

Convex Hull Monte-Carlo Tree-Search

Michael Painter, Bruno Lacerda, Nick Hawes

Oxford Robotics Institute, University of Oxford

Oxford, UK

{mpainter, bruno, nickh}@robots.ox.ac.uk

Abstract

This work investigates Monte-Carlo planning for agents in stochastic environments, with multiple objectives. We propose the *Convex Hull Monte-Carlo Tree-Search* (CHMCTS) framework, which builds upon *Trial Based Heuristic Tree Search* and *Convex Hull Value Iteration* (CHVI), as a solution to multi-objective planning in large environments. Moreover, we consider how to pose the problem of approximating multi-objective planning solutions as a *contextual multi-armed bandits* problem, giving a principled motivation for how to select actions from the view of *contextual regret*. This leads us to the use of *Contextual Zooming* for action selection, yielding *Zooming CHMCTS*. We evaluate our algorithm using the *Generalised Deep Sea Treasure* environment, demonstrating that *Zooming CHMCTS* can achieve a sublinear contextual regret and scales better than CHVI on a given computational budget.

1 Introduction

In *Multi-Objective Planning under uncertainty* (MOPU) an agent has to plan in a stochastic environment while trading off multiple objectives. Thus, we often assume the agent’s environment is known and modeled as a *Multi-Objective Markov Decision Process* (MOMDP). This problem is useful to consider as agents often encounter environments with uncertainty in their dynamics. Additionally, it is often useful for agents to be able to balance different objectives, for which the priorities are not known a priori or may change over time. MOPU has been applied to tackle problems in several domains, such as having a robot trade-off between battery consumption and performance in its primary task (Lahijanian et al. 2018) or between timely achievement of a primary tasks whilst achieving as many soft goals as possible (Lacerda, Parker, and Hawes 2017); or when planning how to lay new electrical lines while trading installation and operational costs for network reliability (Sahoo, Ganguly, and Das 2012). The objective of MOPU is to compute a set of values, and associated policies, known as a *Pareto Front*, that is optimal for some prioritisation over the objectives. Unfortunately, because there is not a single optimum value, as is the case in single-objective planning, there is an additional dimension of complexity to

handle. In single-objective planning under uncertainty two techniques have recently lead to a significant improvement in the scalability of algorithms: *Monte-Carlo Tree-Search* (MCTS) (Kocsis and Szepesvári 2006) and *Value Function Approximation* (Sutton et al. 2000). Note that these techniques do not have to be mutually exclusive, as demonstrated in a system such as Alpha Go (Silver et al. 2017). However, there has been relatively little work in adapting either of these techniques to the multi-objective setting (Perez et al. 2015; Xu and Li 2017; Chen and Lie 2019; Mossalam et al. 2016; Abels et al. 2018).

An additional gap that needs to be addressed in the multi-objective setting is the need for principled approaches to evaluate the *online performance of trial based planning algorithms* (such as MCTS). In this paper, we propose a regret based metric to do so. We consider a sequence of trials, each with a known priority over the different objectives, such that the performance of the algorithm can be mapped to a scalar value that we wish to optimise. This leads us to use the notion of *Contextual Regret* as a measure of the online performance of a MOPU algorithm.

The main contributions of this work are: (i) applying the notion of contextual regret to *multi-objective planning*, and justify that exploration policies that achieve low contextual regret explore the trade-offs between objectives appropriately, as opposed to other metrics proposed in the literature; (2) proposing *Contextual Zooming for Trees*, that outperforms prior work on this metric.

2 Related Work

In the *Multi-Objective Multi-Armed Bandit* (MOMAB) problem, an agent must pick one of K arms round by round such that we optimise a multi-objective payoff vector. The MOMAB problem can be considered a special case of MOPU, as it can be mapped to a finite horizon MOMDP with a single state with K actions, one for each arm. Drugan and Nowe (2013) address the MOMAB problem by defining a set of arms, the *Pareto Front*, that are all considered optimal, in the sense that the performance in one objective cannot be improved without degrading the performance for least one of the other objectives. They extend the well known UCB1 (Auer, Cesa-Bianchi, and Fischer 2002) algorithm to

the MOMAB problem, using a multi-dimensional confidence interval rather than a one-dimensional confidence interval. Because the algorithm was adapted from UCB1, they refer to it as *ParetoUCB1*.

Multi-objective sequential decision making (which includes *multi-objective planning*) extends ideas from single-objective sequential decision making algorithms, to handle a vector of rewards. In such problems, the solution to be computed is a *Pareto Front* or a *Convex Hull*, which is generally accepted to represent every possible trade off between objectives that could be made. For an in depth introduction to the field, see (Rojers et al. 2013).

Prior work in *Multi-Objective Monte-Carlo Tree-Search* (MOMCTS) can be divided into two categories, those that maintain a single scalar value estimate at each node, which we will call *Point-Based MOMCTS*, and those that maintain an estimate of the Pareto front at each node.

In terms of point-based MOMCTS, Wang and Sebag (2012) maintain a global Pareto front for the root node, and during trials, they select successor nodes based on how “close” the value estimate in the child nodes are to the global Pareto front. In (Chen and Lie 2019) the Pareto front for a node is formed from the value estimates of the children nodes, and actions are selected by running ParetoUCB over those points.

In terms of maintaining an estimate of the Pareto front at each node, Perez, Samothrakis, and Lucas; Perez et al.; Perez-Liebana, Mostaghim, and Lucas (2013; 2015; 2016) define an MOMCTS algorithm for games with deterministic transitions. The deterministic assumption simplifies the operations for updating the Pareto fronts. Our algorithm uses generalised versions of these updates, allowing for stochastic settings too. The rule for selecting successor nodes is adapted from the UCB1 algorithm (Auer, Cesa-Bianchi, and Fischer 2002) over the *hypervolumes* of child nodes. The hypervolume can be thought of as the area under the Pareto front. Xu and Li (2017) use the same algorithm that is presented by Perez, Samothrakis, and Lucas (2013), however, they replace the use of the hypervolumes in UCB1 with the *Chebychev scalarization function*.

All the works described above either assume a deterministic environment or do not maintain an estimated Pareto front at each search node. In this paper, we will show why maintaining an estimated Pareto front at each search node is required to fully explore the space of solutions, and introduce an approach that does so for *stochastic* models.

3 Preliminaries

3.1 DP-UCT

Keller and Helmert (2013) introduce the *Trial-Based Heuristic Tree Search* (THTS) framework, that generalises trial-based planning algorithms for (single-objective) Markov Decision Processes (MDPs). In particular this framework can be specialised to give the *Monte-Carlo Tree-Search* (MCTS) algorithm (Browne et al. 2012) and the (UCT) algorithm (Kocsis and Szepesvári 2006), the most used variant of MCTS, which uses *UCB applied to trees* for action selection. THTS builds a search tree from *decision nodes* and *chance nodes* that correspond to state

and state-action pairs in an MDP, respectively. Moreover, THTS is defined modularly, with different variants of algorithms being specified using seven functions, including `selectAction`, `backupDecisionNode` and `backupChanceNode`. For completeness we give an overview of THTS and pseudocode in Appendix A in Painter, Lacerda, and Hawes (2020).

Due to the modularity of THTS, it is easy to arrive at new algorithms by altering one or a few of the seven functions. Keller and Helmert (2013) utilise this technique to arrive at DP-UCT, an adaption of the standard UCT algorithm that replaces Monte-Carlo backups with dynamic programming backups, and is shown empirically to outperform UCT in many domains. Broadly speaking, DP-UCT algorithm is split into trials. Each trial traverses the tree until a leaf node is found. Once a leaf node has been reached, all nodes that were visited during the trial are backed up to update their value estimates. `selectAction` runs UCB1 at each decision node and `selectOutcome` samples successor states for each chance node. The `backupDecisionNode` and `backupChanceNode` perform Bellman backups in each node, using its children as the successor states in the backup.

In this work we, specify *Convex Hull Monte-Carlo Tree-Search* under the THTS framework, by describing the adaptations made to DP-UCT, rather than re-describing the standard parts of the algorithm. In particular, Section 4 details how to replace the `backupDecisionNode` and `backupChanceNode` functions in DP-UCT and Section 5 describes how to alter the `selectAction` function.

3.2 Multi-Objective Planning under Uncertainty

We will model MOPU problems using a Multi-Objective Markov Decision Process:

Definition 1. A finite horizon *Multi-Objective Markov Decision Process* (MOMDP) is a tuple $M = \langle S, A, \mathbf{R}, T, \bar{s}, H \rangle$, where: S is a finite set of *states*; A is a finite set of *actions*; $\mathbf{R} : S \times A \rightarrow \mathbb{R}^D$ is a D -dimensional *reward function*, specifying D immediate rewards for taking action a when in state s ; $T : S \times A \times S \rightarrow [0, 1]$ is a *transition function* specifying for each state, action and next state, the probability that the next state occurs given the current state and action; \bar{s} is the *initial state*; and $H \in \mathbb{N}$ is a finite *horizon*.

In MOPU we are concerned with the optimisation of a vector \mathbf{V}^π , the sum of each reward observed over the H time-steps, by an agent following a policy π . A (stochastic) policy π maps histories of visited states \mathcal{H} to distributions over actions: given a history of visited states $h \in \mathcal{H}$, $\pi(h)(a)$ represents the probability of the policy choosing to execute a , given that the history of visited states is h . Given a MOMDP M , a policy π and a state s , we define the random variable representing the cumulative reward attained in H timesteps starting in s and following π :

$$\mathbf{X}^\pi(s) = \left(\sum_{t=1}^H \mathbf{r}_t \mid \pi, s_0 = s \right), \quad (1)$$

where \mathbf{r}_t is the reward observed at timestep t according to \mathbf{R} .

Definition 2. The *value* of a policy π is a function $\mathbf{V}^\pi : S \rightarrow \mathbb{R}^D$ such that:

$$\mathbf{V}^\pi(s) = (V_1^\pi(s), \dots, V_D^\pi(s)) = \mathbb{E}[\mathbf{X}^\pi(s)]. \quad (2)$$

In the remainder of this paper, when $s = \bar{s}$ we will omit it and simply write \mathbf{X}^π and $\mathbf{V}^\pi = (V_1^\pi, \dots, V_D^\pi)$. For a set of policies Π , we define the set of points $\hat{\mathbf{V}}(\Pi) \subseteq \mathbb{R}^D$ such that $\hat{\mathbf{V}}(\Pi) = \{\mathbf{V}^\pi \mid \pi \in \Pi\}$. Often there will be no single best policy that can be chosen, as one may encounter a situation where we have $V_i^\pi > V_i^{\pi'}$, but, $V_j^\pi < V_j^{\pi'}$ for some policies π, π' and $i, j \in \{1, \dots, D\}$. However, we can define a partial ordering over multi-objective values:

Definition 3. We say that a point $\mathbf{u} \in \mathbb{R}^D$ *weakly Pareto dominates* another point $\mathbf{v} \in \mathbb{R}^D$ (denoted $\mathbf{u} \succeq \mathbf{v}$), if $\forall i. u_i \geq v_i$; \mathbf{u} *Pareto dominates* \mathbf{v} (denoted $\mathbf{u} \succ \mathbf{v}$), if $\mathbf{u} \succeq \mathbf{v} \wedge \mathbf{u} \neq \mathbf{v}$; \mathbf{u} is *incomparable* to \mathbf{v} (denoted $\mathbf{u} \parallel \mathbf{v}$), if $\mathbf{u} \not\succeq \mathbf{v} \wedge \mathbf{v} \not\succeq \mathbf{u}$.

Given Pareto domination as partial ordering over multi-objective values, we can now define an ‘optimal set’ of policies, in which no policy Pareto dominates another.

Definition 4. Let Π be a set of policies. The *Pareto Front* $PF(\Pi)$ is the set of policies in Π that are not Pareto dominated by any other policy in Π :

$$PF(\Pi) = \{\pi \in \Pi \mid \forall \pi' \in \Pi. \mathbf{V}^\pi \not\succeq \mathbf{V}^{\pi'}\}. \quad (3)$$

An alternative method to overcome the lack of a total order for multi-objective values is to project them onto some scalar value that can be compared. First, we define the set $\mathcal{W}_D \subset \mathbb{R}^D$ of *normalised D -dimensional weights* as:

$$\mathcal{W}_D = \{\mathbf{w} = (w_1, \dots, w_D) \mid \forall i. w_i \geq 0 \wedge \sum_{i=1}^D w_i = 1\}. \quad (4)$$

Definition 5. A *scalarisation function* $f(\mathbf{V}; \mathbf{w})$ maps a multi-objective value into a scalar value, where $\mathbf{w} \in \mathcal{W}_D$. The *linear scalarization* is $f_\ell(\mathbf{v}; \mathbf{w}) = \mathbf{w}^\top \mathbf{v}$, which is (strictly) *monotonic* (i.e. $\mathbf{u} \succ \mathbf{v}$ if and only if $\forall \mathbf{w} \in \mathcal{W}_D. f(\mathbf{u}; \mathbf{w}) > f(\mathbf{v}; \mathbf{w})$).

A weight $\mathbf{w} \in \mathcal{W}_D$ and the linear scalarisation function gives a total ordering over policies, as π, π' can be compared by comparing the scalar values $\mathbf{w}^\top \mathbf{V}^\pi$ to $\mathbf{w}^\top \mathbf{V}^{\pi'}$. Therefore, we can define another optimal set of policies as those that are optimal for some weight vector:

Definition 6. The *Convex Hull* $CH(\Pi)$ is the set of policies in Π that are optimal for some linear scalarization.

$$CH(\Pi) = \{\pi \in \Pi \mid \exists \mathbf{w}. \forall \pi' \in \Pi. \mathbf{w}^\top \mathbf{V}^\pi \geq \mathbf{w}^\top \mathbf{V}^{\pi'}\}. \quad (5)$$

If Π is defined as the set of stochastic policies then $PF(\Pi) = CH(\Pi)$. In order to represent convex hulls and Pareto fronts of Π more compactly, we consider the notion of *Convex Coverage Set* of Π .

Definition 7. A *Convex Coverage Set* $CCS(\Pi)$ is any set such that:

$$\hat{\mathbf{V}}(CCS(\Pi)) = \hat{\mathbf{V}}(CH(\Pi)). \quad (6)$$

For computational reasons, one typically wants to maintain a minimal set of policies that is still a $CCS(\Pi)$. We refer the reader to (Roijers et al. 2013) for more details on the relation between $PF(\Pi)$, $CH(\Pi)$ and $CCS(\Pi)$.

To compare sets of points it is common to consider the *hypervolume*, due to its monotonicity with respect to Pareto domination (Vamplew et al. 2011).

Definition 8. The *hypervolume* of a set of points $\hat{\mathbf{P}} \subseteq \mathbb{R}^D$ can be defined with respect to a reference point \mathbf{o} as follows:

$$HV(\hat{\mathbf{P}}, \mathbf{o}) = \lambda_D(\{\mathbf{v} \in \mathbb{R}^D \mid \exists \mathbf{p} \in \hat{\mathbf{P}}. \mathbf{p} \succeq \mathbf{v} \succeq \mathbf{o}\}), \quad (7)$$

where λ_D is the D -dimensional Lebesgue measure. For example, in two dimensions with $\mathbf{o} = (0, 0)$, this equates to the area between $\hat{\mathbf{P}}$ and the x and y axes.

3.3 Convex Hull Value Iteration

Barrett and Narayanan (2008) proposed an algorithm that extends Value Iteration (Bellman 1957) to handle multiple objectives, named *Convex Hull Value Iteration* (CHVI). CHVI computes $\hat{\mathbf{V}}(PF(\Pi))$ for an infinite-horizon MOMDP by computing a finite set of deterministic policies Π_d that is a $CCS(\Pi)$, where Π is the set of stochastic policies. The same algorithm can be adapted to compute a $CCS(\Pi)$ for a finite-horizon MOMDP, by computing the values for each timestep in the horizon.

As the first step to arriving at CHVI, we define arithmetic rules over a set of points. For sets of points $\hat{\mathbf{P}}, \hat{\mathbf{P}}' \subseteq \mathbb{R}^D$, point $\mathbf{b} \in \mathbb{R}^D$, and scalar $k \in \mathbb{R}$, we define:

$$\hat{\mathbf{P}} + \hat{\mathbf{P}}' = \{\mathbf{p} + \mathbf{p}' \mid \mathbf{p} \in \hat{\mathbf{P}}, \mathbf{p}' \in \hat{\mathbf{P}}'\}, \quad (8)$$

$$\mathbf{b} + k\hat{\mathbf{P}} = \{\mathbf{b} + k\mathbf{p} \mid \mathbf{p} \in \hat{\mathbf{P}}\}. \quad (9)$$

Then, to arrive at the CHVI algorithm, we can replace the fixed point equations used in Value Iteration as follows:

$$\hat{\mathbf{V}}(s) = \begin{cases} \hat{\mathbf{o}} & \text{if } s \text{ terminal,} \\ \text{prune}(\hat{\mathbf{Q}}(s, a)) & \text{otherwise,} \end{cases} \quad (10)$$

$$\hat{\mathbf{Q}}(s, a) = \mathbb{E}[\mathbf{R}(s, a) + \hat{\mathbf{V}}(s') \mid s, a], \quad (11)$$

where $\hat{\mathbf{V}}(s)$ and $\hat{\mathbf{Q}}(s, a)$ are sets of points, $\hat{\mathbf{o}} = \{(0, \dots, 0)\} \subset \mathbb{R}^D$ and *prune* is an operation that removes Pareto-dominated points from a set. For details on how to compute *prune* see the survey by Roijers et al. (2013). The expectation is taken with respect to the next state $s' \sim T(s, a, \cdot)$, and can be computed using equations (8) and (9) in the standard definition of expectation for a discrete random variable.

3.4 Contextual Regret

The *Contextual Bandit* framework extends the standard *Multi-Armed Bandit* (MAB) problem (Auer, Cesa-Bianchi, and Fischer 2002). In this framework we are given a context $x_k \in X$ in round k and have to pick an arm $y_k \in Y$. Subsequently a scalar reward $r_k \sim U(x_k, y_k)$ is received, where $U(x_k, y_k)$ is a reward distribution for arm y_k in context x_k , with expected value $\mu(x_k, y_k)$. We denote the shared context-arm space as $\mathcal{P} \subseteq X \times Y$, and denote the optimal value of context $x \in X$ as $\mu^*(x) = \sup_{\{y \in Y \mid (x, y) \in \mathcal{P}\}} \mu(x, y)$.

Definition 9. The *contextual regret* for N many rounds, is defined as follows:

$$\text{reg}(N) = \sum_{k=1}^N \mu^*(x_k) - \mu(x_k, y_k), \quad (12)$$

where y_k is the arm chosen at round k .

Typically in contextual MABs, the aim is to maximise the *total payoff* of $\sum_{k=1}^N r_k$, which corresponds to minimizing the contextual regret. Algorithms usually aim to achieve *sub-linear regret*, i.e. pick arms such that $\text{reg}(N) \in o(N)$, where $f(x) \in o(x)$ if $\frac{f(x)}{x} \rightarrow 0$ as $x \rightarrow \infty$.

As the problem is very general, works often make additional assumptions, or assumptions of additional knowledge to make the problem more tractable. In particular, one assumption made by Slivkins (2014) is having access to a metric space (\mathcal{P}, d) called the *similarity space*, such that the following Lipschitz condition holds:

$$|\mu(x, y) - \mu(x', y')| \leq d((x, y), (x', y')). \quad (13)$$

This allows us to reason about the expected values of contexts that have not been seen in the past.

4 Convex Hull Monte-Carlo Tree-Search

To motivate the design of our tree-search algorithm, we consider an example MOMDP with a horizon of length three. We will see that any Monte-Carlo tree-search algorithm that maintains a *sample average* at each node, i.e. the average reward obtained from each action, will perform suboptimally.

Example 1. Consider state s_0 and actions a_1, a_2 and a_3 in the MOMDP in Fig. 1 with a horizon of 2. It is clearly optimal to take action a_3 in s_0 for both objectives (and mixes thereof), and the true Pareto front is $F^* = \{(0, 6), (6, 0)\}$. The sample averages for a_1, a_2 and a_3 are $\bar{Q}(s_0, a_1) = (0, 4)$, $\bar{Q}(s_0, a_2) = (4, 0)$ and $\bar{Q}(s_0, a_3) = (6\lambda, 6(1 - \lambda))$, for some $\lambda \in [0, 1]$ (note that choosing action a_3 from s_0 will have a return of either $(6, 0)$ or $(0, 6)$, and $(6\lambda, 6(1 - \lambda))$, $\lambda \in [0, 1]$ represents a weighted average of the two possible returns). If one used these sample averages, the estimated Pareto front at s_0 for a given λ would be $F(\lambda) = \text{prune}(\{(0, 4), (4, 0), (6\lambda, 6(1 - \lambda))\}) \neq F^*$. Therefore, if we try to use the sample averages set $F(\lambda)$ to extract a policy, then it may be suboptimal. For example, if $\lambda = \frac{1}{2}$ then $\bar{Q}(s_0, a_3) = (3, 3)$ and the policy π extracted to maximise only the second objective would have $\pi(s_0) = a_1$, since $\bar{Q}(s_0, a_1) = (0, 4)$, i.e. the sample average for the second objective is higher for a_1 than it is for a_3 .

Given Example 1, it is clear that we need something more than just a single point in each node, so we will consider sets of points to approximate a Pareto front at each node. In the remainder of this section, we describe the backup functions that can be used as part of the THTS schema (Keller and Helmert 2013). Any algorithm that makes use of these backups in a Monte-Carlo Tree-Search we will refer to as a *Convex Hull Monte-Carlo Tree-Search* (CHMCTS).




Figure 1: A deterministic MOMDP with six states, two-dimensional rewards, and initial state s_0 . As it is deterministic, all transition probabilities are one, and so are omitted.

4.1 Backup Functions

For every THTS decision node, corresponding to some $s \in S$, we store a set of points approximating $\hat{V}(s)$, and for every chance node, corresponding to some $(s, a) \in S \times A$, we store a set of points approximating $\hat{Q}(s, a)$. The `backupDecisionNode` function updates the set approximating $\hat{V}(s)$ using Equation (10), but replacing each $\hat{Q}(s, a)$ with its approximation stored in the corresponding child chance node. Similarly, `backupChanceNode` updates the approximation of $\hat{Q}(s, a)$ using Equation (11), replacing each $\hat{V}(s')$ with the approximation stored in the corresponding child decision node.

5 Action Selection

Now we consider how to select actions from decision nodes (recalling decision nodes correspond to states in an MOMDP). We present the problem of policy selection (i.e. selecting all of the actions for a trial) framed as a Contextual Multi-Armed Bandit problem (Section 3.4):

Definition 10. (*Linear*) *Contextual Policy Selection* problem, is a special case of the Contextual Bandit problem. Let M be a MOMDP and Π the corresponding set of policies. For N rounds (or *trials*), we perform the follow sequence of operations: (1) receive a context $w_k \in \mathcal{W}_D$; (2) select a policy π_k to follow for this trial; (3) receive a cumulative reward $x^{\pi_k} \in \mathbb{R}^D$, where $x^{\pi_k} \sim X^{\pi_k}$ as defined in Equation (1) – recall that $V^{\pi_k} = \mathbb{E}[X^{\pi_k}]$. The objective over N rounds is to select a sequence of policies $\{\pi_k\}_{k=1}^N$ that maximise the *expected cumulative payoff* $\sum_{k=1}^N w_k^\top V^{\pi_k}$.

5.1 Design by Regret

In THTS, designing action selection using regret metrics typically has two main benefits it is a direct measure of the *online* performance of the action selection; and it is a good way to balance the exploration-exploitation trade-off, as it selects arms proportionally to how good the performance of each arm has been in the past. Previous multi-objective works (Drugan and Nowe 2013; Chen and Lie 2019) have considered the notion of Pareto regret:

Definition 11. The *Pareto Suboptimality Gap* (PSG) for selecting policy π is defined as:

$$\Delta(\pi) = \inf\{\epsilon \in \mathbb{R} \mid \forall \pi' \in \Pi. V^{\pi'} \not\leq V^\pi + \epsilon \mathbf{1}\}, \quad (14)$$

where $\mathbf{1}$ is a vector of ones. Intuitively, we can think of $\Delta(\pi)$ as how much needs to be added to \mathbf{V}^π so that it is *Pareto optimal*, i.e. it is not dominated by any other policy. Let N_π be the number of times that π was selected for the N trials. The *Pareto regret* is defined as:

$$\text{reg}_P(N) = \sum_{\pi \in \Pi} N_\pi \Delta(\pi). \quad (15)$$

To demonstrate why Pareto regret is not the most suitable regret metric, we consider Example 2, which demonstrates that optimising for the Pareto regret does not correspond to our objective of computing the CCS for a MOMDP.

Example 2. Consider the algorithm that uses the (single-objective) UCT algorithm (Kocsis and Szepesvári 2006) to find $\pi_i^* = \arg \max_{\pi} \mathbf{V}_i^\pi$, the optimal policy for the i th objective and then continues to follow π_i^* . UCT is known to have sublinear regret for the i th objective and, because π_i^* is Pareto-optimal, it also has sublinear Pareto regret. However, this algorithm does not align well with the objective of computing a CCS, as it focuses only on one Pareto-optimal policy and does not explore the rest of the CCS.

Following from this, we introduce *Linear Contextual Regret*, a special case of *contextual regret*, that is well correlated with approximating the CCS.

Definition 12. Let $\{\mathbf{w}_k\}_{k=1}^N$ be a *context weight vector*, i.e. a sequence of weights sampled uniformly from \mathcal{W}_D . The *Linear Contextual Regret* (LCR) for the policy selection $\{\pi_k\}_{k=1}^N$ is defined by:

$$\text{reg}_C(N, \{\mathbf{w}_k\}_{k=1}^N) = \sum_{k=1}^N \max_{\pi' \in \Pi} \mathbf{w}_k^\top \mathbf{V}^{\pi'} - \mathbf{w}_k^\top \mathbf{V}^{\pi_k}, \quad (16)$$

We consider LCR because it will penalise any algorithm that cannot find a *CCS*(Π): if there exists some $\mathbf{V}^\pi \in \hat{\mathbf{V}}(PF(\Pi))$ such that no policy π' with $\mathbf{V}^{\pi'} = \mathbf{V}^\pi$ was found, then a weight close to $\mathbf{w}^* = \arg \max_{\mathbf{w}} \mathbf{w}^\top \mathbf{V}^\pi$ is sampled, it will accumulate regret.

Recall from Definition 10, we aim to maximise expected cumulative payoff, which, is equivalent to minimizing the expected LCR. Note that for any algorithm that achieves a sublinear LCR, the average regret will tend to zero, and in the limit of the number of trials the algorithm must almost surely act optimally for all weight vectors.

5.2 Exploration Policies

We now formally define exploration policies:

Definition 13. An *exploration policy* for MOMDP M is a function of the form $\sigma : \mathcal{H} \times \mathcal{W}_D \rightarrow A$.

In essence, an exploration policy maps a history and a context weight vector to an action, i.e. it extends a policy to consider the context. Algorithms for the contextual policy selection problem must specify a sequence of exploration policies $\{\sigma_k\}_{k=1}^N$, where the MOMDP policy followed on trial k will be $\pi_k = \sigma(\cdot; \mathbf{w}_k)$. The set of exploration policies Π_e for a MOMDP can be divided into two broad classes:

Definition 14. A *context-free* exploration policy σ is one such that $\forall \mathbf{u}, \mathbf{w} \in \mathcal{W}_D, \sigma(\cdot; \mathbf{u}) = \sigma(\cdot; \mathbf{w})$. An exploration policy is *context-aware* if it is not context-free.




Figure 2: A deterministic MOMDP with three states, two-dimensional rewards and a horizon of one.




Figure 3: Visualisation of a snapshot of the CZ algorithm, running over three actions with two objectives, where w_0 is the context weighting for one objective.

Theorem 1 shows any sequence of context-free exploration policies can suffer linear LCR.

Theorem 1. For some MOMDP M , for any sequence of context-free policies $\{\sigma_k\}_{k=1}^N$, the expected LCR over N trials is $\Omega(N)$.

Proof Outline. (See Appendix B in Painter, Lacerda, and Hawes (2020) for full proof.) Consider the MOMDP from Fig. 2. If we follow a context-free exploration policy σ_k on the k th trial, then the action selected at s_0 is independent of the weight context vector. Assume, wlog, that $\sigma_t(s_0; \mathbf{w}) = a_1$ for all \mathbf{w} . As \mathbf{w} is sampled uniformly from \mathcal{W}_D , a_1 will be the suboptimal action with probability $\frac{1}{2}$. If a_1 is suboptimal, the expected regret suffered is $\frac{1}{2}$, because if the weight vector is varied from $(\frac{1}{2}, \frac{1}{2})$ to $(1, 0)$, then the contextual regret suffered varies from 0 to 1. On average, the context-free policies will suffer and expected regret of $\frac{1}{4}$ per trial, and thus the cumulative regret over N is $\Omega(N)$.

5.3 Context-Aware Action Selection

Extending UCB1 to handle a contextual MAB can be hard, as we have an uncountably infinite set of contexts, the weight vectors \mathcal{W}_D . If one has knowledge of a metric that satisfies Equation (13), the problem can be made more tractable. The metric allows contexts to be grouped (in sets with a fixed radius, i.e. balls), and maintain an average value over all contexts in the ball. Intuitively, smaller balls allow each context vector to have a more accurate value estimate maintained. These ideas underline the *Contextual Zooming* (CZ) algorithm (Slivkins 2014), that modifies UCB1 to run over balls of contexts, and introduces balls of smaller radii as required.

We present CZ by defining the *similarity space* (i.e. the metric over the context-arm space be $\mathcal{P} = \mathcal{W}_D \times \Pi$ satisfying Equation (13)); presenting CZ for *policy selection*; and then adapting CZ to be used in `selectAction` from THTS. A snapshot of CZ is visualised in Fig. 3, with three actions. The initial three balls of radii one are blue, and balls of radii $\frac{1}{2}$

are green. a_1 has been covered by green balls, and will only use value estimates from these green balls, whereas a_3 will use the value of the one green ball that it has, *if it is relevant*, and otherwise use the blue ball.

Similarity Space. Consider the linear contextual policy selection problem and let $d : \mathcal{W}_D \times \Pi \rightarrow \mathbb{R}_{\geq 0}$ such that:

$$d((w, \pi), (w', \pi')) = \begin{cases} C(\pi) \|w - w'\|_\infty & \text{if } \pi = \pi', \\ U & \text{otherwise,} \end{cases} \quad (17)$$

where $U > 0$ and $2U \geq C(\pi) \geq \|V^\pi\|_\infty \geq 0$, i.e. $C(\pi)$ is a value that overestimates the infinity norm of the expected vector reward. Furthermore, U is an upper bound on the maximum scalarised reward that can be achieved in a single trial, i.e. $U \geq \sup_{w \in \mathcal{W}_D, \pi \in \Pi} w^\top V^\pi$. For example, if $R(s, a) \in [0, 1]^D$ then these values can be set to $C(\pi) = U = H$. Additionally, observe the Lipschitz property (Equation (13)) holds when $\pi \neq \pi'$ because $|\mu(w, \pi) - \mu(w', \pi)| \leq U$ and it holds when $\pi = \pi'$ because:

$$|\mu(w, \pi) - \mu(w', \pi)| = \|(w^\top - w'^\top) V^\pi\|_\infty \quad (18)$$

$$\leq \|V^\pi\|_\infty \|w^\top - w'^\top\|_\infty \quad (19)$$

$$\leq C(\pi) \|w^\top - w'^\top\|_\infty \quad (20)$$

$$= d((w, \pi), (w', \pi)), \quad (21)$$

where in the first line we use the definition of μ and the fact that the modulus and infinity-norm operations are identical on scalar values. We also used the result that for any matrices $\|AB\|_\infty \leq \|A\|_\infty \|B\|_\infty$.

Contextual Zooming. CZ is an algorithm that achieves a contextual regret of $O(N^{1-1/(2+c)} \log N)$, where c is the *covering dimension*. The *covering dimension* is related to how many balls of a fixed radius are needed to cover the similarity space. For a full explanation and derivation of the regret bound we refer the reader to (Slivkins 2014).

Throughout the algorithm a finite set of balls called the set of *active balls* is maintained. Let \mathcal{A}_k be the set of active balls at the start of trial k . For us, each ball corresponds to a set of context vectors and has an associated arm (i.e. policy). Whenever CZ needs to select an arm, it will find a set of *relevant balls* in \mathcal{A}_k , compute an *upper confidence bound* for each relevant ball and select the arm associated with the largest bound. The ball B is *relevant* for the weight vector $w \in \mathcal{W}_D$ if there is some arm $\pi \in \Pi$ such that $(w, \pi) \in \text{dom}_k(B)$, where:

$$\text{dom}_k(B) = B - \left(\bigcup_{\{B' \in \mathcal{A}_k \mid r(B') < r(B)\}} B' \right), \quad (22)$$

with $r(B)$ the radius of B . $\text{dom}_k(B)$ is used to decide which balls are relevant to consider when making a choice for some context vector. The upper confidence bound $I_k(B)$ for each B in the set of relevant balls during round k is defined using the following equations:

$$\text{conf}_k(B) = 4\sqrt{\frac{\log k}{1 + n_k(B)}}, \quad (23)$$

$$I_k^{\text{pre}}(B) = \nu_k(B) + r(B) + \text{conf}_k(B), \quad (24)$$

$$I_k(B) = r(B) + \min_{B' \in \mathcal{A}_k} (I_k^{\text{pre}}(B') + \mathcal{D}(B, B')), \quad (25)$$

where $n_k(B)$ is the number of times that ball B has been selected in the previous $k - 1$ rounds, $\nu_k(B)$ is the average (scalarised) reward for ball B in the previous $k - 1$ rounds, $r(B)$ is the radius of the ball and $\mathcal{D}(B, B')$ is the distance between the centers of balls B and B' .

The algorithm then proceeds by repetitively applying the following two rules. *Selection rule:* On round k select the ball B , from those that are relevant, that has the maximum index $I_k(B)$. From that ball, select an arm arbitrarily such that $(w_k, \pi) \in B$. *Activation Rule.* If the ball B that was selected satisfies $\text{conf}_k(B) \leq r(B)$ after this round, then a new ball is added to the active set, otherwise we set $\mathcal{A}_{k+1} = \mathcal{A}_k$. When adding a new ball, if π was the arm selected by the selection rule in round k , then a new ball B_{new} with center (w_k, π) is introduced with radius $\frac{1}{2}r(B)$, and we set $\mathcal{A}_{k+1} = \mathcal{A}_k \cup \{B_{\text{new}}\}$. To initialise the set of active balls we add one ball per arm π , with center (\bar{w}, π) and radius $C(\pi)$ to \mathcal{A}_0 , where $\bar{w} = \frac{1}{d}1$.

Contextual Zooming for Trees. Running CZ directly over policies is infeasible for MOMDPs, because the number of policies grows exponentially in the size of the state space, and additionally the distance metric does not allow two different policies to be close in the similarity space. Instead, we run CZ for the `selectAction` method in THTS at every decision node. So now we consider a new similarity space (\mathcal{P}_s, d_s) associated with the state s of the decision node, where $\mathcal{P}_s = \mathcal{W}_D \times A$, and d_s is defined as:

$$d_s((w, a), (w', a')) = \begin{cases} C_s(a) \|w - w'\|_\infty & \text{if } a = a', \\ U & \text{otherwise,} \end{cases} \quad (26)$$

where $C_s(a) = \sup_{\{\pi \mid \pi(s)=a\}} C(\pi)$. Given this similarity space, the CZ algorithm operates as before, using actions for the arms instead of policies in the contextual MAB problem. Note that for each decision node we, in fact, have a *non-stationary* contextual multi-armed bandits problem, similar to (Kocsis, Szepesvári, and Willemson 2006). We refer to this action selection method as *Contextual Zooming for Trees* (CZT), and when we use CZT for action selection in CHMCTS we call the algorithm *Zooming CHMCTS*.

6 Experiments

To validate Zooming CHMCTS we will evaluate its performance on a variable-sized grid world problem, the *Generalised Deep Sea Treasure* (GDST) problem (Vamplew et al. 2011; 2017). Moreover, we will consider its performance in both online and offline planning, where in online planning we assume that the agent follows each policy selected, and we want to maximise the cumulative payoff over many trials. In contrast, in offline planning, we do not care about the performance during the planning phase, but only the quality of policies that can be extracted afterwards.

6.1 Experimental Setup

In the GDST(c, p) problem we consider a two-dimensional grid world, consisting of c columns and a transition noise of $p \in [0, 1]$. The submarine can move left, right, up or down each time step, with the submarine remaining stationary if it would otherwise leave the grid world. The submarine starts in




Figure 4: A visualisation of a GDST(5,p) problem, gray squares represent the treasure that can be obtained.

the top left corner on each trial. On every timestep, transition noise p indicates the probability that the submarine is instead swept by a current, moving it in a random direction. The seafloor becomes increasingly deep from left to right (with depth increasing by a small amount between zero and three inclusive for each column), but also holds increasing amounts of treasure at greater depths, ranging in values from one to 1000. There are two objectives, one is to collect the maximal amount of treasure, while the other is to minimise the time cost of reaching the treasure, where a cost of one is incurred for each timestep. Each trial concludes after either $100c$ steps or as soon as the submarine arrives at a piece of treasure. A visualisation of an environment for $c = 5$ is given in Fig. 4. During planning, we normalise rewards to the range $[0, 1]$ to give equal priority over the different objectives.

6.2 Practical Considerations

Because the CHVI backups from Equations (10) and (11) can be computationally intensive, optimisations can be made to minimise the number of these backup operations performed. Where appropriate we use these following optimisations.

Firstly, Lizotte, Bowling, and Murphy; Lizotte, Bowling, and Murphy (2010; 2012) demonstrate that performing the relevant computations in *weight space*, rather than *value space*, is computationally more efficient. We additionally follow the *backup optimisation* considered by Perez et al. (2015). We can identify when either a `backupChanceNode` or a `backupDecisionNode` does not change the value at a node, and subsequently prune the remaining backups for the trial. Also, it is possible to incorporate the ideas from UCD Saffidine, Cazenave, and Méhat (2012) that allow nodes to be re-used when a state is re-visited. Finally, for offline planning, we can use the idea of *labelling* nodes, where a node is labelled if its value has converged. Nodes can be labelled when they are backed up and all of their children are labelled. Leaf nodes are always labelled. By using labelling, we can avoid searching parts of the tree that have already converged.

6.3 Alternative Action Selection Methods

We now state some other action selection methods used in the literature. Let s be the state of the decision node that an action is being selected at, which has been visited $N(s)$ times. For any action a , let $N(s, a)$ be the number of times that the corresponding child node has been visited and let $\hat{Q}(s, a)$




Figure 5: The LCR $\text{reg}_U(N)$ with respect to the number of trials N of different CHMCTS algorithms. Each algorithm was run five times.

denote the Pareto front stored in the child node. Most prior works use an exploration policy of the form:

$$\nu(s; \mathbf{w}) = \arg \max_a \zeta(\hat{Q}(s, a); \mathbf{w}) + C \sqrt{\frac{\log(N(s))}{N(s, a)}}, \quad (27)$$

where C is some appropriate constant and ζ is some scalarization function that maps a Pareto front to a scalar value. Perez, Samothrakis, and Lucas; Perez et al.; Perez-Liebana, Mostaghim, and Lucas (2013; 2015; 2016) set

$$\zeta(\hat{Q}(s, a); \mathbf{w}) = \frac{HV(\hat{Q}(s, a), \mathbf{o})}{N(s)}, \quad (28)$$

where HV is the hypervolume function (Definition 8). This action selection is *context-free* and we refer to the algorithm that results from this exploration policy as *Hypervolume CHMCTS*. Xu and Li (2017) use the *Chebychev scalarization* function instead, setting:

$$\zeta(\hat{Q}(s, a); \mathbf{w}) = \min_{\mathbf{p} \in \hat{Q}(s, a)} \max_i w_i |p_i - z_i|, \quad (29)$$

where $\mathbf{z} = (z_1, \dots, z_D)$ is called the *utopian point*, defined by $z_i = \max_{\pi} V_i^{\pi}$. When we use this action selection we call the algorithm *Chebychev CHMCTS*. This action selection is actually *context-aware*, and Xu and Li (2017) select the weight vector \mathbf{w} randomly each trial.

The ParetoUCB1 algorithm is used by Chen and Lie (2019) for action selection in a MOMAB problem. In our approach, we run ParetoUCB1 over the set $\bigcup_{a \in A} \hat{Q}(s, a)$, rather than using vectors from child nodes as in point-based MOMCTS. When using this action selection, we refer to our algorithm as *Pareto CHMCTS*.

6.4 Regret Comparison

In this section, we consider a GDST(7,0.01) instance, which is small enough for us to feasibly compute the true CCS, and thus the regret is computable. Fig. 5 shows a plot of the cumulative LCR over 100,000 trials, on the GDST(7,0.01) instance.




Figure 6: The hypervolume with respect to the number of CHVI operations. We compare each variant of CHMCTS with the baseline CHVI. 95% Confidence intervals are plotted from 5 runs each.

As we can see in the figure, only Zooming CHMCTS manages to achieve a sublinear regret. These results are consistent with Theorem 1, with context-free action selection methods suffering a linear regret. Interestingly, Chebychev CHMCTS suffers a linear regret too, despite being context-aware. From these results, we can hypothesise that a regret bound could be found for CZT. This demonstrates that Zooming CHMCTS outperforms all other variants for online planning (i.e. in the contextual policy selection problem).

6.5 Comparing Action Selections

To compare CHMCTS algorithms for offline planning, we compare the hypervolume of the CCS estimated in the root node, with respect to the number of backups required. Hypervolume is considered a suitable metric for the quality of a CCS (Vamplew et al. 2011). We compare algorithms on a GDST(30,0.01) environment, using CHVI as a baseline in Fig. 6. We note that CHVI will compute the optimal CCS in $|S| \cdot H$ many backups if S is the state space and H the horizon. $|S| \cdot H$ is the least number of backups required to guarantee that the optimal value has been computed, so we expect CHVI to converge first. We see that Zooming CHMCTS continues to make steady progress towards the optimal hypervolume, while the other methods appear to quickly plateau. Also, we can see that CHMCTS algorithm would outperform CHVI if given a small computational budget. In the next section, we will also see that as the environment size is increased the performance of CHVI severely deteriorates.

6.6 Scalability Analysis

To understand how scalable CHMCTS is, we compare Zooming CHMCTS with CHVI on a range of different sized environments. In Fig. 7 we estimate how much of the true CCS is discovered given a budget of 25000 backups. We vary c between three and 80, for both $p = 0$ and $p = 0.01$.

Let $h\nu(c, p)$ denote the optimal hypervolume of a




Figure 7: The approximate ratio of the hypervolume found, $\frac{eh\nu(c, p)}{h\nu(c, 0)}$, with respect to problem size, i.e. c , with a budget of 25000 backups. 95% Confidence intervals are plotted from 3 runs each.

GDST(c, p) instance, and let $eh\nu(c, p)$ be the estimated hypervolume resulting from some algorithm on the same GDST(c, p) instance. The ratio of $\frac{eh\nu(c, p)}{h\nu(c, 0)}$ indicates what proportion of the CCS was found by the algorithm. However, $h\nu(c, p)$ is infeasible to compute for $p > 0$. Therefore we instead plot the lower bound $\frac{eh\nu(c, p)}{h\nu(c, 0)} \leq \frac{eh\nu(c, p)}{h\nu(c, p)}$, which for GDST instances will always be in the range $[0, 1]$.

In Fig. 7 we see that CHMCTS outperforms CHVI on instances with $c \geq 40$. We only compare Zooming CHMCTS because of computational constraints. Our results suggest that in the presence of a budget on the number of backups, there will be some threshold in the sizes of MOMDPs from which CHMCTS outperforms CHVI.

7 Conclusion

In this work, we posed planning in Multi-Objective Markov Decision Processes as a *contextual multi-armed bandits problem*. We then discussed why one should maintain approximations of the Pareto front in each tree node, to produce the Convex Hull Multi-Objective Tree-Search framework. By considering contextual regret, we introduced a novel action selection method and empirically verified it can achieve a sublinear regret and outperforms other state-of-the-art approaches.

Future work includes proving the relevant regret bounds and applying these algorithms in larger, real-world settings.

Acknowledgements

This work is funded by Bossa Nova Robotics, Inc located at 610 22nd Street, Suite 250, San Francisco, CA 94107, USA, UK Research and Innovation and EPSRC through the Robotics and Artificial Intelligence for Nuclear (RAIN) research hub [EP/R026084/1], and the European Union Horizon 2020 research and innovation programme under grant agreement No 821988 (ADE).

References

- Abels, A.; Roijers, D. M.; Lenaerts, T.; Nowé, A.; and Steckelmacher, D. 2018. Dynamic Weights in Multi-Objective Deep Reinforcement Learning. *arXiv preprint arXiv:1809.07803*.
- Auer, P.; Cesa-Bianchi, N.; and Fischer, P. 2002. Finite-time analysis of the multiarmed bandit problem. *Machine learning* 47(2-3):235–256.
- Barrett, L., and Narayanan, S. 2008. Learning all optimal policies with multiple criteria. In *Proceedings of the 25th international conference on Machine learning*, 41–47. ACM.
- Bellman, R. 1957. *Dynamic Programming*. Princeton, NJ, USA: Princeton University Press, 1 edition.
- Browne, C. B.; Powley, E.; Whitehouse, D.; Lucas, S. M.; Cowling, P. I.; Rohlfshagen, P.; Tavener, S.; Perez, D.; Samothrakis, S.; and Colton, S. 2012. A survey of Monte Carlo tree search methods. *IEEE Transactions on Computational Intelligence and AI in games* 4(1):1–43.
- Chen, W., and Lie, L. 2019. Pareto Monte Carlo Tree Search for Multi-Objective Informative Planning.
- Drugan, M. M., and Nowe, A. 2013. Designing multi-objective multi-armed bandits algorithms: A study. In *The 2013 International Joint Conference on Neural Networks (IJCNN)*, 1–8. IEEE.
- Keller, T., and Helmert, M. 2013. Trial-Based Heuristic Tree Search for Finite Horizon MDPs. In *ICAPS*.
- Kocsis, L., and Szepesvári, C. 2006. Bandit based monte-carlo planning. In *European conference on machine learning*, 282–293. Springer.
- Kocsis, L.; Szepesvári, C.; and Willemson, J. 2006. Improved monte-carlo search. *Univ. Tartu, Estonia, Tech. Rep* 1.
- Lacerda, B.; Parker, D.; and Hawes, N. 2017. Multi-objective policy generation for mobile robots under probabilistic time-bounded guarantees. In *Proc. of the 27th Int. Conf on Automated Planning and Scheduling (ICAPS)*.
- Lahijanian, M.; Svorenova, M.; Morye, A. A.; Yeomans, B.; Rao, D.; Posner, I.; Newman, P.; Kress-Gazit, H.; and Kwiatkowska, M. 2018. Resource-performance tradeoff analysis for mobile robots. *IEEE Robotics and Automation Letters* 3(3):1840–1847.
- Lizotte, D. J.; Bowling, M. H.; and Murphy, S. A. 2010. Efficient reinforcement learning with multiple reward functions for randomized controlled trial analysis. In *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*, 695–702. Citeseer.
- Lizotte, D. J.; Bowling, M.; and Murphy, S. A. 2012. Linear fitted-q iteration with multiple reward functions. *Journal of Machine Learning Research* 13(Nov):3253–3295.
- Mossalam, H.; Assael, Y. M.; Roijers, D. M.; and Whiteson, S. 2016. Multi-Objective Deep Reinforcement Learning. *CORR*.
- Painter, M.; Lacerda, B.; and Hawes, N. 2020. Convex Hull Monte-Carlo Tree Search. *arXiv preprint*.
- Perez, D.; Mostaghim, S.; Samothrakis, S.; and Lucas, S. M. 2015. Multiobjective monte carlo tree search for real-time games. *IEEE Transactions on Computational Intelligence and AI in Games* 7(4):347–360.
- Perez-Liebana, D.; Mostaghim, S.; and Lucas, S. M. 2016. Multi-objective tree search approaches for general video game playing. In *2016 IEEE Congress on Evolutionary Computation (CEC)*, 624–631. IEEE.
- Perez, D.; Samothrakis, S.; and Lucas, S. 2013. Online and offline learning in multi-objective Monte Carlo Tree Search. In *IEEE Conference on Computational Intelligence and Games, CIG*, 1–8. IEEE.
- Roijers, D. M.; Vamplew, P.; Whiteson, S.; and Dazeley, R. 2013. A survey of multi-objective sequential decision-making. *Journal of Artificial Intelligence Research* 48:67–113.
- Saffidine, A.; Cazenave, T.; and Méhat, J. 2012. UCD: Upper Confidence bound for rooted Directed acyclic graphs. *Knowledge-Based Systems* 34:26–33.
- Sahoo, N.; Ganguly, S.; and Das, D. 2012. Multi-objective planning of electrical distribution systems incorporating sectionalizing switches and tie-lines using particle swarm optimization. *Swarm and Evolutionary Computation* 3:15–32.
- Silver, D.; Schrittwieser, J.; Simonyan, K.; Antonoglou, I.; Huang, A.; Guez, A.; Hubert, T.; Baker, L.; Lai, M.; Bolton, A.; and Others. 2017. Mastering the game of go without human knowledge. *Nature* 550(7676):354.
- Slivkins, A. 2014. Contextual bandits with similarity information. *The Journal of Machine Learning Research* 15(1):2533–2568.
- Sutton, R. S.; McAllester, D. A.; Singh, S. P.; and Mansour, Y. 2000. Policy gradient methods for reinforcement learning with function approximation. In *Advances in neural information processing systems*, 1057–1063.
- Vamplew, P.; Dazeley, R.; Berry, A.; Issabekov, R.; and Dekker, E. 2011. Empirical evaluation methods for multiobjective reinforcement learning algorithms. *Machine learning* 84(1-2):51–80.
- Vamplew, P.; Webb, D.; Zintgraf, L. M.; Roijers, D. M.; Dazeley, R.; Issabekov, R.; and Dekker, E. 2017. MORL-Glue: A benchmark suite for multi-objective reinforcement learning. In *29th Benelux Conference on Artificial Intelligence November 8–9, 2017, Groningen*, 389.
- Wang, W., and Sebag, M. 2012. Multi-objective monte-carlo tree search. In *Asian conference on machine learning*, volume 25, 507–522.
- Xu, X., and Li, G. 2017. Chebyshev metric based multi-objective Monte Carlo tree search for combat simulations. In *2017 21st International Conference on System Theory, Control and Computing (ICSTCC)*, 607–612. IEEE.