

A Framework for Reinforcement Learning and Planning

Thomas M. Moerland^{1,2}, Joost Broekens², and Catholijn M. Jonker^{1,2}

¹ Interactive Intelligence, TU Delft, The Netherlands

² LIACS, Leiden University, The Netherlands

Contents

1	Introduction	2
2	Background	3
2.1	Markov Decision Process	3
2.2	The distinctive assumption between planning and learning	4
2.3	Planning	5
2.4	Model-free reinforcement learning	6
2.5	Model-based reinforcement learning	8
3	Framework for Reinforcement learning and Planning	9
3.1	Trials and back-ups	10
3.2	Where to put our computational effort?	10
3.3	Where to make the next trial?	12
3.3.1	Candidate set selection	13
3.3.2	Exploration	14
3.3.3	One versus two phase exploration	18
3.3.4	Reverse trials	19
3.4	How to estimate the cumulative return?	19
3.4.1	Sample depth	20
3.4.2	Bootstrap function	21
3.5	How to back-up?	21
3.5.1	Back-up policy	21
3.5.2	Expectation over the actions	22
3.5.3	Expectation over the dynamics	22
3.6	How to represent the solution?	23
3.6.1	Function type	23
3.6.2	Function class and generalization	24
3.7	How to update the solution?	25
3.7.1	Loss	26
3.7.2	Update rule	27
4	Conceptual comparison of well-known algorithms	30
5	Related Work	34
6	Discussion	35
7	Conclusion	37

Abstract

Sequential decision making, commonly formalized as Markov Decision Process optimization, is a key challenge in artificial intelligence. Two successful approaches to MDP optimization are planning and reinforcement learning. Both research fields largely have their own research communities. However, if both research fields solve the same problem, then we should be able to disentangle the common factors in their solution approaches. Therefore, this paper presents a unifying framework for reinforcement learning and planning (FRAP), which identifies the underlying dimensions on which any planning or learning algorithm has to decide. At the end of the paper, we compare - in a single table - a variety of well-known planning, model-free and model-based RL algorithms along the dimensions of our framework, illustrating the validity of the framework. Altogether, FRAP provides deeper insight into the algorithmic space of planning and reinforcement learning, and also suggests new approaches to integration of both fields.

Keywords: Reinforcement learning, planning, model-based reinforcement learning, framework, conceptual overview, survey.

1 Introduction

Sequential decision making is a key challenge in artificial intelligence research. The problem, commonly formalized as a Markov Decision Process (MDP) (Puterman, 2014), has been studied in different research fields. The two prime research directions are *reinforcement learning* (Sutton and Barto, 2018), a subfield of machine learning, and *planning* (also known as *search*), of which the discrete and continuous variants have been studied in the fields of artificial intelligence (Russell and Norvig, 2016) and control (Bertsekas et al., 1995), respectively. Planning and learning approaches differ with respect to a key assumption: is the dynamics model of the environment known (planning) or unknown (reinforcement learning).

Departing from this distinctive assumption, both research fields have largely developed their own methodology, in relatively separated communities. There has been cross-breeding as well, better known as ‘model-based reinforcement learning’ (recently surveyed by Moerland et al. (2020a)). However, a unifying view on both fields, including how their approaches overlap or differ, currently lacks in literature (see Section 5 for an extensive discussion of related work).

Therefore, this paper introduces the Framework for Reinforcement learning and Planning (FRAP), which identifies the essential algorithmic decisions that any planning or RL algorithm has to make. We identify six main questions: 1) where to put our computational effort, 2) where to make our next trial, 3) how to estimate the cumulative return, 4) how to back-up, 5) how to represent the solution and 6) how to update the solution. As we will see, several of these questions have multiple subquestions. However, the crucial message of this paper is that any RL or planning algorithm, from Q-learning (Watkins and Dayan, 1992) to A* (Hart et al., 1968), will have to make a decision on each of these dimensions. We illustrate this point at the end of the paper, by formally comparing a variety of planning and RL papers on the dimensions of our framework.

The framework first of all provides a common language to categorize algorithms in both fields. Too often, we see researchers mention ‘they use a policy gradient algorithm’, while this only specifies the choice on one of the dimensions of our framework and leaves many other algorithmic choices unspecified. Second, the framework identifies new research directions, for example by taking inspiration from solutions in the other research field, or by identifying novel combinations of approaches that are still left untried. Finally, it can also serve an educational purpose, both for researchers and students, to get a more systematic understanding of the common factors in sequential decision-making problems.

The remainder of this article is organized as follows. To keep the document self-contained, Section 2 provides a short overview of the problem formulation (MDP optimization) and the main solution approaches: planning, model-free reinforcement learning, and model-based reinforcement learning. Experienced readers can skip this section, although we do advise them to quickly read paragraph 2.2, since a systematic categorization of the types of environment access is to our knowledge missing in literature. The main contribution of this work, the framework, is presented in Section 3. Section 4 is the other key contribution of our paper, since it compares various well-known planning and reinforcement learning algorithms along the dimensions of our framework (in Table 3), thereby illustrating the generality of FRAP. We conclude the paper with Related Work (Sec. 5), Discussion (Sec. 6) and Summary (Sec. 7) sections.

2 Background

In sequential decision-making, formalized as Markov Decision Process optimization, we are interested in the following problem: given a (sequence of) state(s), what next action is best to choose, based on the criterion of highest cumulative pay-off in the future. More formally, we aim for *context-dependent action prioritization based on a (discounted) cumulative reward criterion*. This is a core challenge in artificial intelligence research, as it contains the key elements of the world: there is sensory information about the environment (states), we can influence that environment through actions, and there is some notion of what is preferable, now and in the future. The formulation can deal with a wide variety of well-known problem instances, like path planning, robotic manipulation, game playing and autonomous driving.

We will first formally introduce the Markov Decision Process optimization problem, and subsequently introduce the considered solution approaches: planning, model-free RL, and model-based RL.

2.1 Markov Decision Process

The formal definition of a *Markov Decision Process* (MDP) (Puterman, 2014) is the tuple $\{\mathcal{S}, \mathcal{A}, \mathcal{T}, \mathcal{R}, p(s_0), \gamma\}$. The environment consists of a *transition function* $\mathcal{T} : \mathcal{S} \times \mathcal{A} \rightarrow p(\mathcal{S})$ and a *reward function* $\mathcal{R} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}$. At each timestep t we observe some state $s_t \in \mathcal{S}$ and pick an action $a_t \in \mathcal{A}$. Then, the environment returns a next state $s_{t+1} \sim \mathcal{T}(\cdot | s_t, a_t)$ and associated scalar reward $r_t = \mathcal{R}(s_t, a_t, s_{t+1})$. The first state is sampled from the initial state distribution $p(s_0)$. Finally, $\gamma \in [0, 1]$ denotes a discount parameter.

The agent acts in the environment according to a *policy* $\pi : \mathcal{S} \rightarrow p(\mathcal{A})$. In the search community, a policy is also known as a *contingency plan* or *strategy* (Russell and Norvig, 2016). By repeatedly selecting actions and transitioning to a next state, we can sample a *trace* through the environment. The *cumulative return* of trace through the environment is denoted by: $J_t = \sum_{k=0}^K \gamma^k \cdot r_{t+k}$, for a trace of length K . For $K = \infty$ we call this the infinite-horizon return.

Define the action-value function $Q^\pi(s, a)$ as the expectation of the cumulative return given a certain policy π :

$$Q^\pi(s, a) \doteq \mathbb{E}_{\pi, \mathcal{T}} \left[\sum_{k=0}^K \gamma^k r_{t+k} \mid s_t = s, a_t = a \right] \quad (1)$$

This equation can be written in a recursive form, better known as the *Bellman equation*:

$$Q^\pi(s, a) = \mathbb{E}_{s' \sim \mathcal{T}(\cdot | s, a)} \left[\mathcal{R}(s, a, s') + \gamma \mathbb{E}_{a' \sim \pi(\cdot | s')} \left[Q^\pi(s', a') \right] \right] \quad (2)$$

Our goal is to find a policy π that maximizes our expected return $Q^\pi(s, a)$:




Figure 1: Types of access to the environment dynamics. Columns: On each trial, we may either get access to the exact transition probabilities of each possible transition (analytic or descriptive model), or we may only get a sampled next state (sample or generative model). Rows: Additionally, we may either be able to revert the model and make another trial from the same state (reversible), or we may need to continue from the resulting state (irreversible). Planning algorithms assume a reversible environment, RL algorithms assume an irreversible environment. We could theoretically think of an irreversible analytic environment, in which we do see the probabilities of each transition but can only continue from one drawn realization, but we are unaware of such a model in practice.

$$\pi^* = \arg \max_{\pi} Q^{\pi}(s, a) = \arg \max_{\pi} \mathbb{E}_{\pi, \mathcal{T}} \left[\sum_{k=0}^K \gamma^k r_{t+k} \mid s_t = s, a_t = a \right] \quad (3)$$

There is *at least* one optimal policy, denoted by π^* , which is better or equal than all other policies π (Sutton and Barto, 2018). In the planning and search literature, the above problem is typically formulated as a cost *minimization* problem (Russell and Norvig, 2016). That formulation is interchangeable with our presentation by negating the reward function. The formulation also contains *stochastic shortest path* problems (Bertsekas and Tsitsiklis, 1991), which are MDP formulations with absorbing states at goal states, where we attempt to reach the goal with a little cost as possible.

2.2 The distinctive assumption between planning and learning

A key consideration in the planning and learning field is: what access to the MDP (dynamics \mathcal{T} and reward function \mathcal{R}) are we provided with? We identify three ways in which we can get access to the MDP (Figure 1):

- *Reversible analytic* environments specify the entire probability distribution $\mathcal{T}(s'|s, a)$. In Figure 1 top-left, we see an example with three possible next states, where the probability of each state is fully explicitized. Such access allows for exact evaluation of the Bellman equation.
- *Reversible sample* environments provide a single sample from $s' \sim \mathcal{T}(\cdot|s, a)$, but do not give access to the underlying probabilities. In Figure 1, top-right, we sampled the same state-action three times, which gave two times the first and one time the third next state.
- *Irreversible sample* environments also provide a sample, but introduce another restriction: we need to keep sampling forward. In other words, we cannot consider the same

state twice directly after each other. If we want to get back, then we will have to pick the correct actions to bring us back to the specific state. The key example of an irreversible sampler is the real-world, in which we cannot revert time. For many real world problems it is hard to specify an analytic or reversible sample model, but we can always get irreversible sample data by interacting with the real world.

Note that there is an ordering in these access types. We can always decide to sample from an analytic model, and we can always restrict ourselves to never revert the environment. Therefore, the reversible analytic model gives us most information and freedom. On the other hand, sample models are usually easier to obtain, and irreversible sampling is of course an important property of the real world, in which we ultimately want to apply learning.

The key difference between planning and RL is that RL fundamentally limits itself to irreversible sample environments, frequently referred to as an ‘unknown model’. On the other hand, planning always assumes a reversible environment (either analytic or sample), which is usually referred to as a ‘known model’. Throughout this work, we will refer to model as any form of a reversible dynamics function. Departing from this fundamental difference in environment access, both fields have developed their own methods and preferences for solving the MDP optimization problem, which will be covered in the next section.

2.3 Planning

Planning (or *search*) is a large research field within artificial intelligence (Russell and Norvig, 2016). Following Sutton and Barto (2018), we will define planning as: ‘any process that takes a model as input and produces or improves a policy for interacting with the modeled environment’. We shortly list some important planning approaches. This presentation is by no means exhaustive, but it does establish some common ground of algorithms we consider in our framework:

- **Dynamic programming (DP)** (Bellman, 1966; Howard, 1960): The key idea of Dynamic programming is to break the optimization problem into smaller subproblems given by the 1-step optimal Bellman operator. We then sweep through state-space, repeatedly solving the small subproblem which eventually solves for the optimal policy. DP is a bridging technique between both planning and reinforcement learning. However, the tabular implementation does not scale well to high-dimensional problems, since the size of the required table grows exponentially in the dimensionality of the state space (‘the curse of dimensionality’). To solve for this issue, Real-time Dynamic Programming (RTDP) (Barto et al., 1995) only applies DP updates on traces sampled from some start state distribution.
- **Heuristic search:** These search approach built a forward tree from some start state. Initial research largely focused on uninformed search strategies, like breadth-first search (BFS) (Moore, 1959) and Dijkstra’s shortest path algorithm (Dijkstra, 1959). These approaches track a *frontier*, which is the set of nodes that have themselves been visited, but whose successor states have not all been visited yet. Later approaches successfully incorporated heuristics, which are functions that provide an initial *optimistic* estimate of the return from a particular state. A well-known heuristic search algorithm is A^* (Hart et al., 1968). However, for many problems informative heuristics are not trivial to obtain.
- **Sample-based search:** This group of search algorithms estimates state-action values based on statistical sampling methods. The simplest example is Monte Carlo search (MCS) (Tesauro and Galperin, 1997), where we sample n traces for each currently available action and use their mean return as an estimate of the value of that action. A successful extension of this paradigm is Monte Carlo Tree Search (Kocsis and Szepesvári, 2006; Browne et al., 2012). While MCS only tracks statistics at the root of the tree, MCTS recursively applies the same principle at deeper levels of the tree

search. Exploration and exploitation within the tree are typically based on variants of the upper confidence bounds (UCB) rule (Kocsis and Szepesvári, 2006). Pure MCTS for example showed early success in the game of Go (Gelly and Wang, 2006). MCTS originates in regret minimization (Auer, 2002), which attempts to select the optimal action as often as possible during the search. In contrast, *best-arm identification* (BAI) tries to identify the optimal root action at the end of the search (Kaufmann and Koolen, 2017), which allows for additional exploration during the search itself. Finally, in the robotics path planning community there is another successful branch of sample-based planning algorithms known as *rapidly-exploring random trees* (RRTs) (LaValle, 1998). While MCTS samples in action space to build the tree, RRTs sample in state space, which is only feasible if the state-space is not too large.

- **Gradient-based planning:** This planning approach is especially popular in the robotics and control community. If we have a differentiable dynamics models (either pre-known or learned from data), then we can directly obtain the derivative of the cumulative reward objective with respect to the policy parameters by differentiating through the dynamics function. An especially popular approach in this category applies when we have a linear dynamics model and a quadratic reward function. In that case, we can derive closed-form expressions for the optimal action, known as the linear-quadratic regulator (LQR) (Anderson and Moore, 2007). While most practical problems have non-linear dynamics, this problem can be partly mitigated by iterative LQR (iLQR) (Todorov and Li, 2005), which repeatedly makes local linear approximations to the true dynamics. In RL literature, gradient-based planning is referred to as *value gradients* (Heess et al., 2015).
- **Direct optimization:** We may also treat the planning challenge as a black-box optimization problem. This approach is especially popular in the robotics and control community, better known as *direct optimal control* (Bock and Plitt, 1984). In this approach we reformulate the objective as a non-linear programming problem, in which the dynamics typically enter as constraints on the solution. We then parametrize a trajectory (a local policy), and perform hill-climbing in this parameter space, for example based on finite-differencing. In the next section on RL, we will encounter similar ideas known as *policy search*.

Another direction of planning research that has been popularized in the last decade treats planning as probabilistic inference (Botvinick and Toussaint, 2012; Toussaint, 2009; Kappen et al., 2012), where we use message-passing like algorithms to infer which actions would lead to receiving a final reward. Note that we do leave out some planning fields that depart from the generic MDP specification. For example, *classical planning* (Ghallab et al., 1998) requires a propositional logic structure of the state space. Approaches in this field may plan based on delete relaxations, in which we temporarily ignore attributes in the state that should be removed, and only focus on solving for the ones that should be added. These methods are not applicable to the generic MDP problem, and are therefore not part of the framework.

Finally, planning can be applied in *open-loop* or *closed-loop* form. Open loop planning, which fully specifies a plan before execution, is only feasible in deterministic environments without full observability of the ground truth state. In closed-loop planning we replan at every timestep, depending on the state that we actually reached in the system. For example, *receding horizon control* (Mayne and Michalska, 1990) computes an open-loop policy at every timestep, which makes it a closed-loop planning approach overall.

2.4 Model-free reinforcement learning

Reinforcement learning (Sutton and Barto, 2018; Wiering and Van Otterlo, 2012) is a large research field within machine learning. The defining assumption of RL is that we do not have

access to a reversible model of the environment, and therefore need to continue sampling from the state that we reach (similar to acting in the real world). This section covers model-free RL, where we directly learn a value or policy from interacting with the irreversible environment.

The planning literature (introduced above) is mostly organized in sub-disciplines, where each discipline focuses on its own set of assumptions or particular approach. In contrast, the RL community is less organized in subtopics, but has rather focused on a range of factors that can be altered in algorithms. This already hints at the possibility of a framework, which should disentangle such factors. We will here introduce some important concepts in RL literature:

- **Value and policy:** While many planning algorithms search for a local solution (e.g., a single trajectory, or only a solution for the current state), RL algorithms in principle approximate a solution for the entire state space. Since RL agents can only try an action once and then have to continue, we cannot really learn a local solution, since we do not know when we will be able to return to the current state. Solutions are usually store in the form of a value function (from which the policy is implicitly derived) or a policy function. Some approaches learn both, where the value function aids in updating the policy, better known as *actor-critic* methods.
- **On- and off-policy bootstrapping:** A crucial idea in RL literature is *bootstrapping*, where we plug in the learned estimate of the value of a state to improve the estimate of a state that precedes it. A key concept is the *temporal difference* error, which is the difference between our previous and new estimate of the value of a state (Sutton, 1988). When bootstrapping state-action values, there is an important distinction between *on-policy* learning, where we estimate the value of the policy that we actually follow, and *off-policy* learning, where we create a value estimate of another (usually greedy) policy. Cardinal examples of the on- and off-policy cases are SARSA (Rummery and Niranjan, 1994) and Q-learning (Watkins and Dayan, 1992), respectively.
- **Exploration:** Exploration is a fundamental theme in nearly all optimization research, where we typically store a (set of) current solution(s) and want to explore to a (set of) potentially better candidate solution(s) around the current solution (set). However, exploration is extra relevant in reinforcement learning, because we also need to collect our own data, which makes the process more brittle.
Many RL exploration methods have focused on injecting some form of noise into the action space decision. Some methods, like ϵ -greedy and Boltzmann exploration, use random perturbation, while other approaches, like confidence bounds (Kaelbling, 1993) or Thompson sampling (Thompson, 1933), base exploration decisions on the remaining uncertainty of an action. While these methods explore in action space, we can also explore in policy parameter space (Plappert et al., 2017). There are other exploration approaches based on intrinsic motivation (Chentanez et al., 2005), like curiosity (Schmidhuber, 1991), or by planning ahead over an uncertain dynamics model (Guez et al., 2012).
- **Generalization:** Since RL tends to store global solutions, it is typically infeasible to store them in a table for problems with a higher dimensional state-space (due to the curse of dimensionality, as already mentioned in the section on Dynamic Programming). Therefore, the RL literature has largely focused on learning methods to *approximate* the solution. Note that such approximation is a supervised learning task itself, which frequently creates a nested supervised learning optimization loop within the outer RL optimization.

A plethora of function approximation methods has been applied to RL, including tile coding, (Sutton, 1996), linear approximation (Bradtke and Barto, 1996), and a recent explosion of (deep) neural network (Goodfellow et al., 2016) applications to RL (Mnih et al., 2015). Recent surveys of deep RL methods are provided by François-Lavet et al.

(2018) and Arulkumaran et al. (2017). Learning not only allows a global solution to be stored in memory (in approximate form), but, equally important, its generalization also provides a fundamental way to share information between similar states.

- **Direct policy optimization:** We may also approach MDP optimization as a direct optimization problem in policy parameter space. An important example are policy gradient methods (Williams, 1992; Sutton et al., 2000; Sutton and Barto, 2018), which provide an unbiased estimator of the gradient of the objective with respect to the policy parameters. We will discuss the policy gradient theorem in much greater detail in Sec. 3.7 of our framework. There has been much research on ways to stabilize policy gradients, for example based on trust region optimization methods (Schulman et al., 2015).

Some gradient-free policy search methods only require the ability to evaluate the objective (the expected cumulative return). Example approaches include evolutionary strategies (ES) applied to the policy parameters (Moriarty et al., 1999; Whiteson and Stone, 2006; Salimans et al., 2017), and the use of the cross-entropy method (CEM) (Rubinstein and Kroese, 2013; Mannor et al., 2003). These approaches treat the MDP as a true black box function which they only need to evaluate. Therefore, they use less MDP specific properties, and will also receive less emphasis in our framework.

There are many specific subtopics in RL research, like hierarchy (Barto and Mahadevan, 2003), goal setting and generalization over different goals (Schaul et al., 2015), transfer between tasks (Taylor and Stone, 2009), inverse reinforcement learning (Abbeel and Ng, 2004), multi-agent learning (Busoniu et al., 2008), etc. While these topics are all really important, our framework solely focuses on a single agent in a single MDP optimization task. However, many of the above topics are complementary to our framework. For example, we may use meta-actions (hierarchical RL) to define a new, more abstract MPD, in which all of the principles of our framework are again applicable.

2.5 Model-based reinforcement learning

In model-based reinforcement learning (Moerland et al., 2020a; Sutton, 1990; Hester and Stone, 2012b), the two research fields of planning and reinforcement learning merge. The original idea of model-based RL was to start from an irreversible environment, and then: i) use sampled data to learn a dynamics model, and ii) use the learned model to improve a learned value or policy. This idea is illustrated in Figure 2.

However, more recently we have also seen a surge of techniques that start from a reversible model, but also use learning techniques for the value or policy. An example is AlphaGo Zero (Silver et al., 2017). Since most researchers also consider this model-based RL, we will define model-based RL as: ‘any MDP approach that uses both a *reversible* model (known or learned) and *learning* of a value or policy to act the environment’.

There are two important steps in model-based RL. First, we should learn a dynamics model itself, which is a supervised learning problem. Since our framework focuses on solving the MDP given a dynamics function, we will not further discuss this topic here. The second important step of model-based RL involves usage of the learned reversible model to improve a value or policy. We will list a few successful approaches to integrate planning in global function approximation:

- **Sampling additional data:** The classic idea of model-based RL was to use the model to sample additional data, which can be used for standard model-free updates. This idea was first introduced in the well-known Dyna algorithm (Sutton, 1990).
- **Multi-step approximate dynamic programming:** More complex integrations use a form a multi-step approximate dynamic programming (Efroni et al., 2019, 2018). In this approach, we use the reversible model to make a multi-step planning back-up, which is then used to update a value or policy approximation at the root of the search.




Figure 2: Model-based versus model-free reinforcement learning. In model-free RL, we directly use experience (data) acquired from the environment to improve a value/policy. In model-based RL, we additionally use the sampled data to learn a model, which can then be used to update the value or policy. Figure based on Sutton and Barto (2018).

This approach has received much recent attention, for example in AlphaGo Zero (Silver et al., 2017) and Guided Policy Search (Levine and Koltun, 2013).

- **Backward trials:** While most models have a forward view (which next states may result from a particular state-action pair), we can also learn a backward model (given a particular state, which state-action pairs could bring us there). A backward model allows us to spread new information more quickly over state-space, by identifying all the possible precursors of a changed state-action value estimate. This idea is better known as *prioritized sweeping* (PS) (Moore and Atkeson, 1993).
- **Value gradients:** When the function class of our learned dynamics model is differentiable, then we can apply gradient-based planning (already introduced in Sec. 2.3). In the RL literature, this approach is known as *value gradients* (Fairbank and Alonso, 2012). A successful example is PILCO (Deisenroth and Rasmussen, 2011), which learns a Gaussian Process (GP) transition model, and combines this with gradient-based planning to achieve good data efficiency in real-world robotics tasks.

For an extensive discussion of model-based RL we refer the reader to the recent survey by Moerland et al. (2020a). This concludes our short overview of planning, model-free RL and model-based RL approaches. The next section will present our framework, which disentangles the common factors in all these methods.

3 Framework for Reinforcement learning and Planning

We now introduce the Framework for Reinforcement Learning and Planning (FRAP). One of the key messages of this article is that both planning and reinforcement learning make the exact same algorithmic choices to solve the MDP problem. For example, a MCTS search of 500 traces is conceptually not too different from 500 episodes of a model-free Q-learning agent in the same environment. In both cases, we repeatedly move forward in the environment to acquire new information, make back-ups to store this information, with the goal to make better informed decisions in the next trace/episode. The model-free RL agent is restricted in the order in which it can visit states, but otherwise, the methodology of exploration, back-ups, representation and updates is the same.

We will center our framework around the concept of *trials* and *back-ups*. We will first introduce these in Section 3.1. Afterwards, we will introduce the dimensions of our framework:

- Where to put our computational effort? (Sec. 3.2)
- Where to make the next trial? (3.3)
- How to estimate the cumulative return? (3.4)
- How to back-up? (3.5)
- How to represent the solution? (3.6)
- How to update the solution? (3.7)

Table 1 is crucial, since it summarizes our entire framework, and can be used as a reference point throughout the sections.

3.1 Trials and back-ups

We will first conceptually define a trial and a back-up:

1. **Trial:** A trial consists of a single call to the environment. We have to specify a certain state action pair (s, a) , and the environment gives us either a sample from, or the entire distributions of, $\mathcal{T}(s'|s, a)$ and $\mathcal{R}(s, a, s')$ (depending on what access we have to the environment, see Figure 1). In Figure 3 this is visualized in red.
2. **Back-up:** The second elementary operation is the 1-step back-up, which uses the information on the state-action pairs below it (for example obtained from the last trial) to update the information of the state-action pairs above it. In Figure 3 this is visualized in green. The back-up can involve any type of information, but frequently involves estimates of state-action values.

The central idea of nearly all MDP optimization algorithms is that the information in the back-up allows us to better choose the location of the next trial. Therefore, most algorithms iterate both procedures. However, we are not forced to alternate a trial and a back-up (Figure 3, right). We may for example first make a series of trials ('a roll-out') to go deep in the domain, and then make a series of back-ups to propagate the information all the way up to the root node.

3.2 Where to put our computational effort?

To make the MDP optimization tractable, the first question that any algorithm implicitly asks is: are there states that we can completely ignore? Fundamentally, we can identify four sets of states, as graphically illustrated in Figure 4:

1. **All states:** i.e., \mathcal{S} .
2. **Reachable states:** all states reachable from any start state under *any* policy.
3. **Relevant states:** all states reachable from any start state under the *optimal* policy.
4. **Start states:** all states with non-zero probability under $p(s_0)$.

Some algorithms find a solution for all states, the most noteworthy example being Dynamic Programming (DP). Such approaches tend to break down in larger problems, as the number of unique states grows exponentially in the dimensionality of the state space. As an illustration, imagine we apply DP to video game playing, where the input is a low-resolution 200x200 pixel greyscale image, with each pixel taking values between 0 and 255. Then the state space consists of $256^{(200 \cdot 200)}$ unique states, a quantity without any meaningful interpretation. However, this state space contains all possible screen configurations, including enormous amounts of noise images that will never occur in the game.

Table 1: Overview of dimensions in the Framework for Reinforcement learning and Planning (FRAP). For any algorithm, we should be able to identify the decision on each of the dimensions. The relevant considerations and possible options on every dimension are shown in the right column. Examples for several algorithms are shown in Table 3. Note: a sample depth of ∞ is better known as a Monte Carlo (MC) roll-out. IM = Intrinsic Motivation.

Dimension	Consideration	Choices
1. Comp. effort (3.2)	- State set	All \leftrightarrow reachable \leftrightarrow relevant
2. Trial selection (3.3)	- Candidate set	Step-wise \leftrightarrow frontier
	- Exploration	Random \leftrightarrow Value-based \leftrightarrow State-based -For value: mean value, uncertainty, priors -For state: ordered, priors (shaping), novelty, knowledge IM, competence IM
	- Phases	One-phase \leftrightarrow two-phase
	- Reverse trials	Yes \leftrightarrow No
3. Return estim. (3.4)	- Sample depth	1 \leftrightarrow n \leftrightarrow ∞
	- Bootstrap func.	Learned \leftrightarrow heuristic \leftrightarrow none
4. Back-up (3.5)	- Back-up policy	On-policy \leftrightarrow off-policy
	- Policy expec.	Expected \leftrightarrow sample
	- Dynamics expec.	Expected \leftrightarrow sample
5. Representation (3.6)	- Function type	Value \leftrightarrow policy \leftrightarrow both (actor-critic) - For all: generalized \leftrightarrow not generalized
	- Function class	Tabular \leftrightarrow function approximation - For tabular: local \leftrightarrow global
6. Update (3.7)	- Loss	- For value: e.g., squared -For policy: e.g., (det.) policy gradient \leftrightarrow value gradient \leftrightarrow cross-entropy, etc.
	- Update	Gradient-based \leftrightarrow gradient-free - For gradient-based, special cases: replace & average update




Figure 3: Trials and back-ups. Left: Grey nodes visualize a search tree, consisting of all the state-action pairs evaluated in the domain so far. In red we visualize the next trial, which picks a state-action pair and queries the environment for either a sample from, or the entire distributions of, \mathcal{T} and \mathcal{R} . The green dotted arrows visualize a back-up, in which the newly acquired information is used to update our value estimates of the state(s) above it. Right: Key procedure in FRAP consists of iterated trials and back-ups.

Therefore, nearly all planning and RL methods start updating states from some start state, thereby only considering states that they can actually reach. Without additional information, this is the only practical way to identify reachable states. However, the reachable state set still tends to be large, and ultimately we are only really interested in the policy in those states that we will encounter under the optimal policy (the relevant states). As the optimal policy is not known in advance, nearly all algorithms beside Dynamic Programming start from the reachable set, and try to gradually narrow this down to the relevant state set. We will discuss approaches to gradually focus on the relevant set in the next section (on exploration).

Some specifications do provide additional information, for example in the form of explicit goal states. This frequently happens in path-planning problems, where we for example want to navigate to a certain destination. This is a form of prior knowledge on the form of the reward function, which peaks at the goal. In such cases, we can also include backwards planning from the goal state, which identifies the reachable state set from two directions. This principle was first introduced as *bidirectional search* (Pohl, 1969), and for example also part of some RRT approaches (LaValle, 1998).

3.3 Where to make the next trial?

A trial is the fundamental way to obtain new information about the environment. The crucial question then becomes: at which state-action pair should we make our next trial? In the previous section we already established that in larger problems, our best chance is to start making trials from the start state (distribution). There are two considerations we need to make for trials selection. First, we need to decide on a *candidate set* of state-action pairs that will be considered for the next trial (Sec. 3.3.1). Then, we need to actually decide which candidate from the set to select, which needs to incorporate some amount of *exploration* (Sec. 3.3.2). At the end of the section, we also briefly touch upon two additional generic concepts in trial selection (phases and reverse trials, in Sec. 3.3.3 and 3.3.4, respectively).

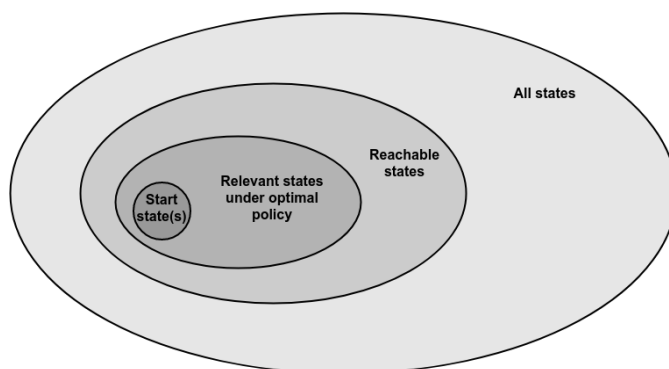


Figure 4: Four sets of states. The reachable state set, a subset of the entire state space, consist of states that are reachable from any start state under any policy. A subset of the reachable states are the relevant states, which are reachable from any start state under the optimal policy. The start states are by definition a subset of the relevant states.

3.3.1 Candidate set selection

The first step is to determine the set of state-action pairs that are candidates for the next trial (in the current iteration). There are two main approaches:

- **Step-wise:** The most frequent approach is to explore *step-wise* on traces from some start state. At each step in the trace, the candidate set consist of all available actions in the current state. After finishing a sequence of step-wise candidate sets, we typically reset to a start state, and repeat the same procedure. This is the standard approach for most RL approaches (Sutton and Barto, 2018) and also for many planning algorithms, like MCTS (Browne et al., 2012). Note that methods that explore by perturbation in (policy) parameter space (Plappert et al., 2017) can be seen as a special form of step-wise perturbation, where the perturbation for all steps is already fixed at the beginning of the episode.
- **Frontier:** The second type of candidate set is a frontier, illustrated in Figure 5. A frontier (or open list) (Dijkstra, 1959) consists of the set of states at the border of the explored region, i.e. those states who have themselves been visited, but who's child states have not all been visited yet. In a search tree, the frontier consists of all leaf nodes, with duplicates states removed (only keeping the occurrence with the lowest cost to reach). The cardinal value-based frontier exploration algorithm is the heuristic search algorithm A* (Hart et al., 1968).

The key difference between step-wise and frontier candidate sets is the moment at which they start exploration (next section). Step-wise methods have a new candidate set at every step in the trace. In contrast, frontier methods only have a single candidate set per episode, fixing a new target at the horizon, and only starting exploration once they are on the frontier.

There a pros and cons for both step-wise and frontier-based candidate sets. A benefit of frontier exploration is that it will by definition explore a new node. By storing the edges of the region you have already visited, you are guaranteed to make a step into new territory. In contrast, step-wise exploration has a risk to repeatedly trial around in an already explored region of state space. This is especially pronounced in tasks with bottleneck states (a narrow passage which brings the agent to another region of the problem). As Ecoffet et al. (2019) mentions, step-wise exploration methods already apply exploration pressure while getting back to the frontier, while we actually want to get back to a new region first, and only then explore. In the long run, (random) step-wise exploration methods will of course also hit the frontier, but this may take a long time of wandering around in known territory.




Figure 5: Illustration of the frontier. Green-shaded nodes have been completely visited, in a sense that either all their children have been visited, or they are terminal. Orange nodes are part the the search tree, but have unvisited child nodes left. Together the orange nodes constitute the frontier (black line), from which we want to continue exploring. White nodes are still unexplored.

Frontier candidate sets also have their challenges. First, frontier exploration assumes that we can always get back to a particular frontier node, which is not guaranteed in stochastic domains (although we may also opt to *approximately* reach the same node (P  r   et al., 2018)). Moreover, in larger problems, the frontier may become very large, let alone all the paths towards it. In those cases, we can no longer store the frontier as a list, or the paths towards it as a tree. We then need to use representation learning for storing both the frontier (Ecoffet et al., 2019) and the paths towards it (P  r   et al., 2018), which may generate instability, and make it hard to actually reach the frontier. Step-wise exploration methods do not have to deal with these issues.

3.3.2 Exploration

Once we have defined the candidate set, we need to decide which state-action pair in the set we will select for the next trail. The *exploitation* decision is to greedily select the action with the highest value estimate. However, as discussed before, this will lead to suboptimal performance. We need to add *exploration* pressure to the greedy policy. We identify three main ways to achieve this: i) random perturbation, ii) value-based perturbation, and iii) state-based perturbation.

- **Random exploration:** In this category we simply inject random exploration noise to the greedy policy. The classic example is ϵ -greedy exploration (Sutton and Barto, 2018), which (in a step-wise candidate set) with small probability randomly selects one of the other actions, independently of its current value estimate or any other characteristics. In continuous action space the noise can for example be Gaussian. We may also inject the noise in (policy) parameter space (Plappert et al., 2017), which may help to correlate it over timesteps.

A benefit of random exploration approaches is that they can guarantee to retain positive exploration pressure throughout learning, and may therefore escape a local optimum when given (a lot of) time. However, they have serious drawbacks as well. Random exploration is undirected, which may lead to *jittering* behaviour, where we undo an exploratory step in the next step (Osband et al., 2016). Moreover, there is no good measure of progress (when should exploration stop), and these methods therefore typically require tedious hyperparameter tuning.

- **Value-based exploration:** A second approach is to use value-based information to better direct the perturbation. There are several approaches:

- *Mean action values:* We may potentially improve over random exploration by incorporating the mean estimates of all the available actions. The general idea is that actions with a higher value estimate also deserve more exploration pressure. In discrete action space, the cardinal example of this approach is Boltzmann exploration (Sutton and Barto, 2018):

$$\pi(a|s) = \frac{\exp(Q(s, a)/\tau)}{\sum_{b \in \mathcal{A}} \exp(Q(s, b)/\tau)}, \quad (4)$$

where $\tau \in \mathbb{R}$ denotes a temperature parameter that scales exploration pressure. For continuous action spaces, we may achieve a similar effect through entropy regularization (Peters et al., 2010; Mnih et al., 2016). These methods usually optimize an adjusted reward function of the form:

$$r(s_t, a_t, s_{t+1}) + \alpha \cdot H(\pi(\cdot|s_t)), \quad (5)$$

where $H(\cdot)$ denotes the entropy of a distribution, and $\alpha \in \mathbb{R}$ is a hyperparameter that scales exploration pressure. The entropy term prevents the policy from converging to a narrow distribution, unless a narrow policy distribution can achieve large gains in expected cumulative return. Thereby, it applies a similar principle as Boltzmann exploration, gradually weighting exploration based on the returns of competing actions.

Compared to random perturbation, the benefit of this approach is that it gradually starts to prefer actions with better returns. On the downside, it does not track any measure of progress (or remaining uncertainty), and therefore cannot assess whether learning has converged, or whether we need additional information. It also depends on the relative scale of the rewards, and can therefore involve tedious tuning of hyperparameters.

- *Action value uncertainty:* A popular approach to exploration uses the remaining uncertainty in the value estimates of the available actions. With high uncertainty around our estimate there is still reason to explore, while reducing uncertainty should gradually shift our policy towards exploitation. A popular uncertainty-based approach are upper confidence bound (UCB) methods (Kaelbling, 1993; Auer, 2002; Kocsis and Szepesvári, 2006; Silver et al., 2017), which for example explore like:

$$\pi(a|s) = Q(s, a) + c \cdot \sqrt{\frac{\ln(n(s))}{n(s, a)}}, \quad (6)$$

where $n(\cdot)$ denotes the number of visits to a state or state-action pair, and $c \in \mathbb{R}$ is a hyperparameter that scales exploration. A popular Bayesian approach to select actions under uncertainty is Thompson sampling (Thompson, 1933). Again, we may want to correlate noise over timesteps, for example by sampling from the value function posterior once at the beginning of a new episode (Osband et al., 2016).

We also consider *pruning* an uncertainty-based method. In certain scenarios, we can completely eliminate an action from the candidate set because we are absolutely certain that it can never outperform an already visited action. This is a form of ‘hard uncertainty’. It for example occurs in two-player games with a minimax (Edwards and Hart, 1961; Knuth and Moore, 1975) structure. Indeed, the *soft pruning* techniques developed in the search community in the early ’80 (Berliner, 1981) can be regarded as early confidence bound methods.

Finally, note that due to the sequential nature of the MDP problem, value uncertainty in a MDP is more complicated than in the bandit setting. In particular, the remaining uncertainty is not only a function of the number of trials at a state-action pair, but also depends on the remaining uncertainty in the value estimates of the state-action pairs that follow it. See Dearden et al. (1998); Osband et al. (2016); Moerland et al. (2017, 2018) for a further discussion of this topic.

- *Priors*: In some cases we may have access to specific prior information about the value function, which then implicitly encodes exploration pressure. The prime example is an *admissible* heuristics. An admissible heuristic provides for every state an optimistic estimate of the cumulative return under the optimal policy. The closer the heuristic is to the true action value, the more prior information about exploration potential we get.

The classic example of a good admissible heuristic is the Euclidean distance to the goal in a path planning task, which can for example be used in A^* (Hart et al., 1968). An admissible heuristic actually provides informative exploration information, as it directly gives an estimate of the remaining value of a node, which may therefore become promising for exploration (or not). However, in most problems an admissible heuristic is not easy to obtain.

- **State-based exploration**: The third main approach to exploration uses state-dependent properties to inject exploration noise, i.e., independently of the value of a particular state action. We again list the most important approaches:

- *Ordered*: First of all, we may simply give every state-action a finite probability of selection. This is better known as a *sweep*. In Dynamic Programming (Bellman, 1966) the sweep is ordered based on the state-space structure, while in exhaustive search (Russell and Norvig, 2016) the sweep is ordered based on a tree structure. Note that a DP sweep is fully exploratory, since it visits every state-action pair in a fixed order, independently of greedy policies.

We also consider *random starts* to be part of this category. Random starts, where our agents starts each new trial at a random state, is part of several classic RL convergence proofs (Watkins and Dayan, 1992; Barto et al., 1995). It ensures that we visit every state eventually infinitely often. Although randomized, it is conceptually close to the DP sweeps, since it ensures that we visit every state-action infinitely often in the limit. We therefore consider it a state-based exploration method, with random ordering.

- *Priors*: The state-based variant of prior information is better known as *shaping*. The best known example are *shaping rewards* (Ng et al., 1999), which are additional rewards placed at states that are likely part of the optimal policy. Recent examples that include this approach are AlphaStar (Vinyals et al., 2019) and For The Win (Jaderberg et al., 2019), who use intermediate game scores as shaping rewards, and optimize the relative weight of each shaping reward in a meta-optimization loop.

Another form of state-dependent priors that guide exploration are *expert demonstrations*, which help to initialize to a good policy. This approach was for example used in the first version of AlphaGo (Silver et al., 2016). While these examples use completely task-specific shaping, we can also find more generic shaping priors. For example, *objects* are generally salient in a task, and children are indeed able to discriminate objects in early infancy. Kulkarni et al. (2016) equips the RL agent with a pre-trained object recognizer, and subsequently places shaping rewards at all detected objects in the scene, which is a more generic form of reward shaping.

- *Novelty*: As discussed before, uncertainty and novelty can be important primitives for exploration. While value-based uncertainty methods use the uncertainty around a value, there are also approaches that use the novelty of the state itself,

independently of its value. An example is *optimistic value initialization* (Sutton and Barto, 2018), where we initialize every state-action estimate to a value higher than the maximum achievable return, which ensures that we initially prefer unvisited actions.

A more formal approach to novelty is the *Probably Approximately Correct* in MDP (PAC-MDP) framework (Kakade et al., 2003). These approaches provide sample complexity guarantees for a RL algorithm, usually based on notions of novelty, ensuring that every reachable state-action pair gets visited enough times. A well-known example is R-max (Brafman and Tenenbholz, 2002), which assumes that every transition in the MDP has maximum reward until it has been visited at least n times. Note that such approaches are generally not computationally feasible in large, high-dimensional state spaces.

- *Knowledge-based intrinsic motivation*: A large group of state-based exploration approaches is knowledge-based intrinsic motivation (Chentanez et al., 2005; Oudeyer et al., 2007). Novelty, as discussed above, is also part of this group, but knowledge-based IM contains a broader set of concepts. The general idea is to provide rewards for events that are intrinsically motivating to humans. Examples include curiosity, novelty, surprise, information gain, reduced model prediction error, etc. These are state dependent properties, independent of the external reward function. For example, depending on the interaction history of an agent, a certain transition can be surprising or not, the model prediction can be correct or completely off, etc. Knowledge-based IM approaches then provide an intrinsic reward for such events, based on the idea that good exploration requires us to decrease novelty, surprise and prediction error over the entire state-space. There is a plethora of different intrinsic motivation approaches (Sutton, 1990; Bellemare et al., 2016; Stadie et al., 2015; Pathak et al., 2017; Lopes et al., 2012; Achiam and Sastry, 2017; Sun et al., 2011; Houthoofd et al., 2016; Dilokthanakul et al., 2019; Mohamed and Rezende, 2015; Hester and Stone, 2012a)
- *Competence-based intrinsic motivation*: In RL, frontier-based exploration has been popularized under the name of competence-based intrinsic motivation (Oudeyer et al., 2007; P  r   et al., 2018). Competence-based IM approaches try to explore by setting their own new goals at the border of their current abilities (i.e., their frontier). This approach typically involves three steps. The first step (goal space learning, for example based on variational auto encoders (P  r   et al., 2018; Laversanne-Finot et al., 2018)) and third step (planning towards the sampled goal) are less relevant from an exploration perspective. The second step involves the exploration decision. We may for example select a new goal based on *learning progress* (Baranes and Oudeyer, 2013; Matiisen et al., 2017), selecting the goal which has shown the largest recent change in our ability to reach it. Otherwise, we can also train a generative model on the state that were of intermediate difficulty to reach, and sample a next goal from this model (Florensa et al., 2018). In any case, the prioritization is dependent on which states we managed to reach so far, and is therefore a form of state-based prioritization.

Note that the above groups are not mutually exclusive. For example, Ecoffet et al. (2019) introduces two methods to prioritize a frontier, one that estimates the amount of progress in the overall task (a value-based prior), and one that uses the visitation frequency of the state (a state-based novelty approach). In summary, we discussed two types of candidate sets (step-wise and frontier) and three approaches to exploration (random, value-based, state-based). Table 2 summarizes our discussion, by displaying common approaches on each of the possible combinations.

Table 2: Schematic overview of common trial selection methods. The columns display the *candidate set* selection method (Sec. 3.3.1), the rows display the way to inject *exploration* pressure to the greedy policy (Sec. 3.3.2). Each cell shows some illustrative example papers. IM = Intrinsic Motivation. PAC-MDP = Probably Approximately Correct in Markov Decision Process (Kakade et al., 2003).

	Step-wise	Frontier
Random	- <i>Random perturbation</i> , e.g., ϵ -greedy (Mnih et al., 2015) Gaussian noise	- <i>Random sampling on frontier</i>
Value-based	- <i>Mean value</i> : e.g., Boltzmann (Mnih et al., 2015), entropy regularization (Peters et al., 2010) - <i>Uncertainty</i> : e.g., confidence bounds (Kaelbling, 1993), posterior sampling (Thompson, 1933)	- <i>Priors</i> , e.g. A* (Hart et al., 1968)
State-based	- <i>Ordered</i> : e.g., DP (Bellman, 1966) - <i>Priors</i> : e.g., reward shaping (Ng et al., 1999) - <i>Novelty</i> : e.g., optim.init. (Sutton and Barto, 2018), PAC-MDP (Brafman and Tenenbholz, 2002) - <i>Knowledge-based IM</i> , e.g., (Achiam and Sastry, 2017)	- <i>Competence-based IM</i> , e.g. (P��r�� et al., 2018)

3.3.3 One versus two phase exploration

The straightforward implementation of the above ideas is to select one method and repeatedly apply it. This is what we call ‘one phase exploration’, where every step of trial selection uses the same method. However, some approaches extend this idea to two distinct phases. It is inspired by the way humans plan and act in the real world, where we typically first plan in our head, and then decide on an action in the real world. The two phases therefore are:

1. *Plan*: repeatedly plan ahead from the *same* state, which is the root of the plan.
2. *Real step*: decide on an action at the root, move forward to the next state, and repeat planning with the resulting state as the new root.

The first step is of course only feasible when we have a reversible model, and two-phase exploration is therefore not applicable to model-free RL. When our goal is to act in a real environment, real steps are enforced by the environment. But we also see the above scheme voluntarily being chosen, for example in AlphaGo Zero (Silver et al., 2017). The reason is that the real step is actually a hard exploration decision itself, since it completely eliminates all the actions that are not chosen. This ensures that we (after a heuristically chosen budget) go deeper in the domain, instead of always searching from the same (root) node, which eventually reduces to exhaustive search.

We may ask ourselves whether Dyna (Sutton, 1990) uses one- or two-phase exploration? Between trials in the real environment, Dyna samples additional data from its learned, reversible model. Typically, it uses the same type of exploration policy and back-up for the additional samples as for data acquired from the real environment. Therefore, it is clearly one-phase in our definition. Multiple phases refers to the use of *different* exploration policies

from the *same* state within one algorithm, but does not depend on the order in which we update states.

When a reversible model is not available but should be approximated from data, two-phase exploration is primarily studied as *Bayes-adaptive exploration* (Guez et al., 2012). In this approach, we first learn a Bayesian dynamics model from the available data. Then, we plan to solve for the optimal action, while we average out over all uncertainty in the dynamics model. We then execute this action, collect new real world data, and repeat the above procedure. This is a provably optimal approach to achieve high data efficiency in the real environment (Guez et al., 2012), but comes at the expense of high computational burden.

3.3.4 Reverse trials

Finally, there is a different approach to trials selection based on the idea of reverse trials. All previous approaches take a forward view on trials, utilizing information about the state-actions in the candidate set obtained from previous trials. However, we can also identify a promising state-action pair for the next trial based on a change in the value of its child state. For this section, we will denote a child of (s, a) as (s', a') . If we reached (s', a') through another trace (not including (s, a)), and the estimate of (s', a') changed a lot, then it is likely that our estimate of (s, a) should be updated as well. In other words, if we learned that a certain state-action pair is good, then we can look back at all the state-action pairs that could bring us here, and update their estimates as well. This idea is better known as *prioritized sweeping* (Moore and Atkeson, 1993).

Prioritized sweeping is actually a special form of a candidate set, but since it is so conceptually different from the rest of the discussion (it requires a reverse model), we nevertheless discuss it separately. The key of prioritized sweeping is a *reverse* model $T^{-1}(s, a|s')$, which tells us which (s, a) can lead to s' . Given a change in some $Q(s', a')$, prioritized sweeping evaluates the *priorities* $\rho(s, a)$ of all possible precursors of s' based on the one-step temporal difference:

$$\rho(s, a) = T(s'|s, a) \cdot \left| R(s, a, s') + \gamma \max_{a'} Q(s', a') - Q(s, a) \right| \quad \forall (s, a) \sim T^{-1}(s, a|s'). \quad (7)$$

Here, we use T and R for the learned transition and reward functions, although the principle equally applies to the ground-truth functions \mathcal{T} and \mathcal{R} . When the priority $\rho(s, a)$ exceeds some small ϵ , we add it to the queue. We then update the state action pair on the top of the queue, and repeat the above procedure for a fixed budget, after which we make a new forward trial.

While forward trials try to figure out where good pay-off may be present further ahead in the MDP, backward trial try to spread the information about an obtained reward as quickly as possible over the nearby states in state space in reverse order. This is graphically illustrated in Figure 6. Note that the difference between multi-step methods, which quickly propagate rewards along the same forward trace, and prioritized sweeping, which spreads to all possible precursors.

3.4 How to estimate the cumulative return?

Once we have selected a trial, we obtain a sampled next state (or a distribution over possible next states) and its associated reward. However, we are not interested in only the single reward of the transition, but actually in the *cumulative return*. The quantity that we need is actually visible in the one-step Bellman equation:

$$Q(s_t, a_t) = \mathbb{E}_{s_{t+1} \sim \mathcal{T}} \left[r_t + \gamma \mathbb{E}_{a_{t+1} \sim \pi} [Q(s_{t+1}, a_{t+1})] \right]. \quad (8)$$




Figure 6: Prioritized sweeping. Regular back-ups are applied in reverse direction of the forward trials (red solid arrows). Prioritized sweeping acknowledges that new reward/value information may also affect other states that lead to a specific outcome. By learning a reverse model, $\hat{T}^{-1}(s, a|s')$, we may identify states in the reverse direction that are candidates for updating (green dashed arrows). The visualization shows that prioritized sweeping can be interpreted as building a new tree in the reverse direction, to spread the obtained reward information more quickly.

In the next section (on back-ups) we will discuss how to deal with the two expectations in Eq. 8. However, we will discuss how get an estimate of $Q(s_{t+1}, a_{t+1})$, i.e., the remaining cumulative reward after the trial. Likely, there is a large subtree below (s_{t+1}, a_{t+1}) , which we can not fully enumerate.

The general form of the cumulative reward estimate takes the following form:

$$\hat{Q}_{K\text{-step}}(s_t, a_t) = \sum_{k=0}^{K-1} \gamma^k \cdot r_{t+k} + \gamma^K B(s_{t+K}), \quad (9)$$

where $K \in \{1, 2, 3, \dots, \infty\}$ denotes the *sample depth* and $B(\cdot)$ is a *bootstrap function*. These are the two key considerations of cumulative reward estimation, which we discuss below.

3.4.1 Sample depth

We first need to determine the sample depth n .

- **$K = \infty$:** A quick way to get an estimate of the cumulative return after the first reward r_t is to sample a deep sequence of trials, and add all the rewards in the trace. In this case ($K \rightarrow \infty$, better known as a *Monte Carlo roll-out*) we do not bootstrap. Although a Monte Carlo roll-out gives an unbiased estimate of the value of the entire remaining subtree, it does have high variance, as we sampled only one realization of all the possible traces. Monte Carlo targets are for example commonly used in MCTS (Browne et al., 2012).
- **$K = 1$:** On the other extreme we directly bootstrap after the trial. One-step targets have low variance but are biased, since the bootstrap estimate can have bias. The bootstrapping function will be discussed in the next section. Well-known algorithms that bootstrap after a single step are for example Q-learning (Watkins and Dayan, 1992) and A^* (Hart et al., 1968).
- **$K = n$:** We can also use an intermediate value for K , which is known as an n -step target, for $1 < n < \infty$.

- **Reweighted:** We can also combine/reweight targets of different depths. Examples include eligibility traces (Sutton and Barto, 2018; Schulman et al., 2016) and more sophisticated reweighting schemes based on importance sampling (Munos et al., 2016).

3.4.2 Bootstrap function

When we stop sampling, we can plug in a fast estimate of the value of the remaining subtree, denoted by $B(\cdot)$ in Eq. 9. This idea is called bootstrapping. There are two main functions to bootstrap from:

- **Learned value function:** We can learn the function to bootstrap from. The ideal candidate is the state value function $V(s)$ or state-action value function $Q(s, a)$ function. These value function may also serve as the solution representation (see Sec. 3.6), in which case they serve two purposes. But also when we represent the solution with a policy, we may still want to learn a value function to bootstrap from.
- **Heuristic:** The second bootstrap approach uses a heuristic (value) function (Pearl, 1984), which is a form of prior information. An admissible heuristic $H(s)$ or $H(s, a)$ gives an optimistic estimate of the cumulative return from a particular state or state-action pair. In many tasks it is hard to obtain a good admissible heuristic, since it should always be optimistic, but should not overestimate the return by too much, as otherwise it is of little benefit. In some planning settings we can obtain a good heuristic by first solving a simplified version of the problem, for example by making a stochastic problem deterministic.

In summary, we need to choose both a sample depth and a bootstrap function to obtain a cumulative reward estimate. Note that a Monte Carlo roll-out is actually a deep sequence of trials. We can of course form value estimates for other state-actions in the trace as well, but our framework focuses on one particular state-action pair that we want to update. Note that we can also make a combined value estimate, for example from two Monte Carlo roll-outs, or from a depth- d limited search. However, these methods simply repeatedly apply the above principle, for example at the leaves of the depth-limited search. How to combine multiple cumulative reward estimates is part of the next section, on the back-up.

3.5 How to back-up?

The trial at s_t, a_t gave us a reward r_t , a next state s_{t+1} (or distribution over next states), and an estimate of the cumulative return. We now wish to back-up this information to improve our estimate of the value at s_t, a_t . In Eq. 8, we still need to specify i) which policy to specify for the back-up, ii) how to deal with the expectation over the actions, and iii) how to deal with the expectation over the dynamics. We will discuss each of them.

3.5.1 Back-up policy

We can in principle specify any back-up policy $\pi^{\text{back}}(a|s)$, which may differ from the forward policy $\pi^{\text{for}}(a|s)$ which we used for trial selection. When $\pi^{\text{back}}(a|s)$ equals $\pi^{\text{for}}(a|s)$ we call the back-up *on-policy*. In all other cases, the back-up is *off-policy*. The cardinal example of an off-policy back-up is the *greedy* or *max* back-up policy, which greedily selects the best action. A benefit of greedy back-ups is that they learn the optimal policy, but they can be unstable in combination with function approximation and bootstrapping (Sutton and Barto, 2018). Some authors study other forms of off-policy back-ups (Keller, 2015; Coulom, 2006), for example more greedy than the exploration policy, but less greedy than the max operator.

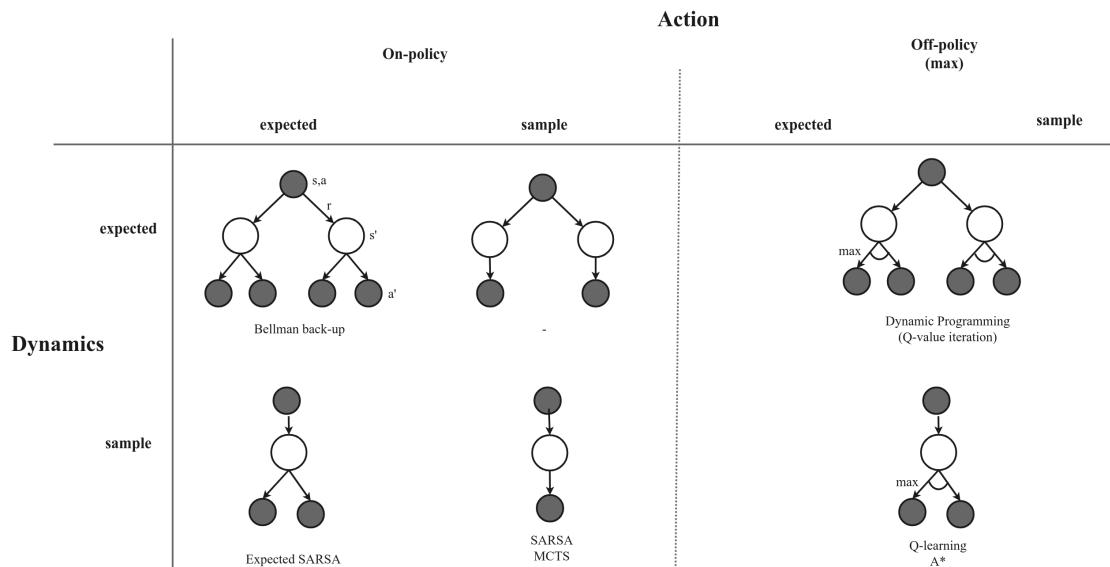


Figure 7: Variants of 1-step back-up. The associated equations are listed in the main text. Mentioned algorithms/back-ups include Value Iteration (Sutton and Barto, 2018), Bellman back-up (Bertsekas et al., 1995), Q-learning (Watkins and Dayan, 1992), Expected SARSA (Van Seijen et al., 2009), SARSA (Rummery and Niranjan, 1994), A* (Hart et al., 1968) and MCTS (Kocsis and Szepesvári, 2006).

3.5.2 Expectation over the actions

Given the back-up policy, we can either make a *sample* or *expected* back-up. A sample back-up samples from the policy, and backs up the value behind this particular action. Sample back-ups are computationally cheap, but need multiple samples to converge to the true value. In contrast, expected back-ups exactly evaluate the expectation over the actions. Sample back-ups are for example used in SARSA (Rummery and Niranjan, 1994), while expected back-ups are used in Expected Sarsa (Van Seijen et al., 2009) and (off-policy) in Tree Backup (Precup, 2000).

3.5.3 Expectation over the dynamics

Like the expectation over the actions, there are two main ways to deal with the expectation over the dynamics: sample, or expected. When the exact transition probabilities are available, then we can exactly evaluate the expectation. Otherwise, when we only have access to an irreversible environment or to a generative model (given or learned), we make a small step in the direction of the sampled value, which will converge to the true value in over multiple back-ups. Although sample-based back-ups provide less information, they can actually be more efficient when many next states have a very small probability (Sutton and Barto, 2018). A special case are deterministic dynamics functions, for which the expected and sample update are equivalent.

The three categories together give rise to several back-up types, as visualized in Figure 7. The vertical axis shows the back-up over the dynamics, while the horizontal axis shows the back-up policy and (nested) the way to deal with the action expectation. In the example, the off-policy back-up (right column) is illustrated by the greedy policy. For the off-policy greedy back-up, the sample and expected action methods are the same, so the right column shows only a single graph centered in the column. For completeness, we list the associated back-up equations below:

- Bellman back-up: $\hat{Q}(s, a) = \mathbb{E}_{s' \sim \mathcal{T}(s'|s, a)}[\mathcal{R}(s, a, s') + \gamma \cdot \mathbb{E}_{a' \sim \pi(a'|s')}Q(s', a')]$
- Q-value iteration: $\hat{Q}(s, a) = \mathbb{E}_{s' \sim \mathcal{T}(s'|s, a)}[\mathcal{R}(s, a, s') + \gamma \cdot \max_{a' \in \mathcal{A}} Q(s', a')]$
- Q-learning: $\hat{Q}(s, a) = \mathcal{R}(s, a, s') + \gamma \cdot \max_{a' \in \mathcal{A}} Q(s', a')$, for $s' \sim T(\cdot|s, a)$
- SARSA: $\hat{Q}(s, a) = \mathcal{R}(s, a, s') + \gamma \cdot Q(s', a')$, for $s' \sim T(\cdot|s, a)$, $a' \sim \pi(\cdot|s')$
- Expected SARSA: $\hat{Q}(s, a) = \mathcal{R}(s, a, s') + \gamma \cdot \mathbb{E}_{a' \sim \pi(a'|s')} [Q(s', a')]$, for $s \sim T(s, a, s')$

We can now look back at the cumulative reward estimation methods from the previous section, and better interpret the Monte Carlo estimate. A MC roll-out is effectively a long sequence of trials, followed by sample-transition, sample-action, on-policy back-ups. So these trials indeed have their own specific back-ups to aid in the update of the root state-action pair under consideration.

3.6 How to represent the solution?

The back-up gave us an improved estimate of the value or policy at some (s, a) . However, we have not discussed yet how the solution will actually be represented. There are two main considerations. First, we need to decide what function we will represent. Second, we need to decide how to represent it in memory.

3.6.1 Function type

There are several ways in which we can represent the solution:

- **Value:** A common solution form is a value function, typically in the form of a state-action values $Q : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$. This function estimates the value of the current or optimal policy at all considered state-action pairs. For action selection, a value function representation requires a mapping from action values to selection probabilities, like an ϵ -greedy or Boltzmann policy.
- **Policy:** A policy $\pi : \mathcal{S} \rightarrow p(\mathcal{A})$ maps every state to a probability distribution over actions. A benefit of learning a policy is that we can directly act in the environment by sampling from the policy distribution. Therefore, this is usually the preferred approach for continuous action spaces, since a max over a continuous action value space requires a nested optimization before we can act.
- **Both:** Some approaches store both a value and a policy function, better known as *actor-critic* methods. The value function is typically used to aid the policy update (see Sec. 3.7).

Generalized value and policy We may extend the above functions by also incorporating a goal that we attempt to reach, better known as *generalized* value or policy functions (Schaul et al., 2015). Generalized value and policy functions take as input the current state and some target/goal state that we would like to reach, and output the value or policy to reach that particular goal state. For example, $Q_g : \mathcal{S} \times \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ would for a particular current state s_0 , goal state s_g , and candidate action a , estimate the value $Q_g(s_0, s_g, a)$, under some reward function that increases when we get closer to s_g . The same principle applies to a generalized policy π_g . The main benefit of generalized value functions is our ability to return to any desired state in state space. The key underlying idea is that we may generalize in goal space, since nearby goals likely share much of their value functions and optimal policy to reach them. We further discuss the concept of generalization in the next section.

There are some others examples of solution representations. For example, some MCTS approaches make their real step decision based on the counts at the root, rather than the value estimates. Counts could be considered a separate function type as well. However, since it is a close derivative of the value, we treat it as a special case of a value function in our framework.




Figure 8: Representation of the solution. Example for a value function $V(s)$ on an one-dimensional, discretized state space. **Left:** Tabular representation. We have obtained estimates for states 2, 4 and 5, while 1, 3 and 6 have not been visited yet. **Right:** Function approximation. Dots indicate the observed data points, the solid line shows a neural network fit on these points. This function generalizes the information in the observations to other (nearby) states, like 1, 3 and 6.

3.6.2 Function class and generalization

Once we determined which function to store, we have to decide how to actually represent it in memory. This topic is closely intertwined with the concept of *generalization*. At the top level, we can discriminate two groups of approaches: i) tabular methods, which do not generalize at all, and ii) function approximation methods, which do generalize. These two approaches are illustrated in Figure 8, and will be discussed in greater detail below:

- **Tabular:** The tabular representation, also known as *atomic*, *symbolic* or *list* representation, treats each input (e.g., state) as a unique element for which we store an individual estimate (e.g., value) (Fig. 8). Tabular representations are the dominant approach in the planning literature, as the nodes in a tree are essentially tables. Note that in a tree, the same state may appear multiple times, like a table with duplicate entries. In the planning literature, a solution to this problem are *transposition* tables, which share information between multiple occurrences of the same state.

For tables, there is an additional important distinction based on the storage duration. On the one hand, *global* tables store the value or policy function for the entire state-action space. However, such a table usually does not fit in memory in large problems. On the other hand, we find *local* tables in planning methods like MCTS (Browne et al., 2012). The local table is temporarily built during the search, but after making a real step we throw away the part of the tree that is no longer relevant in this episode. Another example of local tables are trajectory optimization methods, as used in optimal control, which output a single trajectory (discretized in time).

- **Function approximation:** The second representation approach is function approximation, which builds on the concept of generalization. Generalization implies that similar inputs (states) to a function will in general also have approximately similar output (policy or value) predictions. Thereby, function approximation allows us to share information between near similar states, which a table can not do. A second benefit of function approximation is that due to the approximation we can store a global solution for large state spaces, for which a table would grow too large.

There is a variety of function approximation methods in the machine learning literature. At a high-level we can discriminate parametric methods, like neural networks, and non-parametric methods, like k -nearest neighbours. We will focus on parametric methods here, since these have received most attention and shown most success re-

cently. The most popular groups of parametric function approximation methods are deep neural networks (Goodfellow et al., 2016). For example, if we decide to learn a state-action value function, then we specify the function class $Q_\theta : \mathcal{S} \times \mathcal{A} \times \Theta \rightarrow \mathbb{R}$ with parameters $\theta \in \Theta$. Our aim is to choose the parameters θ in such a way that they accurately predict the state-action estimates obtained from the back-up (as discussed in the previous section).

Note that generalization is actually a spectrum, with complete over-generalization on one extreme, and no generalization at all on the other end. Tabular methods are one extreme, since they do not generalize at all. But once we enter function approximation methods, we still need to balance the amount of generalization to the actual data, better known as balancing overfitting and underfitting.

Reinforcement learning methods have mostly focused on function approximation, while planning has mostly emphasized tabular representation. Tabular methods are easy to implement, stable (since an update to one state action pair does not affect other pairs), and provide good separation between neighboring states. However, their biggest limitation is the memory requirement, which scales exponentially in the dimensionality of the state and action space. Therefore, global solution tables are not feasible in large problems, while local tables cannot profit from offline learning of a global solution.

The benefits of function approximation (generalization, and lower memory requirements) were already mentioned above. Especially generalization can be crucial in large state spaces, where we seldom visit exactly the same state twice, but often encounter approximate similarity. A problem of function approximation is instability, since a local training step also affects the predictions of state-action pairs around it. This is less of a problem in supervised learning with a fixed training set, but since RL collects its own data, a deviation in the policy or value may cause the agent to never explore a certain area of the state-space again. *Replay databases* (Lin and Mitchell, 1992; Mnih et al., 2015) are a way to battle this instability, by reducing the correlation between training data points.

There are some preliminary indications that the combination of function approximation and local tabular methods may actually provide the best of both worlds (Wang et al., 2019; Moerland et al., 2020b). These ideas are inspired by for example AlphaGo Zero (Silver et al., 2017) and Guided Policy Search (Levine and Koltun, 2013), which nest a local tabular planning method in a learning loop. The hypothesis here is that local tabular planning smooths out errors in the global value approximation, while the global approximation provides the necessary information sharing and generalization for the planning method to be effective.

Finally, note that the representations also need to be initialized. The most common approaches are random (for function approximation) or uniform (for tables) initialization. Optimistic initialization, where we initialize all state action value estimates above the maximum achievable return, actually adds an exploration aspect to the initialization, and was already discussed in Sec. 3.3.2.

3.7 How to update the solution?

The last step of our framework involves updating the solution representation (from Sec. 3.6) based on the backed-up value estimates (from Sec. 3.5). While we have mostly discussed value estimation so far, we will now suddenly see policy functions appear more often. The section is split up in two parts. First, we discuss *losses*, which define the way in which our solution can be improved based on the backed-up value. Second, we discuss update rules, which can either be gradient-based or gradient-free. The first part on losses only applies to the gradient-based updates.

3.7.1 Loss

A *loss* is the principled approach in learning to specify the objective. We choose a loss in such a way that when we minimize it, our solution improves. Therefore, it is usually a function of both the solution ($Q_\theta(s, a)$ or $\pi_\theta(a|s)$) and the back-up estimate ($\hat{Q}(s, a)$). Below we discuss some common losses for both value and policy.

Value loss

- **Squared loss:** The most common loss for value function representations is the *mean squared error*. The squared error loss is

$$L(\theta|s_t, a_t) = \frac{1}{2} \left(Q_\theta(s_t, a_t) - \hat{Q}(s_t, a_t) \right)^2. \quad (10)$$

We may also use other losses than the squared loss, as long as minimization of the objective moves our solution closer to the new estimate. For example, Hamrick et al. (2020) recently successfully used a cross-entropy loss between the softmax of the Q-values from the search and the softmax of the Q-values from the value approximation. However, the squared loss is by far most common. In the next section we will show that many common planning updates implicitly use the squared loss as well.

Policy loss There are various ways in which we may specify a policy loss:

- **Policy gradient:** One way to update the policy from a backed-up value estimate is based on the *policy gradient theorem* (Williams, 1992; Sutton et al., 2000; Sutton and Barto, 2018). The theorem gives an unbiased estimator of the gradient of our overall objective (the cumulative reward achieved from the start state):

$$\nabla_\theta V(s_0) = \mathbb{E}_{\pi_\theta, \mathcal{T}} \left[\sum_{t=0}^{\infty} Q(s_t, a_t) \cdot \nabla_\theta \ln \pi_\theta(a_t|s_t) \right], \quad (11)$$

where the expectation runs over all traces induced by π_θ and \mathcal{T} . In practice, the above gradient implicitly specifies a loss. For example, when we use automatic differentiation software, we would implement the policy gradient by sampling traces and minimizing, at every visited state-action pair, the loss

$$L(\theta|s_t, a_t) = -\hat{Q}(s_t, a_t) \cdot \ln \pi_\theta(a_t|s_t) \quad (12)$$

with respect to θ . This last equation clearly shows what the policy gradient equation actually does. It specifies a relation between value estimates $\hat{Q}(s_t, a_t)$ and the policy $\pi_\theta(a_t|s_t)$. If we minimize the above objective, we effectively ensure that actions with a high value also get a high policy probability assigned (since the policy needs to integrate to 1). Note that the policy gradients in not a first-order derivative, but rather a zero-order gradient that tells whether we should move the probability of a certain action up or down. The close relationship between value back-ups and policy gradients is also illustrated by Schulman et al. (2017a).

- **Deterministic policy gradient:** Another popular way to improve a policy based on value estimates is based on deterministic policy gradients (Silver et al., 2014; Lillicrap et al., 2015). These approaches first train a value function based on the methods of the previous paragraph. When we ensure that the learned value function is differentiable with respect to the input action, then we can update the policy by differentiation through the policy action. The associated loss is simply

$$L(\theta|s_t, a_t) = -Q_\psi(s, \pi_\theta(a|s)), \quad (13)$$

where we introduced ψ for the value function parameters, to make clear that the loss is with respect to the policy parameters.

- **Value gradient:** When we have a differentiable reward, transition and policy function, then we can treat our back-up value as a single computational graph, which we can directly optimize (illustrated in Figure 9). This approach is typically applied to sampled traces, for example in PILCO (Deisenroth and Rasmussen, 2011). After sampling a trace, our loss is simply the negative cumulative return:

$$L(\theta|s_t, a_t, \dots, s_\infty) = -\hat{Q}(s, a) = -\sum_{t=0}^{\infty} r_t. \quad (14)$$

We call this objective the value gradient loss. The associated update will be discussed in the next section. Note that the above objective uses an on-policy, sample-action, sample-dynamics back-up, with a sample depth of ∞ and no bootstrapping. However, with a differentiable value function we could also use bootstrapping, and differentiate through the value function as well.

- **Cross-entropy policy loss:** Again, we can in principle can up with any type of policy loss that increase the probability of action that have comparatively higher value estimates. For example, AlphaGo Zero (Silver et al., 2017) makes a heuristic decision for the policy loss. Their MCTS planning procedure returns value estimates and visitation counts at the root of the search. The counts are closely related to the backed-up value estimates in the search, as nodes with higher value estimates get more visits. They propose to normalize the visitation counts to a probability distribution, and train the policy network on a cross-entropy loss with this distribution:

$$L(\theta|s_t) = -\sum_{a \in \mathcal{A}} \log \pi_\theta(a|s_t) \left(\frac{n(s_t, a)}{\sum_b n(s_t, b)} \right), \quad (15)$$

where $n(s_t, a)$ denotes the number of visits to action a at the MCTS root s_t .

The last example illustrates that heuristically motivated losses can work well in practice. The choice for a particular loss may also depend on the setting. For example, value gradients work well in tasks with relatively smooth transition and reward functions, like robotics and control tasks, but have trouble in sparse reward tasks. In short, there is a variety of possible losses for both value and policy targets.

3.7.2 Update rule

The final step of our framework is to actually update our representation. We identify two main approaches: gradient-based (which uses one of the losses of the previous section) and gradient-free optimization.

Gradient-based updates Most learning approach perform gradient-based optimization. The general idea of gradient-based optimization is to repeatedly update our parameters in the direction of the negative gradient of the loss with respect to the parameters:

$$\theta \leftarrow \theta - \alpha \cdot \frac{\partial L(\theta)}{\partial \theta}, \quad (16)$$

where $\alpha \in \mathbb{R}^+$ is a learning rate. We will illustrate some examples:

- *Value update on table:* For a tabular value representation, the θ are simply all the individual table entries $Q_\theta(s, a)$. The derivative of the squared loss (Eq. 10) then becomes

$$\frac{\partial L(\theta)}{\partial \theta} = 2 \cdot \frac{1}{2} \left(Q_\theta(s, a) - \hat{Q}(s, a) \right) = Q_\theta(s, a) - \hat{Q}(s, a). \quad (17)$$

Plugging this into Eq. 16 and reorganizing terms gives the well-known tabular learning rule:

$$Q_\theta(s, a) \leftarrow (1 - \alpha) \cdot Q_\theta(s, a) + \alpha \cdot \hat{Q}(s, a). \quad (18)$$

Note again that this update rule is actually the gradient update of a squared error loss on a value table. This update also makes intuitive sense: we move our table entry $Q_\theta(s, a)$ a bit in the direction of our new estimate $\hat{Q}(s, a)$. Therefore, for the tabular case, we want to keep $\alpha \in [0, 1]$.

We shortly discuss two special cases of the tabular learning rule, which both frequently occur in the planning community:

- *Replace update*: The *replace update* completely replaces the table entry with the new back-up estimate. In Eq. 18, this happens when we set $\alpha = 1$. In that case, it reduces to

$$Q_\theta(s, a) \leftarrow \hat{Q}(s, a). \quad (19)$$

This effectively overwrites the solution with the new estimate obtained from the back-up. We can only afford to do this when we have some guarantees that our new estimate will always improve over our previous estimate. This does specifically happen when we have prior information, like an admissible heuristic. The replace update is for example used in A^* (Hart et al., 1968) planning. When such information is available, replace updates can be much faster than learning updates, which are relatively slow to converge. For example, for route planning on a map (where the euclidean distance in a good admissible heuristic), we would always prefer A^* over Q-learning (Watkins and Dayan, 1992).

- *Averaging update*: The *averaging update*, the second special case of the tabular learning update, ensures that our table entry will remain equal to the *the mean of all previous back-up estimates*. We introduce n to index the update iteration. Then the update rule at every iteration that tracks the average is

$$Q_\theta(s, a) \leftarrow \frac{n-1}{n} Q_\theta(s, a) + \frac{1}{n} \hat{Q}(s, a). \quad (20)$$

Comparing the above to Eq. 18, we see that the averaging update is actually a learning update with $\alpha = \frac{1}{n}$. In other words, we make the learning rate a function of the iteration number. The averaging update is for example the standard approach in MCTS (Browne et al., 2012).

The benefit of the averaging update is that it quickly moves the table entry to a reasonable estimate. After the first iteration, the estimate directly equals the first back-up estimate (while a learning update takes many small steps to move our predictions towards the true estimate). On the downside, fixed learning rates do eventually wash out the effect of the initial estimates, which are typically less reliable. In contrast, averaging updates will always give the initial estimate as much contribution to the table entry as the most recent back-up estimate.

- *Value update with function approximation*: The same principles apply for gradient-based updates in the context of function approximation. If our function approximator is differentiable, then we can simply apply the chain rule to again find the derivative of the loss with respect to the parameters. For example, training a value function approximation on a squared loss (Eq. 10) would have a gradient of

$$\frac{\partial L(\theta|s, a)}{\partial \theta} = (Q_\theta(s, a) - \hat{Q}(s, a)) \cdot \frac{\partial Q_\theta(s, a)}{\partial \theta}, \quad (21)$$

where $\frac{\partial Q_\theta(s, a)}{\partial \theta}$ are for example the derivatives in a neural network.




Figure 9: Illustration of value gradients. Black arrows show the forward specification of an MDP, with a reward function $\mathcal{R}(s, a)$, transition function $\mathcal{T}(s'|s, a)$, and our policy $\pi_\theta(a|s)$ to act in the MDP. If all of these functions are differentiable, then we can update the policy parameters θ by taking the gradient of the cumulative payoff $V(s_0) = \mathbb{E}[\sum_{t=0}^T r(s_t, a_t)]$, with respect to these parameters, as indicated by the red dotted lines. This bears similarity to the way recurrent neural networks are trained with *backpropagation through time*.

- *Policy update with function approximation:* The same chain rule principles apply to the policy gradient loss, and also to the deterministic policy gradient. For example, for the deterministic policy gradient we have:

$$\frac{\partial L(\theta|s, a)}{\partial \theta} = -\frac{\partial Q_\psi(s, a)}{\partial \theta} = -\frac{\partial Q_\psi(s, a)}{\partial a} \frac{\partial \pi_\theta(a|s)}{\partial \theta}, \quad (22)$$

where we again write ψ for the value parameters to distinguish them from the policy parameters.

- *Policy update for value gradient:* Gradient-based planning, better known as value gradients (Heess et al., 2015), is a special case of a policy update. When we have a differentiable dynamics and reward model, and specify a differentiable policy, then we can actually directly differentiate the cumulative reward estimate with respect to the policy parameters (Figure 9).

We will here show the update equations for the gradient of the expected cumulative return $V(s) = \mathbb{E}[\sum_{t=0}^T r(s_t, a_t)|s_0 = s]$. To keep the update equations readable, we will for this equation abbreviate partial differentiation with subscripts, i.e., $V_s = \partial V(s)/\partial s$. The gradient of the sampled trace is given by the following set of recursive relations:

$$\begin{aligned} \hat{V}_\theta &= \mathcal{R}_a \pi_\theta + \gamma \hat{V}'_s \mathcal{T}_a \pi_\theta + \gamma \hat{V}'_\theta, \quad \text{with} \\ \hat{V}_s &= \mathcal{R}_s + \mathcal{R}_a \pi_s + \gamma \hat{V}'_s (\mathcal{T}_s + \mathcal{T}_a \pi_s). \end{aligned} \quad (23)$$

For every trace, the above gradient effectively sums over all paths in Figure 9. In practice we sample a single trace or finite set of traces to compute the gradients with respect to θ . Note the additional V'_θ term in the first equation, which appears since we need to sum the gradients with respect to θ at all timesteps.

Well-known examples of gradient-based planning are PILCO (Deisenroth and Rasmussen, 2011), which achieved high data-efficiency on real-world (small) robotic control tasks, and the linear-quadratic regulator (LQR) (Anderson and Moore, 2007; Todorov and Li, 2005). Gradient-based planning does rely on smooth, differentiable dynamics functions, which makes it mostly applicable to continuous control tasks. Moreover, gradient propagation may suffer from vanishing and exploding gradients, as is also well-known for recurrent neural network (RNN) training.

There are two final remarks for gradient-based updates. First, all above methods have analytic gradients, but we may also use finite differencing to numerically approximate the gradient of our objective. This for example common in optimal control. Second, we have not defined yet how to choose the learning rate in Eq. 16. We neither want to progress too quickly nor too slowly. Most methods use a *line search* with manually tuned learning rate, but other approaches have been popularized in RL as well. A successful approach is to first determine a *trust region*, a region around the current solution in which we aim to search for the next solution, which is for example used in trust region policy optimization (TRPO) (Schulman et al., 2015) and proximal policy optimization (PPO) (Schulman et al., 2017b) algorithms.

Gradient-free updates We have extensively covered losses and learning-based updates. We will now also cover the competing approach, which uses gradient-free optimization. These approaches first specify a parametrized policy function. They then repeatedly: i) perturb the parameters in policy space, ii) evaluate the new solution by sampling traces, and iii) decide whether the perturbed solution should be retained. Example applications to MDP optimization include evolutionary strategies Moriarty et al. (1999); Whiteson and Stone (2006); Salimans et al. (2017), simulated annealing (Atiya et al., 2003) and the cross-entropy method Rubinstein and Kroese (2013); Mannor et al. (2003).

These methods largely bypass the other parts of our framework. They do not use any structural knowledge of the MDP, and never form local estimates of values for a particular state. Instead, they only require an evaluation function (sampling a set of traces), and treat the problem as a black-box optimization setting. There is extensive literature on gradient-free optimization methods, but these methods are not specific to planning and learning in MDPs, and therefore fall outside of the scope of this framework.

This concludes our presentation of FRAP. The discussed dimensions, considerations per dimension, and choices per consideration were already summarized in Table 1. The next section will illustrate the general applicability of FRAP, by analyzing a wide variety of planning and RL algorithms along the dimensions of the framework.

4 Conceptual comparison of well-known algorithms

The key point of FRAP is that planning and learning solve exactly the same problem, and therefore (implicitly) have to make decisions on all the dimensions mentioned in the framework. We illustrate this key idea in Table 3. The table shows for a variety of well-known planning (blue), model-free RL (red) and model-based RL (green) algorithms the choices each algorithm makes on the dimensions of FRAP.

The most important observation from the table is that it reads like a patchwork. On most dimensions, we see similar solution ideas appearing both within planning and reinforcement learning. For example, candidate selection is mostly performed step-wise, but there are both planning, model-free RL and model-based RL papers that use a frontier-based candidate set. For the back-up, MCTS uses an on-policy, sample action, sample transition approach, which is for example shared by SARSA. While these algorithms differ on other dimensions, for example the way they represent their solution, they are similar in their back-up method. Monte Carlo targets for the return estimation appear in all three classes, as do 1-step bootstrapping methods.

There seems to be consensus on few dimensions. For the computational effort dimensions, we do see that nearly all papers in the table except for Dynamic Programming focus on the reachable state set, by sampling forward from some start state distribution. This is indeed our best bet if we do not want to suffer from the curse of dimensionality (see Sec. 3.2).

One may wonder why policy gradient methods still use the value back-up dimension. Policy gradients are actually a form of a loss, which specify how the policy should change

based on a new value estimate. But the value estimate should still be obtained, and any of the methods from Sections 3.4 and 3.5 still apply. For example, policy gradient methods can be combined with Monte Carlo estimates (Williams, 1992), but also with bootstrapping (Mnih et al., 2016).

Note that some approaches, like PILCO (Deisenroth and Rasmussen, 2011) and policy gradients (Williams, 1992), completely rely on a stochastic policy to explore, without any additional exploration pressure. This is technically a form of optimistic initialization, since the start policy should broadly cover state space. There is no additional exploration pressure, and for these methods it is crucial that the initial policy hits a non-zero reward region, since otherwise there will be no learning signal at all. Therefore, this approach seems less applicable to large state spaces.

As discussed in Sec. 3.7, the replace and average update types are special cases of the squared loss. Since the squared loss is never explicitly specified in these tabular updates, we have entered 'squared' between brackets in those cases. For Go-Explore (Ecoffet et al., 2019) we have only considered their initial exploration phase in the table, and omitted the second imitation learning phase in which they solidify their own policy into a neural network. Some smaller comments on the table are part of the table caption.

Table 3: (Next Page): Systematic overview of various learning, planning and model-based RL methods, broken up according to FRAP. See Table 1 for an overview of the components, as discussed throughout Chapter 3. Colour coding: blue = planning, red = model-free RL, green = model-based RL. Abbreviations of function approximation types: NN = neural network, GP = Gaussian Process, k-NN = k-nearest neighbour. Notes: †For Go-Explore (Ecoffet et al., 2019) we only describe their primary exploration approach. In a second stage, they solidify their policy with imitation learning. ∘: Real-time DP leaves the sample depth for the back-up unspecified. In the table we show the vanilla choice for DP itself, a sample depth of 0. § Péré et al. (2018) actually stores $s_0, s_g \rightarrow \theta$, i.e., a mapping from start and goal state to policy parameters, which themselves define another parametric policy.

Table 3: Continued

Paper	Environ- ment	Learned model	Comp. effort	Trial selection					
				Candidate Set	Exploration	Sub-category	Phases	Reverse Trials	Description
Dynamic Programming (Bellman, 1966)	Reversible analytic		All	State set	State	Ordered	1		Sweep
Depth-first search (Russell and Norvig, 2016)	Reversible analytic		Reach.	Step	State	Ordered	1		Sweep
Heuristic search (e.g., A* (Hart et al., 1968))	Reversible analytic		Reach.	Frontier	Value	Prior	1		Greedy on heuristic
MCTS (Browne et al., 2012)	Reversible sample		Reach.	Step	Value	Uncertainty	2		Upper confidence bound
Real-time DP (Barto et al., 1995)	Reversible analytic		Reach.	Step	State	Ordered	1		Random starts
Q-learning (Watkins and Dayan, 1992)	Irreversible sample		Reach.	Step	State	Ordered	1		Random starts
SARSA + eligibility trace (Sutton and Barto, 2018)	Irreversible sample		Reach.	Step	Value	Mean values	1		e.g., Boltzmann
REINFORCE (Williams, 1992)	Irreversible sample		Reach.	Step	Random	-	1		Stochastic policy
DQN (Mnih et al., 2015)	Irreversible sample		Reach.	Step	Value	Random	1		ϵ -greedy
PPO (Schulman et al., 2017b)	Irreversible sample		Reach.	Step	Value	Mean values	1		Stochastic policy with entropy regularization
DDPG (Lillicrap et al., 2015)	Irreversible sample		Reach.	Step	Random	-	1		Noise process (Ornstein-Uhlenbeck)
Go-Explore [†] (Ecoffet et al., 2019)	Irreversible sample		Reach.	Frontier	State+val+rand	Novelty+ prior+random	1		Frontier prior.: visit freq. + heuristics. On frontier: random perturbation.
AlphaStar (Vinyals et al., 2019)	Irreversible sample		Reach.	Step	State+Value	Prior+mean values	1		Imitation learning + shaping rewards + entropy regularization
Dyna-Q (Sutton, 1990)	Irreversible sample	✓	Reach.	Step	State+Value	Knowledge+ mean values	1		Novelty bonus + Boltzmann
Prioritized sweeping (Atkeson and Santamaria, 1997)	Irreversible sample	✓	Reach.	Step	State	Novelty	1	✓	Visitation frequency + Reverse trials
PILCO (Deisenroth and Rasmussen, 2011)	Irreversible sample	✓	Reach.	Step	Random	-	2		Stochastic policy on initialization
AlphaGo (Silver et al., 2017)	Reversible Sample	✓	Reach.	Step	Value + random	Uncertainty	2		Upper confidence bound + noise
Knowledge, e.g., surprise (Achiam and Sastry, 2017)	Irreversible sample	✓	Reach.	Step	State	Knowledge	1		Intrinsic reward for surprise
Competence IM, e.g., (Péré et al., 2018)	Irreversible sample	✓	Reach.	Frontier	State	Competence	1		Sampling in learned goal space

Table 3: Continued

Paper	Cumulative return		Back-up		Representation		Update		
	Sample depth	Bootstrap type	Back-up policy	Action expectation	Dynamics Expectation	Function type	Function class	Loss	Update type
Dynamic Programming (Bellman, 1966)	1	Learned	Off-policy	Max	Exp.	Value	Global table	(Squared)	Replace
Depth-first exh. search (Russell and Norvig, 2016)	∞	None	Off-policy	Max	Exp	Value	Global table	(Squared)	Replace
Heuristic search (e.g., A^* (Hart et al., 1968))	1	Heuristic	Off-policy	Max	Determ.	Value	Global table	(Squared)	Replace
MCTS (Browne et al., 2012)	∞	None	On-policy	Sample	Sample	Value	Local table	(Squared)	Average
Real-time DP (Barto et al., 1995)	1 ^o	Learned	Off-policy	Max	Sample	Value	Global table	(Squared)	Replace
Q-learning (Watkins and Dayan, 1992)	1	Learned	Off-policy	Max	Sample	Value	Global table	Squared	Gradient
SARSA + eligibility trace (Sutton and Barto, 2018)	1 - n	Learned	On-policy	Sample	Sample	Value	Global table	Squared	Gradient
REINFORCE (Williams, 1992)	∞	None	On-policy	Sample	Sample	Policy	Func.approx. (NN)	Policy gradient	Gradient
DQN (Mnih et al., 2015)	1	Learned	Off-policy	Max	Sample	Value	Func.approx. (NN)	Squared	Gradient
PPO (Schulman et al., 2017b)	1 - n	Learned	On-policy	Sample	Sample	Policy	Func.approx. (NN)	Policy gradient	Gradient (trust.reg.)
DDPG (Lillicrap et al., 2015)	1	Learned	Off-policy	Max	Sample	Policy+value	Func.approx. (NN)	Determ. policy grad. + squared	Gradient
Go-Explore [†] (Ecoffet et al., 2019)	1	Heuristic	On-policy	Sample	Sample	Policy	Global table	(Squared)	Replace
AlphaStar (Vinyals et al., 2019)	1- n	Learned	On-policy	Sample	Sample	Policy+value	Func.approx. (NN)	Policy gradient + squared	Gradient
Dyna (Sutton, 1990)	1	Learned	On-policy	Sample	Sample	Value	Global table	Squared	Gradient
Prioritized sweeping (Atkinson and Santamaria, 1997)	1	Learned	Off-policy	Max	Exp.	Value	Global table	Squared	Gradient
PILCO (Deisenroth and Rasmussen, 2011)	∞	None	On-policy	Sample	Sample	Policy	Func.approx. (GP)	Value gradient	Gradient
AlphaGo (Silver et al., 2017)	MCTS: 1- n Value: ∞	Learned	On-policy	Sample	Sample	Policy+value	Func.approx. (NN)+ local table	Cross-entropy+ Squared	Average+ Gradient
Knowledge, e.g., surprise (Achiam and Sastry, 2017)	∞	None	On-policy	Sample	Sample	Policy	Func.approx. (NN)	Policy gradient	Gradient
Competence IM, e.g., (P��r�� et al., 2018)	∞	None	On-policy	Sample	Sample	Generalized policy [§]	Func.approx. (k-NN)	k-NN loss	Gradient-free

5 Related Work

There is surprisingly little work on a systematic categorization of either planning or reinforcement learning algorithms. The two main examples are *trial-based heuristic tree search* (THTS) (Keller, 2015; Keller and Helmert, 2013), and the textbook classification of back-up width and depth by Sutton and Barto (2018). We will discuss both.

THTS is closest to our work, specifying a framework to systematically categorize planning methods. It contains six dimensions, which we will each compare to our framework:

- *Initialization*: In THTS, ‘initialization’ refers to the value a new node in the tree gets assigned when it is generated. The framework describes one possible approach, which is initializing the value with a heuristic. In our framework, this idea is captured as one of the options in the ‘bootstrap’ consideration of the cumulative return estimation dimension (Sec. 3.4.2), where we also discuss other methods.
- *Outcome selection*: In THTS, ‘outcome selection’ refers to the way we generate the next state or state distribution depending on the action. It describes one option: ‘Monte Carlo selection’, which samples each next state according to its probability under \mathcal{T} . In our framework, this is equal to the sample-based dynamics back-up, discussed in Sec. 3.5.3). Note that in our framework we treat this consideration as a back-up choice. If we only sample one action, then we can only back-up a single action, while if we consider the probabilities of all next states, then we can also make an expected back-up. Forward and backward over the dynamics model are thereby directly linked.
- *Back-up*: In THTS the ‘back-up’ dimension covers possible choices like ‘Monte Carlo back-up’ (characterized by ‘averaging’ updates in our framework), ‘Temporal Difference back-up’ (characterized by a bootstrap depth of 1 in our framework), ‘Selective back-ups’ (characterized by a ‘off-policy’ back-up in our framework), etc. To our view, the back-up dimension of THTS actually groups together multiple considerations of the cumulative reward, back-up and update dimensions of our framework. FRAP does properly disentangle these aspects.
- *Trial length*: This dimension describes in THTS by how many trials we expand the search graph. This is clearly related to the ‘sample depth’ dimension of cumulative reward estimation (Sec. 3.4.1) in our framework. However, there is an important additional difference. THTS only counts the graph expansions, which for example for MCTS gives a sample depth of 1 per iteration (ignoring the roll-out). We disagree, as the roll-out is actually a sequence of new trials. We also make back-ups along the roll-out path, but we simply do not use these intermediate estimates to update our representation, which is a separate dimension in our framework.
- *Action selection*: This dimension in THTS contains the categories ‘Greedy’, ‘Uniform’, ‘ ϵ -greedy’, ‘Boltzmann’ and ‘Upper Confidence Bound’. In our framework, this equals part of the exploration dimension (Sec. 3.3.2). However, THTS does not further substructure this dimensions like we do, and thereby fails to incorporate a variety of other methods like intrinsic motivation, frontier-based candidate sets, one- versus two-phase exploration, and reverse trials.
- *Recommendation function*: The final dimension in THTS is the recommendation function, which takes in a search graph and returns a probability distribution over the actions at the root. This category is specific to the online search setting, when we are only interested in the policy at the root. Instead, FRAP contains an entire dimension for solution representation. The above recommendation is a form of a local policy

table in our framework, from which we can read the recommendation decision. But FRAP also includes global representations and various kinds of function approximation methods.

THTS was an important inspiration for the current framework, by proposing that there is a common underlying algorithmic space beneath all MDP search algorithms. However, FRAP extends THTS in many ways, by including the entire spectrum of learning methods (and all its associated RL literature), and by adding and splitting several dimensions to overcome the overlap and confusion of some dimensions of THTS.

Sutton and Barto (2018) also discusses a categorization of planning and learning based on the width and depth of the back-ups. Together these lead to four extremes: exhaustive search (full breadth and depth), Dynamic Programming (full breadth, single depth), Monte Carlo estimation (single breadth, full depth), and temporal difference learning (single breadth, single depth). The depth is clearly represented by the sample depth of the cumulative reward estimation in our framework. The breadth is in FRAP split up in the expectation over the actions and dynamics in the back-up. Note that FRAP considers breadth a back-up dimensions, and therefore considers exhaustive search as a long ordered series of 1-step back-ups. Instead, Sutton and Barto (2018) consider exhaustive search as a single, large, broad and deep back-up. Both view can exist next to one another. Our view better fits a systematic framework that disentangles the elementary operations in search and RL, but the view of Sutton and Barto (2018) is conceptually insightful as well, when we think of an entire planning iteration as creating one new value target.

6 Discussion

This article introduced the framework for reinforcement learning and planning (FRAP), as a systematic approach to categorize and compare planning and reinforcement learning approaches. We will now put our work in a broader perspective, and identify possible implications for future work.

First of all, note that we did not include *stopping criteria* in our framework. Nearly all algorithms empirically stop based on a fixed hyperparameter, or based on manual intervention by a human logging the performance. While some algorithms do have convergence guarantees, like DP (Bellman, 1966), MCTS (Browne et al., 2012), A* (Hart et al., 1968) and many RL algorithms with GLIE (greedy in the limit with infinite exploration) assumptions, it is typically infeasible to assess convergence during execution. The only algorithms that do assess convergence need to either make sweeps through the entire state space (like dynamic programming and exhaustive search), which is the only way to guarantee that we have at a certain moment visited all states frequently enough, or require an admissible heuristic, which ensures that we can stop expanding before visiting all states (Hart et al., 1968).

Tractability of MDP optimization The framework also allows us to zoom out and identify the fundamental ways in which a MDP search can be made tractable. The MDP problem essentially specifies an infinitely deep MAX-EXP tree which we can never fully enumerate. On the most fundamental level, without any consideration of two-phase exploration and a real environment versus planning model, there are only four ways in which we can somehow reduce the size of the true underlying MDP tree:

1. *Reachable states*: focus on reachable states instead of all states (Sec. 3.2).
2. *Exploration-exploitation*: gradually focus from reachable to relevant states, i.e., reduce the breadth of the problem through exploration-exploitation balancing (Sec. 3.3).

3. *Generalization*: share relevance information of one state to other appearances of (approximately) the same state (Sec. 3.6).
4. *Priors*: We do not really consider this a fundamental solution approach, as it requires task specific information. It is however a way to solve an otherwise intractable MDP.

A fifth way to make the problem tractable, which was not discussed in our framework, involves compressing the MDP itself, for example through temporal abstraction (better known as hierarchical RL (Barto and Mahadevan, 2003)) This may define a temporally abstract MDP, in which is is easier to solve for the solution. However, this topic falls outside of the scope of this framework.

Differences between research fields One question that arises from FRAP is: what are the true differences between planning and reinforcement learning? The defining difference was already discussed in Section 2.2. Learning algorithm assume an irreversible environment (‘unknown model’), while planning algorithm assume a reversible environment (‘known model’). Therefore, planning algorithms can repeatedly plan forward from the same state, which RL algorithms cannot. On a conceptual level, the difference is mostly about *the order* in which we do the updates. Once again, 100 traces of MCTS or 100 traces of Q-learning conceptually do the same: they walk forward, acquire information, make back-ups, and update a (local) representation, all to better inform next traces/episodes.

All other differences except for the visitation order seem to be based on convention rather than necessity. We will provide some examples of common conventions in both fields. For example, planning algorithms tend to use (local) tabular representations, in the form of a (discrete) search tree. In contrast, RL algorithms tend to use global representations of the solution, and have put much more emphasis on function approximation. The planning community has put more focus on the use of bootstrapping from heuristics, while the RL community has focused on bootstrapping from learned value functions. Uncertainty-bases exploration has been successful in planning approaches like MCTS, but has also appeared for in in RL research (Kaelbling, 1993). Frontiers originate in research on planning, but competence-based intrinsic motivation now applies similar principles in RL. Sample-based back-ups mostly originate in RL research, where we interact with an irreversible environment and have to rely on sample action, sample dynamics back-ups. However, exactly the same back-up has also become popular in the planning approach of MCTS. In short, both fields have emphasized their own elements of the overall problem, but have at the same time invented similar solutions and approaches, which blurs the algorithmic line between both fields.

Future work The framework helped us identify the following directions for future research:

- *Novel integrations*: We have recently seen a vast surge op integrated planning and learning approaches, like AlphaGo Zero (Silver et al., 2017) and Guided Policy Search (Levine and Koltun, 2013). These model-based RL approaches assume a known model, and nest a planning procedure (with a tabular representation) in the learning loop of a global solution (with a function approximation representation). Such integrations may for example combine the benefits of a tabular representation and function approximation, which each have their strengths and weaknesses, as was already discussed in Sec. 3.6. Such combinations of ideas, which originally belong with separate research fields, may provide mutual benefit for the overall solution.

- *Ideas that have received less attention:* A framework may also help identify which research directions have received little attention recently. One example is prioritized sweeping (Sec. 3.3.4), i.e., the idea of reverse trials to more quickly spread a changed value function over state space. This idea has been successful with tabular models, which are trivial to revert. However, it has hardly been studied when we use function approximation of the dynamics. This is an example of a topic that deserves additional attention, for example in the deep reinforcement learning community.

7 Conclusion

This concludes the description of our framework for reinforcement learning and planning (FRAP). We shortly summarize the main ideas:

- We can disentangle planning algorithms, like A^* , and RL algorithms, like Q -learning, into one underlying framework. Any algorithm that solves a MDP optimization (implicitly) makes decisions on: i) the considered state set, ii) trial selection and exploration, iii) cumulative reward estimation, iv) value back-up, v) solution representation and vi) update of the solution. These dimensions, with their relevant considerations, are summarized in Table 1.
- A key conclusion of the framework is that the lines between planning and learning are actually blurry, and frequently based on convention rather than necessity. Both fields share the same underlying algorithmic space.
- MDP optimization can in principle be approached as a black-box optimization problem. However, our framework illustrates the various ways in MDP specific characteristics can be systematically incorporated in the solution approach.
- Altogether, the framework may serve several purposes: i) provide a common language for researchers in both planning and RL to categorize their solution approach, ii) inspire future research, for example through novel combinations of planning and learning, and iii) serve an educational purpose, for students, and for researchers from either planning or RL who consider working at the intersection of both fields.

References

- Abbeel, P. and Ng, A. Y. (2004). Apprenticeship learning via inverse reinforcement learning. In *Proceedings of the twenty-first international conference on Machine learning*, page 1. ACM.
- Achiam, J. and Sastry, S. (2017). Surprise-based intrinsic motivation for deep reinforcement learning. *arXiv preprint arXiv:1703.01732*.
- Anderson, B. D. and Moore, J. B. (2007). *Optimal control: linear quadratic methods*. Courier Corporation.
- Arulkumaran, K., Deisenroth, M. P., Brundage, M., and Bharath, A. A. (2017). Deep reinforcement learning: A brief survey. *IEEE Signal Processing Magazine*, 34(6):26–38.
- Atiya, A. F., Parlos, A. G., and Ingber, L. (2003). A reinforcement learning method based on adaptive simulated annealing. In *2003 46th Midwest Symposium on Circuits and Systems*, volume 1, pages 121–124. IEEE.
- Atkeson, C. G. and Santamaria, J. C. (1997). A comparison of direct and model-based reinforcement learning. In *Proceedings of International Conference on Robotics and Automation*, volume 4, pages 3557–3564. IEEE.
- Auer, P. (2002). Using confidence bounds for exploitation-exploration trade-offs. *Journal of Machine Learning Research*, 3(Nov):397–422.
- Baranes, A. and Oudeyer, P.-Y. (2013). Active learning of inverse models with intrinsically motivated goal exploration in robots. *Robotics and Autonomous Systems*, 61(1):49–73.
- Barto, A. G., Bradtke, S. J., and Singh, S. P. (1995). Learning to act using real-time dynamic programming. *Artificial intelligence*, 72(1-2):81–138.
- Barto, A. G. and Mahadevan, S. (2003). Recent advances in hierarchical reinforcement learning. *Discrete event dynamic systems*, 13(1-2):41–77.
- Bellemare, M., Srinivasan, S., Ostrovski, G., Schaul, T., Saxton, D., and Munos, R. (2016). Unifying count-based exploration and intrinsic motivation. In *Advances in Neural Information Processing Systems*, pages 1471–1479.
- Bellman, R. (1966). Dynamic programming. *Science*, 153(3731):34–37.
- Berliner, H. (1981). The B* tree search algorithm: A best-first proof procedure. In *Readings in Artificial Intelligence*, pages 79–87. Elsevier.
- Bertsekas, D. P., Bertsekas, D. P., Bertsekas, D. P., and Bertsekas, D. P. (1995). *Dynamic programming and optimal control*, volume 1. Athena scientific Belmont, MA.
- Bertsekas, D. P. and Tsitsiklis, J. N. (1991). An analysis of stochastic shortest path problems. *Mathematics of Operations Research*, 16(3):580–595.
- Bock, H. G. and Plitt, K.-J. (1984). A multiple shooting algorithm for direct solution of optimal control problems. *IFAC Proceedings Volumes*, 17(2):1603–1608.
- Botvinick, M. and Toussaint, M. (2012). Planning as inference. *Trends in cognitive sciences*, 16(10):485–488.

- Bradtke, S. J. and Barto, A. G. (1996). Linear least-squares algorithms for temporal difference learning. *Machine learning*, 22(1-3):33–57.
- Brafman, R. I. and Tenenbholz, M. (2002). R-max-a general polynomial time algorithm for near-optimal reinforcement learning. *Journal of Machine Learning Research*, 3(Oct):213–231.
- 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.
- Busoniu, L., Babuska, R., and De Schutter, B. (2008). A Comprehensive Survey of Multi-agent Reinforcement Learning. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 2(38):156–172.
- Chentanez, N., Barto, A. G., and Singh, S. P. (2005). Intrinsically motivated reinforcement learning. In *Advances in neural information processing systems*, pages 1281–1288.
- Coulom, R. (2006). Efficient selectivity and backup operators in Monte-Carlo tree search. In *International conference on computers and games*, pages 72–83. Springer.
- Dearden, R., Friedman, N., and Russell, S. (1998). Bayesian Q-learning. In *AAAI/IAAI*, pages 761–768.
- Deisenroth, M. and Rasmussen, C. E. (2011). PILCO: A model-based and data-efficient approach to policy search. In *Proceedings of the 28th International Conference on machine learning (ICML-11)*, pages 465–472.
- Dijkstra, E. W. (1959). A note on two problems in connexion with graphs. *Numerische mathematik*, 1(1):269–271.
- Dilokthanakul, N., Kaplanis, C., Pawlowski, N., and Shanahan, M. (2019). Feature control as intrinsic motivation for hierarchical reinforcement learning. *IEEE transactions on neural networks and learning systems*.
- Ecoffet, A., Huizinga, J., Lehman, J., Stanley, K. O., and Clune, J. (2019). Go-explore: a new approach for hard-exploration problems. *arXiv preprint arXiv:1901.10995*.
- Edwards, D. J. and Hart, T. (1961). The alpha-beta heuristic.
- Efroni, Y., Dalal, G., Scherrer, B., and Mannor, S. (2018). Beyond the One-Step Greedy Approach in Reinforcement Learning. In *International Conference on Machine Learning*, pages 1386–1395.
- Efroni, Y., Dalal, G., Scherrer, B., and Mannor, S. (2019). How to combine tree-search methods in reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 3494–3501.
- Fairbank, M. and Alonso, E. (2012). Value-gradient learning. In *The 2012 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE.
- Florensa, C., Held, D., Geng, X., and Abbeel, P. (2018). Automatic Goal Generation for Reinforcement Learning Agents. In *International Conference on Machine Learning*, pages 1514–1523.

- François-Lavet, V., Henderson, P., Islam, R., Bellemare, M. G., Pineau, J., et al. (2018). An introduction to deep reinforcement learning. *Foundations and Trends® in Machine Learning*, 11(3-4):219–354.
- Gelly, S. and Wang, Y. (2006). Exploration exploitation in go: UCT for Monte-Carlo go. In *NIPS: Neural Information Processing Systems Conference On-line trading of Exploration and Exploitation Workshop*.
- Ghallab, M., Howe, A., Knoblock, C., McDermott, D., Ram, A., Veloso, M., Weld, D., and Wilkins, D. (1998). PDDL—the planning domain definition language. *AIPS-98 planning committee*, 3:14.
- Goodfellow, I., Bengio, Y., and Courville, A. (2016). *Deep learning*. MIT press.
- Guez, A., Silver, D., and Dayan, P. (2012). Efficient Bayes-adaptive reinforcement learning using sample-based search. In *Advances in neural information processing systems*, pages 1025–1033.
- Hamrick, J. B., Bapst, V., Sanchez-Gonzalez, A., Pfaff, T., Weber, T., Buesing, L., and Battaglia, P. W. (2020). Combining q-learning and search with amortized value estimates. *International Conference on Learning Representations (ICLR)*.
- Hart, P. E., Nilsson, N. J., and Raphael, B. (1968). A formal basis for the heuristic determination of minimum cost paths. *IEEE transactions on Systems Science and Cybernetics*, 4(2):100–107.
- Heess, N., Wayne, G., Silver, D., Lillicrap, T., Erez, T., and Tassa, Y. (2015). Learning continuous control policies by stochastic value gradients. In *Advances in Neural Information Processing Systems*, pages 2944–2952.
- Hester, T. and Stone, P. (2012a). Intrinsically motivated model learning for a developing curious agent. In *2012 IEEE international conference on development and learning and epigenetic robotics (ICDL)*, pages 1–6. IEEE.
- Hester, T. and Stone, P. (2012b). Learning and using models. In *Reinforcement learning*, pages 111–141. Springer.
- Houthoofd, R., Chen, X., Duan, Y., Schulman, J., De Turck, F., and Abbeel, P. (2016). Vime: Variational information maximizing exploration. In *Advances in Neural Information Processing Systems*, pages 1109–1117.
- Howard, R. A. (1960). Dynamic programming and markov processes.
- Jaderberg, M., Czarnecki, W. M., Dunning, I., Marris, L., Lever, G., Castaneda, A. G., Beattie, C., Rabinowitz, N. C., Morcos, A. S., Ruderman, A., et al. (2019). Human-level performance in 3D multiplayer games with population-based reinforcement learning. *Science*, 364(6443):859–865.
- Kaelbling, L. P. (1993). *Learning in embedded systems*. MIT press.
- Kakade, S. M. et al. (2003). *On the sample complexity of reinforcement learning*. PhD thesis, University of London London, England.
- Kappen, H. J., Gómez, V., and Opper, M. (2012). Optimal control as a graphical model inference problem. *Machine learning*, 87(2):159–182.

- Kaufmann, E. and Koolen, W. M. (2017). Monte-carlo tree search by best arm identification. In *Advances in Neural Information Processing Systems*, pages 4897–4906.
- Keller, T. (2015). *Anytime optimal MDP planning with trial-based heuristic tree search*. PhD thesis, University of Freiburg, Freiburg im Breisgau, Germany.
- Keller, T. and Helmert, M. (2013). Trial-based heuristic tree search for finite horizon MDPs. In *Twenty-Third International Conference on Automated Planning and Scheduling*.
- Knuth, D. E. and Moore, R. W. (1975). An analysis of alpha-beta pruning. *Artificial intelligence*, 6(4):293–326.
- Kocsis, L. and Szepesvári, C. (2006). Bandit based monte-carlo planning. In *ECML*, volume 6, pages 282–293. Springer.
- Kulkarni, T. D., Narasimhan, K., Saeedi, A., and Tenenbaum, J. (2016). Hierarchical deep reinforcement learning: Integrating temporal abstraction and intrinsic motivation. In *Advances in neural information processing systems*, pages 3675–3683.
- LaValle, S. M. (1998). Rapidly-exploring random trees: A new tool for path planning.
- Laversanne-Finot, A., Pere, A., and Oudeyer, P.-Y. (2018). Curiosity Driven Exploration of Learned Disentangled Goal Spaces. In *Conference on Robot Learning*, pages 487–504.
- Levine, S. and Koltun, V. (2013). Guided policy search. In *International Conference on Machine Learning*, pages 1–9.
- Lillicrap, T. P., Hunt, J. J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D., and Wierstra, D. (2015). Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*.
- Lin, L.-J. and Mitchell, T. M. (1992). *Memory approaches to reinforcement learning in non-Markovian domains*. Citeseer.
- Lopes, M., Lang, T., Toussaint, M., and Oudeyer, P.-Y. (2012). Exploration in model-based reinforcement learning by empirically estimating learning progress. In *Advances in neural information processing systems*, pages 206–214.
- Mannor, S., Rubinstein, R. Y., and Gat, Y. (2003). The cross entropy method for fast policy search. In *Proceedings of the 20th International Conference on Machine Learning (ICML-03)*, pages 512–519.
- Matiisen, T., Oliver, A., Cohen, T., and Schulman, J. (2017). Teacher-student curriculum learning. *arXiv preprint arXiv:1707.00183*.
- Mayne, D. Q. and Michalska, H. (1990). Receding horizon control of nonlinear systems. *IEEE Transactions on automatic control*, 35(7):814–824.
- Mnih, V., Badia, A. P., Mirza, M., Graves, A., Lillicrap, T., Harley, T., Silver, D., and Kavukcuoglu, K. (2016). Asynchronous methods for deep reinforcement learning. In *International conference on machine learning*, pages 1928–1937.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., et al. (2015). Human-level control through deep reinforcement learning. *Nature*, 518(7540):529.

- Moerland, T. M., Broekens, J., and Jonker, C. M. (2017). Efficient exploration with double uncertain value networks. *Deep Reinforcement Learning Symposium, 31st Conference on Neural Information Processing Systems (NIPS)*.
- Moerland, T. M., Broekens, J., and Jonker, C. M. (2018). The Potential of the Return Distribution for Exploration in RL. *Exploration in Reinforcement Learning Workshop, 35th International Conference on Machine Learning (ICML)*.
- Moerland, T. M., Broekens, J., and Jonker, C. M. (2020a). Model-based Reinforcement Learning: A Survey. *arXiv preprint arXiv:2006.16712*.
- Moerland, T. M., Deichler, A., Baldi, S., Broekens, J., and Jonker, C. M. (2020b). Think Too Fast Nor Too Slow: The Computational Trade-off Between Planning And Reinforcement Learning. *arXiv preprint arXiv:2005.07404*.
- Mohamed, S. and Rezende, D. J. (2015). Variational information maximisation for intrinsically motivated reinforcement learning. In *Advances in neural information processing systems*, pages 2125–2133.
- Moore, A. W. and Atkeson, C. G. (1993). Prioritized sweeping: Reinforcement learning with less data and less time. *Machine learning*, 13(1):103–130.
- Moore, E. F. (1959). The shortest path through a maze. In *Proc. Int. Symp. Switching Theory, 1959*, pages 285–292.
- Moriarty, D. E., Schultz, A. C., and Grefenstette, J. J. (1999). Evolutionary algorithms for reinforcement learning. *Journal of Artificial Intelligence Research*, 11:241–276.
- Munos, R., Stepleton, T., Harutyunyan, A., and Bellemare, M. (2016). Safe and efficient off-policy reinforcement learning. In *Advances in Neural Information Processing Systems*, pages 1054–1062.
- Ng, A. Y., Harada, D., and Russell, S. (1999). Policy invariance under reward transformations: Theory and application to reward shaping. In *ICML*, volume 99, pages 278–287.
- Osband, I., Blundell, C., Pritzel, A., and Van Roy, B. (2016). Deep exploration via bootstrapped DQN. In *Advances in Neural Information Processing Systems*, pages 4026–4034.
- Oudeyer, P.-Y., Kaplan, F., and Hafner, V. V. (2007). Intrinsic motivation systems for autonomous mental development. *IEEE transactions on evolutionary computation*, 11(2):265–286.
- Pathak, D., Agrawal, P., Efros, A. A., and Darrell, T. (2017). Curiosity-driven exploration by self-supervised prediction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 16–17.
- Pearl, J. (1984). Heuristics: intelligent search strategies for computer problem solving.
- Péré, A., Forestier, S., Sigaud, O., and Oudeyer, P.-Y. (2018). Unsupervised learning of goal spaces for intrinsically motivated goal exploration. *arXiv preprint arXiv:1803.00781*.
- Peters, J., Mulling, K., and Altun, Y. (2010). Relative entropy policy search. In *Twenty-Fourth AAAI Conference on Artificial Intelligence*.

- Plappert, M., Houthoofd, R., Dhariwal, P., Sidor, S., Chen, R. Y., Chen, X., Asfour, T., Abbeel, P., and Andrychowicz, M. (2017). Parameter space noise for exploration. *arXiv preprint arXiv:1706.01905*.
- Pohl, I. (1969). Bidirectional and heuristic search in path problems. Technical report.
- Precup, D. (2000). Eligibility traces for off-policy policy evaluation. *Computer Science Department Faculty Publication Series*, page 80.
- Puterman, M. L. (2014). *Markov Decision Processes.: Discrete Stochastic Dynamic Programming*. John Wiley & Sons.
- Rubinstein, R. Y. and Kroese, D. P. (2013). *The cross-entropy method: a unified approach to combinatorial optimization, Monte-Carlo simulation and machine learning*. Springer Science & Business Media.
- Rummery, G. A. and Niranjan, M. (1994). *On-line Q-learning using connectionist systems*, volume 37. University of Cambridge, Department of Engineering Cambridge, England.
- Russell, S. J. and Norvig, P. (2016). *Artificial intelligence: a modern approach*. Malaysia; Pearson Education Limited,.
- Salimans, T., Ho, J., Chen, X., Sidor, S., and Sutskever, I. (2017). Evolution strategies as a scalable alternative to reinforcement learning. *arXiv preprint arXiv:1703.03864*.
- Schaul, T., Horgan, D., Gregor, K., and Silver, D. (2015). Universal value function approximators. In *International Conference on Machine Learning*, pages 1312–1320.
- Schmidhuber, J. (1991). A possibility for implementing curiosity and boredom in model-building neural controllers. In *Proc. of the international conference on simulation of adaptive behavior: From animals to animats*, pages 222–227.
- Schulman, J., Chen, X., and Abbeel, P. (2017a). Equivalence between policy gradients and soft q-learning. *arXiv preprint arXiv:1704.06440*.
- Schulman, J., Levine, S., Abbeel, P., Jordan, M., and Moritz, P. (2015). Trust region policy optimization. In *International conference on machine learning*, pages 1889–1897.
- Schulman, J., Moritz, P., Levine, S., Jordan, M., and Abbeel, P. (2016). High-Dimensional Continuous Control Using Generalized Advantage Estimation. In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. (2017b). Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., Van Den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., et al. (2016). Mastering the game of Go with deep neural networks and tree search. *nature*, 529(7587):484.
- Silver, D., Lever, G., Heess, N., Degris, T., Wierstra, D., and Riedmiller, M. (2014). Deterministic Policy Gradient Algorithms. In *International Conference on Machine Learning*, pages 387–395.
- Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., Hubert, T., Baker, L., Lai, M., Bolton, A., et al. (2017). Mastering the game of go without human knowledge. *Nature*, 550(7676):354.

- Stadie, B. C., Levine, S., and Abbeel, P. (2015). Incentivizing exploration in reinforcement learning with deep predictive models. *arXiv preprint arXiv:1507.00814*.
- Sun, Y., Gomez, F., and Schmidhuber, J. (2011). Planning to be surprised: Optimal bayesian exploration in dynamic environments. In *International Conference on Artificial General Intelligence*, pages 41–51. Springer.
- Sutton, R. S. (1988). Learning to predict by the methods of temporal differences. *Machine learning*, 3(1):9–44.
- Sutton, R. S. (1990). Integrated architectures for learning, planning, and reacting based on approximating dynamic programming. In *Machine Learning Proceedings 1990*, pages 216–224. Elsevier.
- Sutton, R. S. (1996). Generalization in reinforcement learning: Successful examples using sparse coarse coding. In *Advances in neural information processing systems*, pages 1038–1044.
- Sutton, R. S. and Barto, A. G. (2018). *Reinforcement learning: An introduction*. MIT press.
- 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*, pages 1057–1063.
- Taylor, M. E. and Stone, P. (2009). Transfer learning for reinforcement learning domains: A survey. *Journal of Machine Learning Research*, 10(Jul):1633–1685.
- Tesauro, G. and Galperin, G. R. (1997). On-line Policy Improvement using Monte-Carlo Search. In Mozer, M. C., Jordan, M. I., and Petsche, T., editors, *Advances in Neural Information Processing Systems 9*, pages 1068–1074. MIT Press.
- Thompson, W. R. (1933). On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3/4):285–294.
- Todorov, E. and Li, W. (2005). A generalized iterative LQG method for locally-optimal feedback control of constrained nonlinear stochastic systems. In *Proceedings of the 2005, American Control Conference, 2005.*, pages 300–306. IEEE.
- Toussaint, M. (2009). Robot trajectory optimization using approximate inference. In *Proceedings of the 26th annual international conference on machine learning*, pages 1049–1056. ACM.
- Van Seijen, H., Van Hasselt, H., Whiteson, S., and Wiering, M. (2009). A theoretical and empirical analysis of Expected Sarsa. In *2009 IEEE Symposium on Adaptive Dynamic Programming and Reinforcement Learning*, pages 177–184. IEEE.
- Vinyals, O., Babuschkin, I., Czarnecki, W. M., Mathieu, M., Dudzik, A., Chung, J., Choi, D. H., Powell, R., Ewalds, T., Georgiev, P., et al. (2019). Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature*, pages 1–5.
- Wang, T., Bao, X., Clavera, I., Hoang, J., Wen, Y., Langlois, E., Zhang, S., Zhang, G., Abbeel, P., and Ba, J. (2019). Benchmarking Model-Based Reinforcement Learning. *CoRR*, abs/1907.02057.
- Watkins, C. J. and Dayan, P. (1992). Q-learning. *Machine learning*, 8(3-4):279–292.

- Whiteson, S. and Stone, P. (2006). Evolutionary function approximation for reinforcement learning. *Journal of Machine Learning Research*, 7(May):877–917.
- Wiering, M. and Van Otterlo, M. (2012). Reinforcement learning. *Adaptation, learning, and optimization*, 12:3.
- Williams, R. J. (1992). Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256.