Lowrey, Aravind Rajeswaran, Sham M. Kakade, Emanuel Todorov, and Igor Mordatch. 2019. Plan Online, Learn Offline: Efficient Learning and Exploration via Model-Based Control. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6–9, 2019*. OpenReview.net. <https://openreview.net/forum?id=Byey7n05FQ>
- [45] Weichao Mao, Kaiqing Zhang, Qiaomin Xie, and Tamer Basar. 2020. POLY-HOOT: Monte-Carlo Planning in Continuous Space MDPs with Non-Asymptotic Analysis. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6–12, 2020, virtual*, Hugo Larochelle, Marc'Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin (Eds.). <https://proceedings.neurips.cc/paper/2020/hash/30de24287a6d8f07b37c716ad51623a7-Abstract.html>
- [46] Thomas M Moerland, Joost Broekens, Aske Plaat, and Catholijn M Jonker. 2018. A0c: Alpha zero in continuous action space. *ArXiv preprint abs/1805.09613* (2018). <https://arxiv.org/abs/1805.09613>
- [47] Igor Mordatch, Emanuel Todorov, and Zoran Popović. 2012. Discovery of complex behaviors through contact-invariant optimization. *ACM Transactions on Graphics (TOG)* 31, 4 (2012), 1–8.
- [48] Rémi Munos et al. 2014. From bandits to Monte-Carlo Tree Search: The optimistic principle applied to optimization and planning. *Foundations and Trends® in Machine Learning* 7, 1 (2014), 1–129.
- [49] Kourosh Naderi, Joose Rajamäki, and Perttu Hämäläinen. 2017. Discovering and synthesizing humanoid climbing movements. *ACM Transactions on Graphics (TOG)* 36, 4, Article 43 (2017), 11 pages. <https://doi.org/10.1145/3072959.3073707>
- [50] Rudy R Negenborn, Bart De Schutter, Marco A Wiering, and Hans Hellendoorn. 2005. Learning-based model predictive control for Markov decision processes. *IFAC Proceedings Volumes* 38, 1 (2005), 354–359.
- [51] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. 2011. Scikit-learn: Machine learning in Python. *Journal of machine Learning research* 12 (2011), 2825–2830.
- [52] Diego Perez, Spyridon Samothrakis, and Simon Lucas. 2014. Knowledge-based fast evolutionary MCTS for general video game playing. In *2014 IEEE Conference on Computational Intelligence and Games*. IEEE, 1–8.
- [53] Cristina Pinneri, Shambhuraj Sawant, Sebastian Blaes, Jan Achterhold, Joerg Stueckler, Michal Rolinek, and Georg Martius. 2021. Sample-efficient cross-entropy method for real-time planning. In *Conference on Robot Learning*. PMLR, 1049–1065.
- [54] Ricardo Quintero, Francisco S Melo, and Pedro A Santos. 2021. Limited depth bandit-based strategy for Monte Carlo planning in continuous action spaces. *ArXiv preprint abs/2106.15594* (2021). <https://arxiv.org/abs/2106.15594>
- [55] Joose Julius Rajamäki and Perttu Hämäläinen. 2018. Continuous control monte carlo tree search informed by multiple experts. *IEEE transactions on visualization and computer graphics* (2018).
- [56] Reuven Y Rubinstein. 1997. Optimization of computer simulation models with rare events. *European Journal of Operational Research* 99, 1 (1997), 89–112.
- [57] Reuven Y Rubinstein and Dirk P Kroese. 2004. *The cross-entropy method: a unified approach to combinatorial optimization, Monte-Carlo simulation, and machine learning*. Vol. 133. Springer.
- [58] Spyridon Samothrakis, Samuel A Roberts, Diego Perez, and Simon M Lucas. 2014. Rolling horizon methods for games with continuous states and actions. In *Computational Intelligence and Games (CIG), 2014 IEEE Conference on*. IEEE, 1–8.
- [59] Julian Schrittwieser, Ioannis Antonoglou, Thomas Hubert, Karen Simonyan, Laurent Sifre, Simon Schmitt, Arthur Guez, Edward Lockhart, Demis Hassabis, Thore Graepel, et al. 2020. Mastering atari, go, chess and shogi by planning with a learned model. *Nature* 588, 7839 (2020), 604–609.
- [60] David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. 2016. Mastering the game of Go with deep neural networks and tree search. *nature* 529, 7587 (2016), 484–489.
- [61] David Silver, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, Laurent Sifre, Dharrshan Kumaran, Thore Graepel, et al. 2018. A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play. *Science* 362, 6419 (2018), 1140–1144.
- [62] David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. 2017. Mastering the game of go without human knowledge. *nature* 550, 7676 (2017), 354–359.
- [63] Samuel Sokota, Caleb Y Ho, Zaheen Ahmad, and J Zico Kolter. 2021. Monte carlo tree search with iteratively refining state abstractions. *Advances in Neural Information Processing Systems* 34 (2021), 18698–18709.
- [64] Maciej Świechowski, Tomasz Tajmajer, and Andrzej Janusz. 2018. Improving hearthstone ai by combining mcts and supervised learning algorithms. In *2018 IEEE conference on computational intelligence and games (CIG)*. IEEE, 1–8.
- [65] Yunhao Tang and Shipra Agrawal. 2020. Discretizing Continuous Action Space for On-Policy Optimization. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7–12, 2020*. AAAI Press, 5981–5988. <https://aaai.org/ojs/index.php/AAAI/article/view/6059>
- [66] Yuval Tassa, Yotam Doron, Alistair Muldal, Tom Erez, Yazhe Li, Diego de Las Casas, David Budden, Abbas Abdolmaleki, Josh Merel, Andrew Lefrancq, et al. 2018. Deepmind control suite. *ArXiv preprint abs/1801.00690* (2018). <https://arxiv.org/abs/1801.00690>
- [67] Yuval Tassa, Tom Erez, and Emanuel Todorov. 2012. Synthesis and stabilization of complex behaviors through online trajectory optimization. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 4906–4913.
- [68] Saran Tunyasuvunakool, Alistair Muldal, Yotam Doron, Siqi Liu, Steven Bohez, Josh Merel, Tom Erez, Timothy Lillicrap, Nicolas Heess, and Yuval Tassa. 2020. dm_control: Software and tasks for continuous control. *Software Impacts* 6 (2020), 100022. <https://doi.org/10.1016/j.simpa.2020.100022>
- [69] Benjamin Van Roy. 2006. Performance loss bounds for approximate value iteration with state aggregation. *Mathematics of Operations Research* 31, 2 (2006), 234–244.
- [70] Joe H Ward Jr. 1963. Hierarchical grouping to optimize an objective function. *Journal of the American statistical association* 58, 301 (1963), 236–244.
- [71] Ari Weinstein and Michael L. Littman. 2013. Open-Loop Planning in Large-Scale Stochastic Domains. In *Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence, July 14–18, 2013, Bellevue, Washington, USA*, Marie desJardins and Michael L. Littman (Eds.). AAAI Press. <http://www.aaai.org/ocs/index.php/AAAI/AAAI13/paper/view/6341>
- [72] Linjie Xu, Alexander Döckhorn, and Diego Perez-Liebana. 2023. Elastic Monte Carlo Tree Search. *IEEE Transactions on Games* (2023).
- [73] Weirui Ye, Shaohuai Liu, Thanard Kurutach, Pieter Abbeel, and Yang Gao. 2021. Mastering Atari Games with Limited Data. In *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6–14, 2021, virtual*, Marc'Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan (Eds.). 25476–25488. <https://proceedings.neurips.cc/paper/2021/hash/d5eca8dc3820cad9fe56a3bafda65ca1-Abstract.html>
- [74] Timothy Yee, Viliam Lisý, and Michael H. Bowling. 2016. Monte Carlo Tree Search in Continuous Action Spaces with Execution Uncertainty. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 2016, New York, NY, USA, 9–15 July 2016*, Subbarao Kanbampati (Ed.). IJCAI/AAAI Press, 690–697. <http://www.ijcai.org/Abstract/16/104>