Recursive Reward Aggregation

Yuting Tang Yivan Zhang Johannes Ackermann Yu-Jie Zhang Soichiro Nishimori Masashi Sugiyama

Keywords: Markov decision process, reward aggregation, policy preference, Bellman equation, algebraic data type, dynamic programming, recursion scheme, algebra fusion, bidirectional process

Summary

In reinforcement learning (RL), agents typically learn desired behaviors by maximizing the (discounted) sum of rewards, making the design of reward functions crucial for aligning the agent behavior with specific objectives. However, because rewards often carry intrinsic meanings tied to the task, modifying them can be challenging and may introduce complex trade-offs in real-world scenarios. In this work, rather than modifying the reward function itself, we propose leveraging different reward aggregation functions to achieve different behaviors. By introducing an algebraic perspective on Markov decision processes (MDPs), we show that the Bellman equations naturally emerge from the recursive generation and aggregation of rewards. This perspective enables the generalization of the standard discounted sum to other recursive aggregation functions, such as discounted max and Sharpe ratio. We empirically evaluate our approach across diverse environments using value-based and actor-critic algorithms, demonstrating its effectiveness in optimizing a wide range of objectives. Furthermore, we apply our method to a real-world portfolio optimization task, showcasing its potential for practical deployment in decision-making applications where objectives cannot easily be expressed as the discounted sum of rewards.

Contribution(s)

- We provide an algebraic perspective on the recursive structure of MDPs based on fusion. Context: The algebra of recursive functions (Meijer et al., 1991; Bird & de Moor, 1997; Hutton, 1999) is a well-studied topic in functional programming. The fusion technique, explored by Hinze et al. (2010), has been applied to dynamic programming (Bellman, 1966; De Moor, 1994; Bertsekas, 2022). In the context of RL, the recursive structure of the discounted sum of rewards was studied by Hedges & Sakamoto (2022). Our diagrammatic representation of recursive reward generation and aggregation processes is inspired by Gavranović (2022).
- 2. We generalize the Bellman equations and Bellman operators for the standard discounted sum to other recursive aggregation functions, providing greater flexibility in goal specification. Context: The problem of alternative reward aggregations is not entirely new. Prior works have explored objectives such as optimizing the maximum (Quah & Quek, 2006; Gottipati et al., 2020; Veviurko et al., 2024), minimum (Cui & Yu, 2023), top-k (Wang et al., 2020), and Sharpe ratio (Nägele et al., 2024) of rewards. Specifically, the method proposed by Cui & Yu (2023) is a special case of our framework, where the recursive structure is on the original reward space, and the update function is order-preserving.
- We extend existing RL algorithms by incorporating the generalized Bellman operators and empirically demonstrate their effectiveness across various tasks.
 Context: While our method modifies the Bellman operators within the base RL algorithms, the fundamental structures of Q-learning (Watkins, 1989; Watkins & Dayan, 1992), PPO (Schulman et al., 2017), and TD3 (Fujimoto et al., 2018) remain unchanged.

Recursive Reward Aggregation

Yuting Tang^{1, 2}Yivan Zhang^{1, 2}Johannes Ackermann^{1, 2}Yu-Jie Zhang²Soichiro Nishimori^{1, 2}Masashi Sugiyama^{2, 1}tang@ms.k.u-tokyo.ac.jpyivan.zhang@k.u-tokyo.ac.jp

¹The University of Tokyo, Japan ²RIKEN AIP, Japan

Abstract

In reinforcement learning (RL), aligning agent behavior with specific objectives typically requires careful design of the reward function, which can be challenging when the desired objectives are complex. In this work, we propose an alternative approach for flexible behavior alignment that eliminates the need to modify the reward function by selecting appropriate reward aggregation functions. By introducing an algebraic perspective on Markov decision processes (MDPs), we show that the Bellman equations naturally emerge from the recursive generation and aggregation of rewards, allowing for the generalization of the standard discounted sum to other recursive aggregations, such as discounted max and Sharpe ratio. Our approach applies to both deterministic and stochastic settings and integrates seamlessly with value-based and actor-critic algorithms. Experimental results demonstrate that our approach effectively optimizes diverse objectives, highlighting its versatility and potential for real-world applications.

1 Introduction

Reinforcement learning (RL) formalizes sequential decision-making as interaction between an agent and an environment modeled by a Markov decision process (MDP). In standard RL, the objective is to *maximize the discounted sum of rewards* obtained through interaction (Sutton & Barto, 1998; Bowling et al., 2023). This formulation has been widely adopted across various domains, including games (Mnih et al., 2015; Silver et al., 2018; Guss et al., 2019), autonomous driving (Kiran et al., 2021), and stock trading (Wu et al., 2020; Kabbani & Duman, 2022; Liu et al., 2024).

While the discounted sum is standard in RL, many important objectives cannot be expressed in this form. For example, in tasks where stability matters, minimizing some measure of *variability* in rewards is as important as maximizing expected returns (Sobel, 1982; Tamar et al., 2012). In finance, the *Sharpe ratio* (Sharpe, 1966) evaluates risk-adjusted returns by penalizing high volatility, requiring optimization beyond simple returns. Other objectives include (i) maximizing the *peak performance* in drug discovery to identify the most effective compounds (Quah & Quek, 2006; Gottipati et al., 2020), (ii) maximizing the *worst-case outcome* in safety-critical domains like self-driving (Wang et al., 2020) or *bottleneck objective* in network routing (Cui & Yu, 2023), or (iii) maximizing the *average reward* in continuing tasks where future and immediate returns are equally important (Sutton & Barto, 1998, Section 10.3). These cases call for alternative reward aggregation beyond the discounted sum.

Can we simply *modify the reward function* to accommodate these objectives? This is a natural idea, and prior work has explored shaping or redesigning the rewards to reflect alternative criteria (Moody et al., 1998; Ng et al., 1999; Moody & Saffell, 2001; Nägele et al., 2024). However, this approach often requires expanding the state space to encode long-term objectives (Mannor & Tsitsiklis, 2011; Wang et al., 2020), which can change the effective optimization landscape. Moreover, manually redesigning a reward function that induces the desired behavior is notoriously difficult (Leike et al., 2017; Hadfield-Menell et al., 2017; Zhu et al., 2020), which can lead to reward hacking or goal misalignment (Amodei et al., 2016; Christiano et al., 2017; Di Langosco et al., 2022; Ji et al., 2023).



Figure 1: Illustration of three deterministic policies in a simple environment, shown as colored paths with their rewards on edges. The table on the right shows the aggregated rewards for each policy, with the <u>optimal scores</u> (higher \uparrow or lower \downarrow) underlined. We can observe that different aggregation functions lead to different policy preferences.

In this work, we propose a simple yet general alternative: *optimize different reward aggregations* directly, while keeping the state space and reward function fixed. This shifts the focus from what to reward to how to evaluate reward sequences, enabling greater flexibility without increased structural complexity. Choosing the right aggregation is essential because it defines the optimization objective. As shown in Fig. 1, even in a toy environment, different aggregation functions such as sum, mean, or max can lead to different policy preferences. This highlights the need for a general framework that can express and optimize such objectives in a unified way.

Intuition In order to optimize reward aggregations directly, our key insight is that many aggregation functions, including the standard discounted sum, can be computed *recursively*, one reward at a time. For example, the discounted sum aggregation sum_{γ} with a discount factor γ satisfies

$$\operatorname{sum}_{\gamma}[r_1, r_2, r_3, \ldots] := r_1 + \gamma \cdot \operatorname{sum}_{\gamma}[r_2, r_3, \ldots] = r_1 + \gamma r_2 + \gamma^2 r_3 + \ldots,$$
(1)
while the discounted maximum aggregation max obeys

 $\max_{\gamma}[r_1, r_2, r_3, \ldots] := \max(r_1, \gamma \cdot \max_{\gamma}[r_2, r_3, \ldots]) = \max\{r_1, \gamma r_2, \gamma^2 r_3, \ldots\}.$ (2) This shared structure suggests a unifying algebraic view: each aggregation *folds* a reward sequence using an update rule and an initial value. Such recursions naturally induce Bellman-like equations, enabling direct optimization of diverse objectives using standard RL machinery. Concretely, we use a technique known as *algebra fusion* (Hinze et al., 2010) to derive Bellman-style updates for a wide range of recursive reward aggregations, integrating them into standard RL algorithms without altering the state space or reward function.

Related work Several studies have extended Bellman-style updates to optimize non-cumulative objectives. An early method by Quah & Quek (2006) defined a value function for expected discounted maximum rewards, but it lacked rigorous justification and incorrectly interchanged expectation with maximum (Gottipati et al., 2020), effectively optimizing the discounted maximum of expected rewards. Later approaches addressed this issue by augmenting the state space with auxiliary variables (Veviurko et al., 2024). Beyond maximum objectives, Wang et al. (2020); Cui & Yu (2023) analyzed broader classes of objectives, including minimum, harmonic mean, and top-k, but their approaches either required symmetry (precluding discounting), were limited to deterministic systems, or implicitly used order-preserving properties. In contrast, we derive value functions from first principles using algebraic fusion (Hinze et al., 2010), supporting recursive aggregations involving multi-dimensional statistics (e.g., for range, mean, and variance), while unifying deterministic and stochastic cases.

Contributions In this paper, we introduce an *algebraic perspective* on the MDP model, showing that the Bellman equations naturally emerge from the *recursive* generation and aggregation of rewards (Section 2). This perspective allows us to generalize the standard discounted sum to other recursive aggregation functions, such as discounted max and Sharpe ratio (Section 3), while unifying deterministic and stochastic settings within the same framework (Section 4). We provide theoretical justification for our approach, which enables the optimization of various objectives beyond cumulative rewards while maintaining computational efficiency. Finally, we validate the effectiveness of our method in both discrete and continuous environments across various recursive reward aggregation functions, showcasing its flexibility and scalability in handling diverse reward structures (Section 5).¹

¹Code: https://github.com/Tang-Yuting/recursive-reward-aggregation.

2 An algebraic perspective on Bellman equations

In this section, we introduce the standard MDP model (Puterman, 1994) for sequential decisionmaking problems from an algebraic perspective. Using a technique known as *fusion* in algebra and functional programming (Meijer et al., 1991; Hinze et al., 2010), we show that the Bellman equations (Bellman, 1966) naturally arise from the *recursive* generation and aggregation of rewards. This perspective reveals opportunities for generalizing to alternative reward aggregation functions.

In this section, we focus on the standard discounted sum and deterministic transitions and policies. We study other recursive aggregations in Section 3 and stochastic transitions and policies in Section 4.

2.1 Preliminaries

Notation In this section, S is the set of *states*, A is the set of *actions*, and R is the set of *rewards*, which can be finite or infinite. The dynamics of the environment is given by a (deterministic) transition function $p: S \times A \to S$. An agent interacts with the environment by following a (deterministic) policy $\pi: S \to A$ that maps states to actions. A reward function $r: S \times A \to R$ assigns a reward to each state-action pair. Furthermore, we assume that there is an *initial state* $s_0 \in S$ and a subset $S_{\omega} \subset S$ of terminal states, whose indicator function is ω . The horizon Ω of the task can be fixed or varying, depending on the terminal condition ω .

Moreover, $\{*\}$ denotes a singleton (any set with a single element *). [R] denotes the set of *finite lists* of rewards, defined using the *empty list function* nil : $\{*\} \to [R]$, which represents the empty list [], and the *list constructor function* cons : $R \times [R] \rightarrow [R]$, which prepends an element to a list. We have cons(r, []) = [r] and $cons(r_t, [r_{t+1}, \dots, r_{\Omega}]) = [r_t, r_{t+1}, \dots, r_{\Omega}]$, which we abbreviate as $r_{t:\Omega}$.

Composite functions Let us introduce some composite functions that are useful for defining the recursive generation of states, actions, and rewards. Given a policy $\pi: S \to A$, the *pairing function* $(\mathrm{id}_S, \pi) : S \to S \times A = [s \mapsto s, \pi(s)]$ keeps a copy of the current state $s \in S$ and outputs the next action $\pi(s) \in A$.²³ Then, pre-composing this function with the transition function $p: S \times A \to S$ and the reward function $r: S \times A \rightarrow R$ yields two *policy-dependent* functions as follows. We use the subscripts π to explicitly indicate the dependence on the policy π :

■ state transition $p_{\pi} : S \to S := p \circ \langle id_S, \pi \rangle = [s \mapsto p(s, \pi(s))]$ and ■ state reward function $r_{\pi} : S \to R := r \circ \langle id_S, \pi \rangle = [s \mapsto r(s, \pi(s))].$

2.2 Recursive generation of rewards

Using the state transition p_{π} and reward function r_{π} , we can generate states and rewards step by step:

$$\operatorname{step}_{\pi, \mathrm{p}, \mathrm{r}, \omega} : S \to \{*\} + R \times S := \begin{bmatrix} s \mapsto \begin{cases} * & s \in S_{\omega} \\ \mathrm{r}_{\pi}(s), \mathrm{p}_{\pi}(s) & s \notin S_{\omega} \end{bmatrix}.$$
(3)

Let us take a closer look at this step function. The codomain, $\{*\} + R \times S$, is the disjoint union (+) of a singleton $\{*\}$, representing termination, and the Cartesian product $R \times S$ of rewards and states. At each step, the step function either halts by returning the *termination signal* * if the current state s is terminal or continues by returning a pair of the reward $r_{\pi}(s) \in R$ and the next state $p_{\pi}(s) \in S$, both determined by the policy π .

Remark 1 (Terminal condition). By incorporating the terminal condition ω into the step function, we can describe both *episodic* and *continuing* tasks for any reward aggregation, without relying on a special absorbing state and the unit of the aggregation function, e.g., 0 for the discounted sum function. See also Sutton & Barto (1998, Section 3.4).

²For a set C, id_C : $C \to C$ is the *identity function* mapping an element $c \in C$ to itself. For two functions $f : C \to A$ and $g: C \to B$, their pairing $\langle f, g \rangle: C \to A \times B$ is the unique function that applies these two functions to the same input, mapping an input $c \in C$ to a pair $(f(c), g(c)) \in A \times B$ of outputs.

³We write name : domain \rightarrow codomain = [input \mapsto output], assigning an *anonymous function* [input \mapsto output] to a named, typed function name : domain \rightarrow codomain, following Petrov (2020).

Starting from an initial state, by recursively applying this step function and collecting the results, we can obtain a sequence of rewards:

Definition 2.1 (Recursive generation). Given a policy π , a transition function p, a reward function r, and a terminal condition ω , a *recursive reward generation function* $gen_{\pi,p,r,\omega} : S \to [R]$ is defined as follows:

$$\operatorname{gen}_{\pi, \mathrm{p}, \mathrm{r}, \omega} : S \to [R] := \left\lfloor s \mapsto \begin{cases} [\] & s \in S_{\omega} \\ \operatorname{cons}(\mathrm{r}_{\pi}(s), \operatorname{gen}_{\pi, \mathrm{p}, \mathrm{r}, \omega}(\mathrm{p}_{\pi}(s))) & s \notin S_{\omega} \end{cases} \right\rfloor.$$
(4)

2.3 Recursive aggregation of rewards

Given a sequence of rewards, we can aggregate them into a single value using an aggregation function. In the standard MDP setting, the *discounted sum* $\sup_{\gamma} : [R] \to R := \left[r_{1:\Omega} \mapsto \sum_{t=1}^{\Omega} \gamma^{t-1} r_t\right]$ of rewards is a standard choice, where $\gamma \in [0, 1]$ is a *discount factor*.

Note that the discounted sum function can be expressed as a recursive function:

$$\operatorname{sum}_{\gamma} : [R] \to R := \begin{bmatrix} [] & \mapsto & 0 \\ r_{t:\Omega} & \mapsto & r_t + \gamma \cdot \operatorname{sum}_{\gamma}(r_{t+1:\Omega}) \end{bmatrix}.$$
(5)

In other words, the discounted sum function is uniquely defined by two functions: the base case $0 \in R$ and the recursive case "discounted addition" $+_{\gamma} : R \times R \to R := [a, b \mapsto a + \gamma \cdot b]$. This recursive structure has been used, explicitly or implicitly, in prior work on alternative objectives (Quah & Quek, 2006; Hedges & Sakamoto, 2022; Cui & Yu, 2023; Veviurko et al., 2024). In Section 3, we show that many other reward aggregations also admit similar recursive definitions.

2.4 Bellman equation for the state value function

We have introduced the recursive generation and aggregation of rewards in a standard MDP model. The generation function $gen_{\pi,p,r,\omega}: S \to [R]$ is the *producer* of rewards, and the discounted sum function $sum_{\gamma}: [R] \to R$ is the *consumer* of rewards. By composing these two recursive functions, we obtain a *state value function* $v_{\pi}: S \to R$, which can also be calculated recursively:

$$\mathbf{v}_{\pi}: S \to R := \operatorname{sum}_{\gamma} \circ \operatorname{gen}_{\pi, \mathbf{p}, \mathbf{r}, \omega} = \begin{bmatrix} s \mapsto \begin{cases} 0 & s \in S_{\omega} \\ \mathbf{r}_{\pi}(s) + \gamma \cdot \mathbf{v}_{\pi}(\mathbf{p}_{\pi}(s)) & s \notin S_{\omega} \end{bmatrix}.$$
(6)

This recursive calculation of the state value function $v_{\pi} : S \to R$ is known as the *Bellman equation* (Bellman, 1966), which expresses the value of a state s under a policy π as the sum of the immediate reward $r_{\pi}(s)$ and the discounted value of the next state $p_{\pi}(s)$.

Remark 2 (State-action recursion). We can define the state-action transition/step/generation functions and derive a Bellman equation for the *state-action value function* $q_{\pi} : S \times A \to R$ in a similar way, which is omitted here for brevity and discussed in Appendix A.

Remark 3 (Algebra fusion). For readers familiar with algebra and functional programming, we point out that the Bellman equation emerges as a consequence of the *fusion law* for recursive coalgebras (Hinze et al., 2010, Section 4; Yang & Wu, 2022, Section 10), shown in the following diagram:⁴

$$\operatorname{id}_{\{*\}} + \operatorname{id}_R \times \mathbf{v}$$

$$\{*\} + R \times S \xrightarrow{\operatorname{id}_{\{*\}} + \operatorname{id}_R \times \operatorname{gen}_{\pi, p, r, \omega}} \{*\} + R \times [R] \xrightarrow{\operatorname{id}_{\{*\}} + \operatorname{id}_R \times \operatorname{sum}_{\gamma}} \{*\} + R \times R$$

$$\underset{\{*\} \xrightarrow{s_0} S}{\overset{s_0} \longrightarrow} S \xrightarrow{gen_{\pi, p, r, \omega}} [R] \xrightarrow{gen_{\pi, p, r, \omega}} \{R] \xrightarrow{s_0} R$$

$$(7)$$

The left square is the recursive definition of the generation function in Eq. (4), and the right square is the recursive definition of the discounted sum function in Eq. (5). Consequently, the whole rectangle is the Bellman equation for the state value function in Eq. (6). See Appendix B for more details.

⁴For two functions $f : A \to C$ and $g : B \to C$, their *copairing* $[f, g] : A + B \to C$ is the unique function defined by cases, mapping an input $x \in A + B$ to f(x) if $x \in A$, to g(x) if $x \in B$.

	definition post $\circ \operatorname{agg}_{\operatorname{init},\triangleright} : [R] \to R$	initial value of statistic(s) init $\in T$	update function $\triangleright: R \times T \to T$	post-processing $post: T \to R$
discounted sum	$r_1 + \gamma r_2 + \dots + \gamma^{t-1} r_t$	discounted sum $s: 0 \in \mathbb{R}$	$+_\gamma:=[r,s\mapsto r+\gamma\cdot s]$	$\mathrm{id}_{\mathbb{R}}$
discounted min	$\min\{r_1, \gamma r_2, \dots, \gamma^{t-1} r_t\}$	discounted min $n: \infty \in \overline{\mathbb{R}}$	$\min_\gamma := [r, n \mapsto \min(r, \gamma \cdot n)]$	$\mathrm{id}_{\overline{\mathbb{R}}}$
discounted max	$\max\{r_1, \gamma r_2, \dots, \gamma^{t-1} r_t\}$	discounted max $m: -\infty \in \overline{\mathbb{R}}$	$\mathrm{max}_\gamma := [r, m \mapsto \mathrm{max}(r, \gamma \cdot m)]$	$\mathrm{id}_{\overline{\mathbb{R}}}$
log-sum-exp	$\log(e^{r_1} + e^{r_2} + \dots + e^{r_t})$	log-sum-exp $m:-\infty\in\overline{\mathbb{R}}$	$[r, m \mapsto \log(e^r + e^m)]$	$\mathrm{id}_{\overline{\mathbb{R}}}$
range	$\max(r_{1:t}) - \min(r_{1:t})$	$\max_{m \in \mathbb{N}} m \begin{bmatrix} -\infty \\ \infty \end{bmatrix} \in \overline{\mathbb{R}}^2$	$\left[r, \begin{bmatrix}m\\n\end{bmatrix} \mapsto \begin{bmatrix}\max(r,m)\\\min(r,n)\end{bmatrix}\right]$	$\left[\begin{bmatrix} m \\ n \end{bmatrix} \mapsto m - n \right]$
mean	$\overline{r} := \frac{1}{t} \sum_{i=1}^t r_i$	$ \sup_{s \text{ sum } s} n \begin{bmatrix} 0 \\ 0 \end{bmatrix} \in \begin{bmatrix} \mathbb{N} \\ \mathbb{R} \end{bmatrix} $	$\left[r, \begin{bmatrix}n\\s\end{bmatrix} \mapsto \begin{bmatrix}n+1\\s+r\end{bmatrix}\right]$	$\left[\begin{bmatrix} n \\ s \end{bmatrix} \mapsto \frac{s}{n} \right]$
		$ \begin{array}{c} \operatorname{length} n \\ \operatorname{mean} m \\ \end{array} \begin{bmatrix} 0 \\ 0 \end{bmatrix} \in \begin{bmatrix} \mathbb{N} \\ \mathbb{R} \end{bmatrix} $	$\left[r, \begin{bmatrix}n\\m\end{bmatrix} \mapsto \begin{bmatrix}n+1\\\frac{n\cdot m+r}{n+1}\end{bmatrix}\right]$	$\left[\begin{bmatrix} n \\ m \end{bmatrix} \mapsto m \right]$
variance	$\frac{1}{t}\sum_{i=1}^{t}(r_i-\overline{r})^2 = \overline{r^2} - \overline{r}^2$	$ \begin{array}{c} \operatorname{length} n \\ \operatorname{sum} s \\ \operatorname{sum} \operatorname{square} q \end{array} \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix} \in \begin{bmatrix} \mathbb{N} \\ \mathbb{R} \\ \mathbb{R}_{\geq 0} \end{bmatrix} $	$\left[r, \begin{bmatrix}n\\s\\q\end{bmatrix} \mapsto \begin{bmatrix}n+1\\s+r\\q+r^2\end{bmatrix}\right]$	$\left[\begin{bmatrix} n \\ s \\ q \end{bmatrix} \mapsto \frac{q}{n} - \left(\frac{s}{n} \right)^2 \right]$
		$ \begin{array}{c} \operatorname{length} n \\ \operatorname{mean} m \\ \operatorname{variance} v \end{array} \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix} \in \begin{bmatrix} \mathbb{N} \\ \mathbb{R} \\ \mathbb{R}_{\geq 0} \end{bmatrix} $	$\begin{bmatrix} n \\ m \\ v \end{bmatrix} \mapsto \begin{bmatrix} n+1 \\ \frac{n\cdot m+r}{n+1} \\ v + \frac{n(r-m)^2 - (n+1)v}{(n+1)^2} \end{bmatrix} \end{bmatrix}$	$\left[\begin{bmatrix} n \\ m \\ v \end{bmatrix} \mapsto v \right]$
top-k	$k\text{-th}$ largest in $r_{1:t}$	$\begin{array}{ccc} \operatorname{top-1} & \left[-\infty \\ \operatorname{top-2} & \left[-\infty \\ -\infty \\ \operatorname{buffer} & \vdots \end{array} \right] \in \overline{\mathbb{R}}^k \end{array}$	$\begin{bmatrix} r, b \mapsto \begin{cases} \operatorname{insert}(r, b) & r > \min b \\ b & r \le \min b \end{bmatrix}$	$[b\mapsto\min b]$

Table 1: Recursive aggregation functions

3 Recursive reward aggregation functions

In this section, we generalize the discounted sum function in Eq. (5) to other recursive reward aggregation functions that summarize a sequence of rewards into a single value. Our primary goal is to derive a generalized Bellman equation extending Eq. (6) and provide theoretical insights for efficient policy evaluation and optimization with recursive reward aggregation.

3.1 Bellman equation for the state statistic function

First, we observe that many aggregation functions are inherently recursive. However, the recursive structure does not always operate directly within the original space. For instance, we can calculate the arithmetic mean by calculating both the sum and the length recursively and then dividing the sum by the length. Based on this observation, we propose the following definition:

Definition 3.1 (Recursive aggregation). Let T be a set of *statistics*. Given an *initial value* init $\in T$, an *update function* $\triangleright : R \times T \to T$, and a *post-processing function* post $: T \to R$, a *recursive statistic aggregation function* $\operatorname{agg}_{init,\triangleright} : [R] \to T$ of is defined as follows:

$$\operatorname{agg}_{\operatorname{init},\triangleright} : [R] \to T := \begin{bmatrix} [] & \mapsto & \operatorname{init} \\ r_{t:\Omega} & \mapsto & r_t \triangleright \operatorname{agg}_{\operatorname{init},\triangleright}(r_{t+1:\Omega}) \end{bmatrix},$$
(8)

and a *recursive reward aggregation function* post \circ agg_{init,>} : $[R] \rightarrow R$ is the composition of this function with the post-processing function post : $T \rightarrow R$, shown in the following diagram:

By substituting the discounted sum function with a general recursive reward aggregation function, we can generalize the Bellman equation in Eq. (6) as follows:

Theorem 3.2 (Bellman equation for the state statistic function). Given a recursive reward generation function $gen_{\pi,p,r,\omega}$ (Definition 2.1) and a recursive statistic aggregation function $agg_{init,\triangleright}$ (Definition 3.1), their composition, called the state statistic function $\tau_{\pi} : S \to T$, satisfies

$$\tau_{\pi}: S \to T := \operatorname{agg}_{\operatorname{init}, \triangleright} \circ \operatorname{gen}_{\pi, p, r, \omega} = \left[s \mapsto \begin{cases} \operatorname{init} & s \in S_{\omega} \\ r_{\pi}(s) \triangleright \tau_{\pi}(p_{\pi}(s)) & s \notin S_{\omega} \end{cases} \right].$$
(10)



Figure 2: By combining the recursive generation and aggregation of rewards, we can express the state statistic function $\tau_{\pi} : S \to T$ as a composition of *bidirectional processes*. The forward process $S \to R \times S$, parameterized by a policy π , takes a state $s_t \in S$ and generates a reward $r_{t+1} \in R$ and the next state $s_{t+1} \in S$. The backward process $R \times T \to T$ takes a statistic $\tau_{t+1} \in T$ from the future and updates it with the previously generated reward $r_{t+1} \in R$ to produce the current statistic $\tau_t \in T$. These bidirectional processes continue until a terminal state is reached, at which point its statistic is assigned the initial value init $\in T$. See Appendix B for more details.

Definition 3.3 (Value function). The state value function $v_{\pi} : S \to R := \text{post} \circ \tau_{\pi}$ is the composition of the state statistic function $\tau_{\pi} : S \to T$ with the post-processing function post $: T \to R$.

While prior work such as Quah & Quek (2006) defined the recursive structure of the value function directly, our approach derives it from the recursive structure of the reward generation and aggregation processes. Examples of recursive reward aggregation functions are provided in Table 1. An illustration of the recursive structure is given in Fig. 2.

3.2 Policy evaluation: Iterative statistic function estimation

Next, we consider how to estimate the state statistic function $\tau_{\pi}: S \to T$ for an arbitrary policy π , known as the *policy evaluation* problem (Sutton & Barto, 1998, Sections 4.1 and 11.4). We introduce a generalized *Bellman operator* and prove the uniqueness of its fixed points under certain conditions. This result enables iterative statistic/value function estimation used in *policy iteration* and modern *actor-critic* methods (Barto et al., 1983; Mnih et al., 2016; Haarnoja et al., 2018; Fujimoto et al., 2018). Concretely, the Bellman operator is defined as follows:

Definition 3.4 (Bellman operator). Given a policy π , a transition function p, a reward function r, a terminal condition ω , and a recursive statistic aggregation function $\underset{t}{\operatorname{agg}}$ (Definition 3.1), the *Bellman operator* $\mathcal{B}_{\pi} : [S,T] \to [S,T]$ for a function $\tau : S \to T$ is defined by

$$\mathcal{B}_{\pi}\tau: S \to T := \left[s \mapsto \begin{cases} \text{init} & s \in S_{\omega} \\ r_{\pi}(s) \triangleright \tau(\mathbf{p}_{\pi}(s)) & s \notin S_{\omega} \end{cases} \right].$$
(11)

According to the Bellman equation in Theorem 3.2, we have $\mathcal{B}_{\pi}\tau_{\pi} = \tau_{\pi}$, which means that the state statistic function τ_{π} is a fixed point of the Bellman operator. Then, we can generalize the classical fixed point theorem under the following condition:

Definition 3.5 (Contractive update function). An update function $\triangleright : R \times T \to T$ is *contractive* with respect to a premetric d_T on statistics T if $\forall r \in R$. $\forall \tau_1, \tau_2 \in T$. $d_T(r \triangleright \tau_1, r \triangleright \tau_2) \leq k \cdot d_T(\tau_1, \tau_2)$, where $k \in [0, 1)$ is a constant. In other words, $r \triangleright (-) : T \to T$ is a contraction for all $r \in R$.

Theorem 3.6 (Uniqueness of fixed points of Bellman operator). Let $\tau_1, \tau_2 : S \to T$ be fixed points of the Bellman operator \mathcal{B}_{π} (Definition 3.4). If the update function \triangleright is contractive with respect to a premetric d_T on statistics T (Definition 3.5), then $d_T(\tau_1(s), \tau_2(s)) = 0$ for all states $s \in S$. If d_T is a strict premetric, then $\tau_1 = \tau_2 = \tau_{\pi}$.

This result applies to a broad class of recursive aggregation functions beyond the discounted sum. See Appendix C for further discussion on the premetric d_T and the Bellman operator \mathcal{B}_{π} .

3.3 Policy optimization: Optimal policies and optimal value functions

Finally, we consider how to find an *optimal policy* and compute its statistic/value functions recursively based on the Bellman equation in Theorem 3.2:

Definition 3.7 (Optimal policy). Given a preorder \leq_T on statistics T, a policy π_* is an *optimal policy* if $\forall \pi$. $\forall s \in S$. $\tau_{\pi}(s) \leq_T \tau_{\pi_*}(s)$, which has the *optimal state statistic function* $\tau_* : S \to T := \tau_{\pi_*}(s)$ and the optimal state value function $v_*: S \to R := \text{post} \circ \tau_*$.

Theorem 3.8 (Bellman optimality equation for the state statistic function). Given a preorder \leq_T on statistics T, the optimal state statistic function τ_* (Definition 3.7) satisfies

$$\tau_*: S \to T := \left[s \mapsto \begin{cases} \text{init} & s \in S_{\omega} \\ \sup_{a \in A} (\mathbf{r}(s, a) \triangleright \tau_*(\mathbf{p}(s, a))) & s \notin S_{\omega} \end{cases} \right].$$
(12)

Definition 3.7 and Theorem 3.8 are analogous to their classical counterparts (Sutton & Barto, 1998, Section 3.6), but they extend to arbitrary recursive aggregation functions and allow comparisons using a preorder \leq_T on statistics. A *Bellman optimality operator* \mathcal{B}_* can be defined similarly to the Bellman operator in Definition 3.4, and we can prove the uniqueness of its fixed points under certain conditions. This result enables the value iteration algorithm (Sutton & Barto, 1998, Section 4.4), temporal difference methods such as *Q-learning* (Watkins, 1989), and deep Q-network (DQN) based methods (Mnih et al., 2013; Bellemare et al., 2017) to find the optimal policy π_* . See Appendix D for further discussion on the preorder \leq_T and the Bellman optimality operator \mathcal{B}_* .

4 From deterministic to stochastic Markov decision processes

In this section, we briefly discuss the extension of our framework to the stochastic setting. We show that the deterministic and stochastic settings share a fundamental similarity: all recursive structures remain unchanged, except that (deterministic) functions are replaced by stochastic functions, and function composition is replaced by marginalization over the intermediate variable, as described by the Chapman–Kolmogorov equation (Giry, 1982; Puterman, 1994). The main difference is that the stochastic setting allows for a richer class of aggregation functions (Bellemare et al., 2023), where the non-commutativity and non-distributivity of certain operations can lead to more complex behaviors.

Notation Slightly abusing notation, we use the same symbols to denote the *measurable spaces* of states S, actions A, rewards R, and statistics T. For a measurable space C, we write $\mathbb{P}C$ for the measurable space of all *probability measures* on C, and we denote by $\delta_c \in \mathbb{P}C$ the *Dirac measure* concentrated at $c \in C$. An identity stochastic function $id_C : C \to \mathbb{P}C : [c \mapsto \delta_c]$ maps an element $c \in C$ to the Dirac measure $\delta_c \in \mathbb{P}C$. We consider stochastic transition $p: S \times A \to \mathbb{P}S$ and policy $\pi: S \to \mathbb{P}A$, while other functions can be deterministic. We also use the usual conditional distribution notation such as p(s'|s, a) and $\pi(a|s)$.

Stochastic composite functions In the stochastic setting, we can compose two stochastic functions by marginalizing over the intermediate variable. Additionally, we can compose a stochastic function with a deterministic one using the *pushforward* operation, which is equivalent to treating deterministic functions as stochastic functions to Dirac measures. Then, we can define stochastic versions of

- state transition $p_{\pi} : S \to \mathbb{P}S := p \circ \langle id_S, \pi \rangle = [s \mapsto s' \sim \int_A p(s'|s, a)\pi(a|s) da]$ and state reward function $r_{\pi} : S \to \mathbb{P}R := r \circ \langle id_S, \pi \rangle = [s \mapsto r \sim \int_A \delta_{r(s,a)}(r)\pi(a|s) da].$

Stochastic recursive functions Analogous to Theorem 3.2, we can derive the recursive calculation of the stochastic state statistic function $\tau_{\pi}: S \to \mathbb{P}T$, known as the distributional Bellman equation (Morimura et al., 2010a;b; Bellemare et al., 2017), for any recursive aggregation function agg_{init}.

$$\tau_{\pi}: S \to \mathbb{P}T = \begin{bmatrix} s \mapsto \tau \sim \begin{cases} \delta_{\text{init}} & s \in S_{\omega} \\ r \triangleright \tau' \mid r \sim r_{\pi}(r|s), \tau' \sim \int_{S} \tau_{\pi}(\tau'|s') p_{\pi}(s'|s) \, \mathrm{d}s' & s \notin S_{\omega} \end{bmatrix}.$$
(13)



Figure 3: The recursive structures of (left) the *expected discounted sum of rewards* $\mathbb{E}[r+_{\gamma} s]$ and (right) the *discounted sum of expected rewards* $\mathbb{E}[r]+_{\gamma}\mathbb{E}[s]$.

Stochastic aggregation functions Note that this framework also accommodates the traditional *expected discounted sum of rewards* $\mathbb{E}\left[\sum_{t=1}^{\Omega} \gamma^{t-1} r_t\right]$ learning objective, by selecting δ_0 as init, the (pushforward through) discounted addition function $+_{\gamma} : R \times R \to R$ as the update function \triangleright , and the expectation operator $\mathbb{E} : \mathbb{P}R \to R$ as post. The stochastic statistic function $\tau_{\pi} : S \to \mathbb{P}R$ in Eq. (13), refered to as the *value distribution* in Bellemare et al. (2017), outputs the distribution of the discounted sum of rewards, while the value function outputs its expectation. Since the expectation distributes over the discounted addition, by changing the update function and initial value, we can recursively calculate the *discounted sum of expected rewards* $\sum_{t=1}^{\Omega} \gamma^{t-1} \mathbb{E}[r_t]$ instead (see Fig. 3), which is the traditional approach in RL (Sutton & Barto, 1998). In this case, the statistic function and the value function coincide, as no post-processing is required. However, Bellemare et al. (2017) have shown that even in the discounted sum setting, the Bellman operator may be a contraction in some metrics but not in others, while the Bellman optimality operator is a contraction only in expectation and not in any distributional metric, leading to different convergence behaviors. These challenges persist and may become unavoidable when using alternative aggregation functions due to the inconsistency between expected aggregated rewards and aggregated expected rewards. We discuss this further in Appendix E and leave a full investigation for future work.

5 Experiments

In this section, we empirically evaluate the proposed *recursive reward aggregation* technique across a variety of environments and optimization objectives to support the following claims:

- Different aggregation functions significantly influence policy preferences. Selecting an appropriate
 aggregation function is an alternative approach to optimizing policies for specific objectives and
 aligning agent behaviors with task-specific goals without modifying rewards (Sections 5.1 to 5.3).
- In challenging real-world applications such as portfolio optimization, our method can directly
 optimize desired evaluation criteria, demonstrating superior performance compared to existing
 approaches and showcasing its practical effectiveness (Section 5.4).

5.1 Grid-world: Value-based methods for discrete planning

First, we present illustrative experiments in a simple grid-world environment to demonstrate the fundamental impact of different recursive reward aggregation functions on learned policies.

Environment Fig. 4a shows the results for a 3×4 grid environment, where an agent navigates from the top-left corner to a fixed goal at the bottom-right corner. As shown in Fig. 4a, the agent receives a small negative reward at each step, which varies across states, and a positive reward upon reaching the terminal state.

Method For this discrete environment, we modified the Q-learning algorithm (Watkins, 1989; Watkins & Dayan, 1992) using the Bellman optimality operator introduced in Section 3.3 (more specifically, the one for the state-action statistic function in Definition D.9). We used four recursive aggregation functions: discounted sum, discounted max, min, and mean, as detailed in Table 1. The detailed algorithm is provided in Algorithm 1 in Appendix G.



Figure 4: **Grid-world**: Fig. 4a shows the discrete environment and the reward function r(s, a), where the agent starts from the top-left corner \bullet and needs to reach the goal at the bottom-right corner \blacksquare . Figs. 4b to 4e show the optimal state-action value functions $q_*(s, a)$ under different aggregations.



Figure 5: **Wind-world**: Fig. 5a shows the continuous environment, where the agent encounters wind disturbances (visualized with streamlines) and receives higher rewards near the center (depicted with colored contours). Fig. 5b illustrates the trajectories of agents trained using different aggregation functions, while Fig. 5c compares the rewards obtained by each agent. Fig. 5d presents the evaluation metrics, highlighting the impact of aggregation functions on performance.

Results Compared to the standard discounted sum aggregation (Fig. 4b), optimizing for the discounted max reward (Fig. 4c) makes the agent indifferent to intermediate costs, favoring shorter paths to the goal. In contrast, the discounted min (Fig. 4d) encourages risk-averse behavior, while the mean aggregation (Fig. 4e) promotes efficiency by maximizing average reward per step. Overall, these results demonstrate how each aggregation function uniquely impacts reward evaluation and policy preferences.

5.2 Wind-world: Policy improvement methods for trajectory optimization

Next, we show that the recursive reward aggregation technique can also be seamlessly integrated into methods for continuous state and action spaces to optimize trajectories in complex environments.

Environment Inspired by Dorfman et al. (2021); Ackermann et al. (2024), we designed a twodimensional continuous environment where an agent navigates to a fixed goal amidst varying wind disturbances, as shown in Fig. 5a. This setup allows us to evaluate the impact of different aggregation functions on trajectory optimization.

Method For this continuous environment, we utilized the Proximal Policy Optimization (PPO) algorithm (Schulman et al., 2017), which is a widely used policy improvement method. We estimated the value function using the Bellman operator for the state statistic function in Definition 3.4. The detailed algorithm is provided in Algorithm 2 in Appendix G.

Results The results in Figs. 5b to 5d show that different aggregation functions lead to distinct trade-offs in trajectory optimization. Specifically, the max aggregation function prioritizes high-reward paths, while the min function ensures more conservative and consistent behavior. The variance-regularized mean aggregation provide balanced strategies, demonstrating the flexibility of the recursive reward aggregation technique in optimizing diverse objectives.



Figure 6: Lunar Lander Continuous: Comparison of five reward aggregation methods. (Left) Radar plot showing performance across seven evaluation metrics, averaged over four random seeds. (Right) Sample trajectories illustrating the qualitative behaviors induced by each aggregation method.

5.3 Physics simulation: Actor-critic methods for continuous control

Then, we extend our evaluation to more complex physics simulation environments.

Environment We conducted experiments on three continuous control environments: (i) Lunar Lander Continuous (Brockman et al., 2016) from the Box2D environment, (ii) Hopper (Erez et al., 2012), and (iii) Ant (Schulman et al., 2016) simulated using MuJoCo (Todorov et al., 2012). A detailed description of these environments can be found in Appendix H.3.

Method In these experiments, we employed the Twin Delayed Deep Deterministic Policy Gradient (TD3) algorithm (Fujimoto et al., 2018), with a modified recursive version detailed in Algorithm 3 in Appendix G. We considered five different reward aggregation functions: (i) discounted sum $(sum_{0.99})$, (ii) discounted max $(max_{0.99})$, (iii) min (min), (iv) discounted sum plus discounted max $(sum_{0.99} + max_{0.99})$, and (v) discounted sum minus variance $(sum_{0.99} - var)$.

Results The results for Lunar Lander Continuous are provided in Fig. 6, with results for other environments in Appendix H.3. As a goal-reaching task, Lunar Lander Continuous reveals how different aggregation strategies influence landing behavior and overall performance.

With the sum_{0.99} aggregation, which serves as the baseline, the agent learns a balanced landing strategy, effectively managing thrust control to achieve a smooth descent while minimizing fuel consumption. In contrast, the $\max_{0.99}$ aggregation encourages the agent to seek high instantaneous rewards, leading to aggressive thrusting behaviors. As a consequence, the lander may exhibit erratic flight patterns, either applying excessive thrust to maximize immediate reward or failing to decelerate properly, which increases the likelihood of hard landings, instability, or even complete mission failure. This outcome underscores the risk of optimizing for short-term reward spikes at the expense of long-term stability and control. The min aggregation demonstrates its effectiveness in risk-averse tasks, as it prioritizes maximizing the worst-case outcomes rather than accumulate reward. The agent adopts a cautious descent strategy, reducing the likelihood of crashes by avoiding sudden thrust changes. Furthermore, since goal-reaching tasks inherently align cumulative and peak rewards, the sum_{0.99} + max_{0.99} aggregation performs similarly to sum_{0.99}, as both encourage stable and efficient landings without introducing significant behavioral differences. Finally, in the sum $_{0.99}$ – var aggregation, the lander remains airborne, ultimately leading to mission termination. This occurs because both successful and failed landings yield large positive or negative rewards, the agent attempts to avoid these extremes, increasing variance and leading to hesitant and inefficient control. This highlights the conflict between variance minimization and goal-reaching tasks, where effective performance relies on high-reward actions often discouraged by variance penalties. These findings emphasize the need to choose aggregation strategies that align with the specific demands of the task.

Table 2: Performance comparison of different methods for portfolio optimization using the Sharpe ratio. The table reports the mean and standard deviation of the Sharpe ratio across five random seeds during the test period, where a higher value indicates better risk-adjusted returns.

	DiffSharpe	NCMDP	Ours
Sharpe Ratio (Test)	0.29 ± 1.22	0.48 ± 0.79	1.12 ± 0.92

5.4 Real-world application: Sharpe ratio in portfolio optimization

Lastly, we evaluate the practical applicability of our method in a real-world application.

Portfolio optimization (Moody et al., 1998; Sood et al., 2023; Liu et al., 2024) is a real-world financial application where an agent (or investor) determines the optimal allocation of assets across different investment options. It can be framed as a sequential decision-making problem as the agent continuously adjusts the portfolio in response to evolving market conditions, fluctuating asset prices, and shifting risk preferences, rather than setting a static allocation. Each decision not only influences immediate returns but also conditions future decisions.

The *Sharpe ratio* (Sharpe, 1966) is a standard metric for evaluating the performance of investment strategies by quantifying the trade-off between return and risk. It is defined as the ratio of the average return (arithmetic mean) to the volatility of return (standard deviation) (Bodie et al., 2011, Eq. (5.18)):

SharpeRatio
$$(r_{1:t}) \coloneqq \frac{\operatorname{mean}(r_{1:t})}{\operatorname{std}(r_{1:t})},$$
(14)

where $r_t := (P_{t+1} - P_t)/P_t$ represents the simple return, and P_t is the portfolio value at time t. Since the Sharpe ratio is non-cumulative, previous RL approaches have relied on the approximate differential Sharpe ratio (Moody et al., 1998; Moody & Saffell, 2001) as a reward signal to facilitate learning. However, this approach introduces an inconsistency between the learning objective and the actual Sharpe ratio, potentially leading to suboptimal policy learning.

Environment This experiment was conducted in a financial market simulation, where an agent learned to optimize portfolio allocations across 11 different S&P 500 sector indices from 2006 to 2021. The environment is the same as that described by Sood et al. (2023); Nägele et al. (2024), with further details provided in Appendix H.4.

Baselines We considered two baseline methods: (i) DiffSharpe (Moody et al., 1998; Moody & Saffell, 2001), which optimizes an approximate differential Sharpe ratio, and (ii) a non-cumulative Markov decision process (NCMDP) method proposed by Nägele et al. (2024), which maps NCMDPs to standard MDPs and defines per-step rewards based on consecutive differences.

Method As demonstrated in Table 1, since both mean and variance admit recursive computation, the Sharpe ratio can also be expressed and updated in a recursive manner. This property allows our method to address the aforementioned inconsistency, aligning the learning objective with the true Sharpe ratio. Our method is built upon the PPO (Schulman et al., 2017) algorithm, with specific modifications on Bellman equation detailed in Algorithm 2 in Appendix G.

Results We conducted experiments across five random seeds, reporting the mean and standard deviation of test performance. Since a higher Sharpe ratio reflects superior risk-adjusted returns, the results in Table 2 and Fig. 17 in Appendix H.4 indicate that our method often attains improved risk-reward balance relative to the baselines. These results illustrate that modifying either the local reward signal or the global performance measure can create misalignment, leading to inconsistencies in policy training and suboptimal outcomes. Unlike baseline methods, our method maintains the original per-step reward structure while estimating and optimizing the exact Sharpe ratio over the entire trajectory. This design may help maintain alignment between training and evaluation, enabling the agent to focus more on long-term performance and become less sensitive to short-term fluctuations.

6 Conclusion

Summary In this paper, we revealed that the recursive structures in the standard MDP can be generalized to a broader class of recursive reward aggregation functions, resulting in generalized Bellman equations and operators. Our theoretical analysis on the existence and uniqueness of fixed points of the generalized Bellman operators provides a solid foundation for designing RL algorithms based on recursive reward aggregation and understanding their convergence properties. Empirical evaluations across discrete and continuous environments confirmed that different aggregation functions significantly influence policy preferences, and we can align the agent behavior with the task requirements by selecting appropriate aggregation functions. These findings highlight the flexibility of recursive reward aggregation, paving the way for more versatile RL algorithms that can be tailored to complex task requirements.

Scope and limitations Our framework is designed for recursive aggregations. As such, it does not directly support non-recursive objectives such as the median or semivariance, which cannot be computed using a bounded-size accumulator in an online fashion with a single pass. Although approximate solutions may be feasible, e.g., sketching algorithms such as online quantile estimation (Greenwald & Khanna, 2001), they fall outside the exact scope of our algebraic formulation.

Additionally, while our method frees the designer from modifying the reward function itself, it introduces a different axis of design: selecting or constructing a suitable aggregation function. The space of meaningful aggregations is vast and may require domain-specific insight or empirical tuning.

Finally, our work is agnostic to the validity of the *reward hypothesis* (Bowling et al., 2023): the idea that all goals can be expressed as the maximization of expected cumulative scalar rewards. We neither rely on this assumption nor seek to refute it. Instead, we explore an orthogonal dimension of goal specification: how reward signals are aggregated over time. This perspective complements traditional reward design and provides a flexible mechanism for aligning behavior with complex objectives, without requiring any claims about the ultimate expressiveness or limitations of scalar rewards.

Future work Future research could explore several extensions and applications of the proposed recursive reward aggregation framework.

First, since the abstract framework does not require the outputs of the generation function and the inputs of the aggregation function to be real values, one promising direction is to investigate the use of *multi-dimensional objectives* or *non-numerical feedback signals*, enhancing the flexibility and expressiveness of policy preferences, particularly in complex environments with intricate reward structures (Pitis, 2023; Wiltzer et al., 2024) or constraints (Gattami et al., 2021; Wachi et al., 2024).

Second, exploring the properties of the generalized Bellman operators in the *stochastic setting*, especially their contraction behavior under different distributional metrics, is an important area of study (Bellemare et al., 2023).

Third, extending the framework to *approximate non-recursive aggregations* (Greenwald & Khanna, 2001) or to *learn aggregation functions* from data (Zaheer et al., 2017; Ong & Veličković, 2022) could broaden its applicability and automate goal specification.

Finally, applying recursive reward aggregation to real-world settings such as (i) risk-sensitive decisionmaking, (ii) risk-adjusted return optimization and portfolio diversification in finance, and (iii) safe, robust, and multi-objective control in robotics and autonomous driving, presents promising directions (Kober et al., 2013; Kiran et al., 2021; Liu et al., 2024).

Acknowledgments

We are grateful to Tongtong Fang for carefully reviewing the abstract and introduction and offering insightful suggestions. We also thank the anonymous reviewers for their constructive feedback.

YT was supported by Institute for AI and Beyond, UTokyo. SN was supported by JSPS KAKENHI Grant Number JP24KJ0818. MS was supported by Institute for AI and Beyond, UTokyo.

References

- Joshua Achiam, David Held, Aviv Tamar, and Pieter Abbeel. Constrained policy optimization. In *International Conference on Machine Learning*, 2017. URL https://proceedings.mlr.press/v70/achiam17a.html.
- Johannes Ackermann, Takayuki Osa, and Masashi Sugiyama. Offline reinforcement learning from datasets with structured non-stationarity. In *Reinforcement Learning Conference*, 2024. URL https://openreview.net/forum?id=qowNlhKcPw.
- Dario Amodei, Chris Olah, Jacob Steinhardt, Paul Christiano, John Schulman, and Dan Mané. Concrete problems in AI safety. *arXiv preprint*, 2016. URL https://arxiv.org/abs/16 06.06565.
- Marcin Andrychowicz, Filip Wolski, Alex Ray, Jonas Schneider, Rachel Fong, Peter Welinder, Bob McGrew, Josh Tobin, OpenAI Pieter Abbeel, and Wojciech Zaremba. Hindsight experience replay. In *Neural Information Processing Systems*, 2017. URL https://proceedings.neurips. cc/paper/2017/hash/453fadbd8a1a3af50a9df4df899537b5-Abstract.ht ml.
- Saurabh Arora and Prashant Doshi. A survey of inverse reinforcement learning: Challenges, methods and progress. *Artificial Intelligence*, 297:103500, 2021. URL https://doi.org/10.1016/j.artint.2021.103500.
- David H Bailey, Jonathan Borwein, Marcos Lopez de Prado, and Qiji Jim Zhu. The probability of backtest overfitting. *Journal of Computational Finance (Risk Journals)*, 2015. URL https://dx.doi.org/10.2139/ssrn.2326253.
- Leemon C. Baird. Reinforcement learning in continuous time: Advantage updating. In *IEEE International Conference on Neural Networks*, volume 4, pp. 2448–2453. IEEE, 1994. URL https://doi.org/10.1109/ICNN.1994.374604.
- Andrew G. Barto, Richard S. Sutton, and Charles W. Anderson. Neuronlike adaptive elements that can solve difficult learning control problems. *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-13(5):834–846, 1983. URL https://doi.org/10.1109/TSMC.1983.6313077.
- Marc G Bellemare, Will Dabney, and Rémi Munos. A distributional perspective on reinforcement learning. In *International Conference on Machine Learning*, 2017. URL https://proceedings.mlr.press/v70/bellemare17a.html.
- Marc G Bellemare, Will Dabney, and Mark Rowland. *Distributional Reinforcement Learning*. MIT Press, 2023. URL https://doi.org/10.7551/mitpress/14207.001.0001.
- Richard Bellman. Dynamic programming. *Science*, 153(3731):34–37, 1966. URL https://doi.org/10.1126/science.153.3731.34.

Dimitri Bertsekas. Abstract Dynamic Programming. Athena Scientific, 2022.

Richard Bird and Oege de Moor. Algebra of Programming. Prentice Hall, 1997.

Zvi Bodie, Alex Kane, and Alan J Marcus. Investments. McGraw-hill, 2011.

- Michael Bowling, John D Martin, David Abel, and Will Dabney. Settling the reward hypothesis. In *International Conference on Machine Learning*, 2023. URL https://proceedings.mlr.press/v202/bowling23a.html.
- James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. JAX: composable transformations of Python+NumPy programs, 2018. URL http://github.com/jax-ml/jax.
- Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. OpenAI Gym. *arXiv preprint*, 2016. URL https://arxiv.org/abs/1606.01540.
- Paul F Christiano, Jan Leike, Tom Brown, Miljan Martic, Shane Legg, and Dario Amodei. Deep reinforcement learning from human preferences. In *Neural Information Processing Systems*, volume 30, 2017. URL https://proceedings.neurips.cc/paper/2017/hash/d 5e2c0adad503c91f91df240docd4e49-Abstract.html.
- Geoffrey SH Cruttwell, Bruno Gavranović, Neil Ghani, Paul Wilson, and Fabio Zanasi. Categorical foundations of gradient-based learning. In *European Symposium on Programming*, pp. 1–28, 2022. URL https://doi.org/10.1007/978-3-030-99336-8_1.
- Wei Cui and Wei Yu. Reinforcement learning with non-cumulative objective. *IEEE Transactions* on Machine Learning in Communications and Networking, 1:124–137, 2023. URL https: //doi.org/10.1109/TMLCN.2023.3285543.
- Oege De Moor. Categories, relations and dynamic programming. *Mathematical Structures in Computer Science*, 4(1):33–69, 1994. URL https://doi.org/10.1017/S096012950 0000360.
- Lauro Langosco Di Langosco, Jack Koch, Lee D Sharkey, Jacob Pfau, and David Krueger. Goal misgeneralization in deep reinforcement learning. In *International Conference on Machine Learning*, 2022. URL https://proceedings.mlr.press/v162/langosco22a.h tml.
- Ron Dorfman, Idan Shenfeld, and Aviv Tamar. Offline meta reinforcement learning identifiability challenges and effective data collection strategies. In *Neural Information Processing Systems*, 2021. URL https://openreview.net/forum?id=IBdEfhLveS.
- Tom Erez, Yuval Tassa, and Emanuel Todorov. Infinite-horizon model predictive control for periodic tasks with contacts. In *Robotics: Science and Systems VII*. The MIT Press, 2012. URL https://doi.org/10.7551/mitpress/9481.003.0015.
- Brendan Fong and Michael Johnson. Lenses and learners. In *International Workshop on Bidirectional Transformations*, 2019. URL https://arxiv.org/abs/1903.03671.
- Tobias Fritz. A synthetic approach to Markov kernels, conditional independence and theorems on sufficient statistics. *Advances in Mathematics*, 370:107239, 2020. URL https://doi.org/10.1016/j.aim.2020.107239. https://arxiv.org/abs/1908.07021.
- Scott Fujimoto, Herke Hoof, and David Meger. Addressing function approximation error in actor-critic methods. In *International Conference on Machine Learning*, 2018. URL https://proceedings.mlr.press/v80/fujimoto18a.html.
- Ather Gattami, Qinbo Bai, and Vaneet Aggarwal. Reinforcement learning for constrained Markov decision processes. In *International Conference on Artificial Intelligence and Statistics*, 2021. URL https://proceedings.mlr.press/v130/gattami21a.html.

- Bruno Gavranović. Space-time tradeoffs of lenses and optics via higher category theory. *arXiv* preprint, 2022. URL https://arxiv.org/abs/2209.09351.
- Michèle Giry. A categorical approach to probability theory. *Categorical Aspects of Topology and Analysis*, pp. 68–85, 1982. URL https://doi.org/10.1007/BFb0092872.
- David Goldberg. What every computer scientist should know about floating-point arithmetic. ACM computing surveys (CSUR), 23(1):5–48, 1991. URL https://doi.org/10.1145/103162.103163.
- Sai Krishna Gottipati, Yashaswi Pathak, Rohan Nuttall, Raviteja Chunduru, Ahmed Touati, Sriram Ganapathi Subramanian, Matthew E Taylor, and Sarath Chandar. Maximum reward formulation in reinforcement learning. *arXiv preprint*, 2020. URL https://arxiv.org/ab s/2010.03744.
- Michael Greenwald and Sanjeev Khanna. Space-efficient online computation of quantile summaries. *ACM SIGMOD Record*, 30(2):58–66, 2001. URL https://doi.org/10.1145/376284 .375670.
- William H. Guss, Brandon Houghton, Nicholay Topin, Phillip Wang, Cayden Codel, Manuela Veloso, and Ruslan Salakhutdinov. MineRL: A large-scale dataset of Minecraft demonstrations. In Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, 2019. URL https://doi.org/10.24963/ijcai.2019/339.
- Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft Actor-Critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International Conference on Machine Learning*, 2018. URL https://proceedings.mlr.press/v80/haarnoja18b.html.
- Dylan Hadfield-Menell, Anca Dragan, Pieter Abbeel, and Stuart Russell. The off-switch game. In Workshop on AI, Ethics, and Society at the Thirty-First AAAI Conference on Artificial Intelligence, 2017. URL https://arxiv.org/abs/1611.08219.
- Charles R. Harris, K. Jarrod Millman, Stéfan J. van der Walt, Ralf Gommers, Pauli Virtanen, David Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J. Smith, Robert Kern, Matti Picus, Stephan Hoyer, Marten H. van Kerkwijk, Matthew Brett, Allan Haldane, Jaime Fernández del Río, Mark Wiebe, Pearu Peterson, Pierre Gérard-Marchant, Kevin Sheppard, Tyler Reddy, Warren Weckesser, Hameer Abbasi, Christoph Gohlke, and Travis E. Oliphant. Array programming with NumPy. *Nature*, 585(7825):357–362, 2020. URL https://doi.org/10.1038/s415 86-020-2649-2. https://numpy.org.
- Jules Hedges and Riu Rodríguez Sakamoto. Value iteration is optic composition. In *International Conference on Applied Category Theory*, 2022. URL https://arxiv.org/abs/2206.04547.
- Ralf Hinze, Thomas Harper, and Daniel W. H. James. Theory and practice of fusion. In *Symposium on Implementation and Application of Functional Languages*, pp. 19–37, 2010. URL https://doi.org/10.1007/978-3-642-24276-2_2.
- Ahmed Hussein, Mohamed Medhat Gaber, Eyad Elyan, and Chrisina Jayne. Imitation learning: A survey of learning methods. *ACM Computing Surveys (CSUR)*, 50(2):1–35, 2017. URL https://doi.org/10.1145/3054912.
- Graham Hutton. A tutorial on the universality and expressiveness of fold. *Journal of Functional Programming*, 9(4):355–372, 1999. URL https://doi.org/10.1017/S09567968990 03500.

- Jiaming Ji, Tianyi Qiu, Boyuan Chen, Borong Zhang, Hantao Lou, Kaile Wang, Yawen Duan, Zhonghao He, Jiayi Zhou, Zhaowei Zhang, Fanzhi Zeng, Kwan Yee Ng, Juntao Dai, Xuehai Pan, Aidan O'Gara, Yingshan Lei, Hua Xu, Brian Tse, Jie Fu, Stephen McAleer, Yaodong Yang, Yizhou Wang, Song-Chun Zhu, Yike Guo, and Wen Gao. AI alignment: A comprehensive survey. *arXiv preprint*, 2023. URL https://arxiv.org/abs/2310.19852.
- Taylan Kabbani and Ekrem Duman. Deep reinforcement learning approach for trading automation in the stock market. *IEEE Access*, 10:93564–93574, 2022. URL https://doi.org/10.1109/ACCESS.2022.3203697.
- B Ravi Kiran, Ibrahim Sobh, Victor Talpaert, Patrick Mannion, Ahmad A Al Sallab, Senthil Yogamani, and Patrick Pérez. Deep reinforcement learning for autonomous driving: A survey. *IEEE Transactions on Intelligent Transportation Systems*, 23(6):4909–4926, 2021. URL https: //doi.org/10.1109/TITS.2021.3054625.
- Jens Kober, J Andrew Bagnell, and Jan Peters. Reinforcement learning in robotics: A survey. *The International Journal of Robotics Research*, 32(11):1238–1274, 2013. URL https://doi.org/10.1177/0278364913495721.
- Robert Tjarko Lange. gymnax: A JAX-based reinforcement learning environment library, 2022. URL http://github.com/RobertTLange/gymnax.
- Jan Leike, Miljan Martic, Victoria Krakovna, Pedro A Ortega, Tom Everitt, Andrew Lefrancq, Laurent Orseau, and Shane Legg. AI safety gridworlds. *arXiv preprint*, 2017. URL https: //arxiv.org/abs/1711.09883.
- Xiao-Yang Liu, Ziyi Xia, Hongyang Yang, Jiechao Gao, Daochen Zha, Ming Zhu, Christina Dan Wang, Zhaoran Wang, and Jian Guo. Dynamic datasets and market environments for financial reinforcement learning. *Machine Learning*, 113(5):2795–2839, 2024. URL https://doi.org/10.1007/S10994-023-06511-w. https://arxiv.org/abs/2304.13174.
- Shie Mannor and John Tsitsiklis. Mean-variance optimization in Markov decision processes. In International Conference on Machine Learning, 2011. URL https://dl.acm.org/doi/a bs/10.5555/3104482.3104505. https://icml.cc/2011/papers/156_icml paper.pdf.
- Erik Meijer, Maarten Fokkinga, and Ross Paterson. Functional programming with bananas, lenses, envelopes and barbed wire. In *Conference on Functional Programming Languages and Computer Architecture*, pp. 124–144, 1991. URL https://doi.org/10.1007/3540543961_7.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing Atari with deep reinforcement learning. *arXiv preprint*, 2013. URL https://arxiv.org/abs/1312.5602.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, Martin Riedmiller, Andreas K. Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dharshan Kumaran, Daan Wierstra, Shane Legg, and Demis Hassabis. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015. URL https://doi.org/10.1038/nature14236.
- Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In *International Conference on Machine Learning*, 2016. URL https://proceedings.mlr.press/v48/mniha16.html.
- John Moody and Matthew Saffell. Learning to trade via direct reinforcement. *IEEE transactions on neural Networks*, 12(4):875–889, 2001. URL https://doi.org/10.1109/72.935097.

- John Moody, Lizhong Wu, Yuansong Liao, and Matthew Saffell. Performance functions and reinforcement learning for trading systems and portfolios. *Journal of forecasting*, 17(5-6):441–470, 1998. URL https://doi.org/10.1002/(SICI)1099-131X(1998090)17: 5/6%3C441::AID-FOR707%3E3.0.C0;2-%23.
- Tetsuro Morimura, Masashi Sugiyama, Hisashi Kashima, Hirotaka Hachiya, and Toshiyuki Tanaka. Nonparametric return distribution approximation for reinforcement learning. In *International Conference on Machine Learning*, 2010a. URL https://dblp.org/rec/conf/icml/M orimuraSKHT10.html. https://icml.cc/2010/papers/652.pdf.
- Tetsuro Morimura, Masashi Sugiyama, Hisashi Kashima, Hirotaka Hachiya, and Toshiyuki Tanaka. Parametric return density estimation for reinforcement learning. In *Conference on Uncertainty in Artificial Intelligence*, 2010b. URL https://dblp.org/rec/conf/uai/MorimuraSK HT10.html. https://event.cwi.nl/uai2010/papers/UAI2010_0115.pdf.
- Jean-Michel Muller, Nicolas Brunie, Florent De Dinechin, Claude-Pierre Jeannerod, Mioara Joldes, Vincent Lefèvre, Guillaume Melquiond, Nathalie Revol, and Serge Torres. *Handbook of floating-point arithmetic*, volume 1. Springer, 2018. URL https://doi.org/10.1007/978-3-3 19-76526-6.
- Kevin Murphy. Reinforcement learning: An overview. *arXiv preprint*, 2024. URL https://arxiv.org/abs/2412.05265.
- Maximilian Nägele, Jan Olle, Thomas Fösel, Remmy Zen, and Florian Marquardt. Tackling decision processes with non-cumulative objectives using reinforcement learning. *arXiv preprint*, 2024. URL https://arxiv.org/abs/2405.13609.
- Andrew Y. Ng, Daishi Harada, and Stuart J. Russell. Policy invariance under reward transformations: Theory and application to reward shaping. In *International Conference on Machine Learning*, 1999. URL https://dl.acm.org/doi/10.5555/645528.657613.
- Euan Ong and Petar Veličković. Learnable commutative monoids for graph neural networks. In *Learning on Graphs Conference*, 2022. URL https://proceedings.mlr.press/v198 /ong22a.html.
- Aleksandar Petrov. Compositional computational systems. Master's thesis, ETH Zurich, 2020. URL https://doi.org/10.3929/ethz-b-000463467.
- Silviu Pitis. Consistent aggregation of objectives with diverse time preferences requires non-Markovian rewards. *Advances in Neural Information Processing Systems*, 2023. URL https: //proceedings.neurips.cc/paper/2023/hash/08342dc6ab69f23167b4123 086ad4d38-Abstract.html.
- Martin L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, Inc., 1994. URL https://doi.org/10.1002/9780470316887.
- Kian Hong Quah and Chai Quek. Maximum reward reinforcement learning: A non-cumulative reward criterion. *Expert Systems with Applications*, 31(2):351–359, 2006. URL https://doi.org/10.1016/j.eswa.2005.09.054.
- Antonin Raffin, Ashley Hill, Adam Gleave, Anssi Kanervisto, Maximilian Ernestus, and Noah Dormann. Stable-Baselines3: Reliable reinforcement learning implementations. *Journal of Machine Learning Research*, 22(268):1–8, 2021. URL http://jmlr.org/papers/v22/20-1364.html. https://github.com/DLR-RM/stable-baselines3.
- Mitchell Riley. Categories of optics. *arXiv preprint*, 2018. URL https://arxiv.org/abs/ 1809.00738.

- John Schulman, Philipp Moritz, Sergey Levine, Michael I. Jordan, and Pieter Abbeel. Highdimensional continuous control using generalized advantage estimation. In *International Conference on Learning Representations*, 2016. URL https://arxiv.org/abs/15 06.02438.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint*, 2017. URL https://arxiv.org/abs/1707.06347.
- William F. Sharpe. Mutual fund performance. The Journal of Business, 39(1):119–138, 1966. URL http://www.jstor.org/stable/2351741.
- David Silver, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, Laurent Sifre, Dharshan Kumaran, Thore Graepel, Timothy Lillicrap, Karen Simonyan, and Demis Hassabis. A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play. *Science*, 362(6419):1140–1144, 2018. URL https://doi.org/10.1126/science.aar6404.
- Toby St. Clere Smithe. Bayesian updates compose optically. *arXiv preprint*, 2020. URL https://arxiv.org/abs/2006.01631.
- Matthew J. Sobel. The variance of discounted Markov decision processes. *Journal of Applied Probability*, 19(4):794–802, 1982. URL https://doi.org/10.2307/3213832.
- Srijan Sood, Kassiani Papasotiriou, Marius Vaiciulis, and Tucker Balch. Deep reinforcement learning for optimal portfolio allocation: A comparative study with mean-variance optimization. *FinPlan*, pp. 21, 2023. URL https://icaps23.icaps-conference.org/papers/finplan /FinPlan23_paper_4.pdf.
- Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction.* The MIT Press, 1998. URL http://incompleteideas.net/book/the-book.html.
- Aviv Tamar, Dotan Di Castro, and Shie Mannor. Policy gradients with variance related risk criteria. In *International Conference on Machine Learning*, 2012. URL https://dl.acm.org/doi /10.5555/3042573.3042784. https://icml.cc/2012/papers/489.pdf.
- Emanuel Todorov, Tom Erez, and Yuval Tassa. MuJoCo: A physics engine for model-based control. In 2012 IEEE/RSJ International Conference on Intelligent Robots and Systems, 2012. URL https://doi.org/10.1109/IROS.2012.6386109.
- Grigorii Veviurko, Wendelin Böhmer, and Mathijs de Weerdt. To the max: Reinventing reward in reinforcement learning. In *International Conference on Machine Learning*, 2024. URL https://proceedings.mlr.press/v235/veviurko24a.html.
- Akifumi Wachi, Xun Shen, and Yanan Sui. A survey of constraint formulations in safe reinforcement learning. In *International Joint Conference on Artificial Intelligence*, 2024. URL https://doi.org/10.24963/ijcai.2024/913.
- Ruosong Wang, Peilin Zhong, Simon S Du, Russ R Salakhutdinov, and Lin Yang. Planning with general objective functions: Going beyond total rewards. In *Neural Information Processing Systems*, 2020. URL https://proceedings.neurips.cc/paper/2020/hash/a6a 767bbb2e3513233f942e0ff24272c-Abstract.html.
- Christopher J. C. H. Watkins and Peter Dayan. Q-learning. *Machine Learning*, 8:279–292, 1992. URL https://doi.org/10.1007/BF00992698.
- Christopher John Cornish Hellaby Watkins. *Learning from delayed rewards*. PhD thesis, King's College, Cambridge United Kingdom, 1989. URL http://www.cs.rhul.ac.uk/~chris w/new_thesis.pdf.

- Barry Payne Welford. Note on a method for calculating corrected sums of squares and products. *Technometrics*, 4(3):419–420, 1962. URL https://doi.org/10.1080/00401706.196 2.10490022.
- Harley Wiltzer, Jesse Farebrother, Arthur Gretton, and Mark Rowland. Foundations of multivariate distributional reinforcement learning. In *Neural Information Processing Systems*, 2024. URL https://proceedings.neurips.cc/paper/2024/hash/b76bec34ef5e0c0ce edff6edfbefc9f5-Abstract.html.
- Xing Wu, Haolei Chen, Jianjia Wang, Luigi Troiano, Vincenzo Loia, and Hamido Fujita. Adaptive stock trading strategies with deep reinforcement learning methods. *Information Sciences*, 538: 142–158, 2020. URL https://doi.org/10.1016/j.ins.2020.05.066.
- Zhixuan Yang and Nicolas Wu. Fantastic morphisms and where to find them: A guide to recursion schemes. In *International Conference on Mathematics of Program Construction*, pp. 222–267, 2022. URL https://doi.org/10.1007/978-3-031-16912-0_9.
- Manzil Zaheer, Satwik Kottur, Siamak Ravanbakhsh, Barnabas Poczos, Russ R Salakhutdinov, and Alexander J Smola. Deep sets. In *Neural Information Processing Systems*, 2017. URL https://proceedings.neurips.cc/paper/2017/hash/f22e4747da1aa27e3 63d86d4off442fe-Abstract.html.
- Meixin Zhu, Yinhai Wang, Ziyuan Pu, Jingyun Hu, Xuesong Wang, and Ruimin Ke. Safe, efficient, and comfortable velocity control based on reinforcement learning for autonomous driving. *Transportation Research Part C: Emerging Technologies*, 117:102662, 2020. URL https://doi.org/10.1016/j.trc.2020.102662.
- Brian D. Ziebart, Andrew Maas, J. Andrew Bagnell, and Anind K. Dey. Maximum entropy inverse reinforcement learning. In *National Conference on Artificial Intelligence*, 2008. URL https://dl.acm.org/doi/abs/10.5555/1620270.1620297.

Supplementary Materials

The following content was not necessarily subject to peer review.

Contents

1	Intro	oduction 1
2	An a	lgebraic perspective on Bellman equations 3
	2.1	Preliminaries
	2.2	Recursive generation of rewards
	2.3	Recursive aggregation of rewards
	2.4	Bellman equation for the state value function
3	Recu	Irsive reward aggregation functions5
	3.1	Bellman equation for the state statistic function
	3.2	Policy evaluation: Iterative statistic function estimation 6
	3.3	Policy optimization: Optimal policies and optimal value functions
4	Fron	n deterministic to stochastic Markov decision processes 7
5	Expo	eriments 8
	5.1	Grid-world: Value-based methods for discrete planning
	5.2	Wind-world: Policy improvement methods for trajectory optimization 9
	5.3	Physics simulation: Actor-critic methods for continuous control 10
	5.4	Real-world application: Sharpe ratio in portfolio optimization 11
6	Con	clusion 12
Bil	bliogr	raphy 13
А	State	e-action recursion 22
	A.1	State-action transition
	A.2	State-action step function and generation function
	A.3	State-action statistic function and value function
	A.4	Relationship between state and state-action statistic functions
	A.5	Advantage function
B	Alge	braic structures in Markov decision process 27
	B.I	Algebra fusion
	B.2	Bidirectional process
G	В.3	Non-uniqueness of update function and post-processing function
C	Met	Preliminarian 29
	C.1	Preniminaries
	C.2	Metrics on statistics and rewards 29 Ballman operators 30
	C.5	Existence of fixed points of Bellman operators 30
	C.5	Uniqueness of fixed points of Bellman operators
D	Ord	ers and Bellman ontimality onerators 31
2	D.1	Preliminaries
	D.2	Orders on statistics and rewards
	D.3	Bellman optimality operators
	D.4	Existence of fixed points of Bellman optimality operators
	D.5	Uniqueness of fixed points of Bellman optimality operators
Е	Stoc	hastic Markov decision process 34
	E.1	Composition of stochastic functions
	E.2	Stochastic recursion
	E.3	Relationship between stochastic state and state-action statistic functions 35
	E.4	Expected aggregated rewards vs. aggregated expected rewards

F	Proofs	36
G	Learning algorithms with recursive reward aggregation	43
	G.1 Q-learning	43
	G.2 PPO	44
	G.3 TD3	45
H	Experiments	46
	H.1 Grid-world environment	46
	H.2 Wind-world environment	46
	H.3 Continuous control environments	46
	H.4 Portfolio environment	49
I	Discussion	51

List of Figures

1	Different aggregation functions lead to different policy preferences	2
2	State statistic bidirectional process	6
3	Expected discounted sum of rewards vs. discounted sum of expected rewards	8
4	Grid-world	9
5	Wind-world	9
6	Lunar Lander	10
7	State statistic bidirectional process $\tau_{\pi}^{S}: S \to T$	25
8	State statistic bidirectional process (with different behavior and target policies)	25
9	State statistic bidirectional process (with state as the residual)	25
10	State-action statistic bidirectional process $\tau_{\pi}^{S \times A} : S \times A \to T$	26
11	State-action statistic bidirectional process (with different behavior and target policies)	26
12	State-action statistic bidirectional process (with state-action as the residual)	26
13	Relationship between state and state-action statistic functions	41
14	$\max - \lambda \operatorname{range} = \lambda \min + (1 - \lambda) \max \dots \dots \dots \dots \dots \dots \dots \dots \dots $	46
15	Hopper	48
16	Ant	48
17	Sharpe ratio	50

List of Tables

1	Recursive aggregation functions	5
2	Sharpe ratio	11
3	Properties of metrics	29
4	Properties of orders	31
5	Fixed points of the Bellman operators and the Bellman optimality operators	33
6	Expected aggregated rewards vs. aggregated expected rewards: maximum as an example	35

List of Algorithms

1	Q-learning (Watkins & Dayan, 1992) with recursive reward aggregation	43
2	PPO (Schulman et al., 2017) with recursive reward aggregation	44
3	TD3 (Fujimoto et al., 2018) with recursive reward aggregation	45

State-action recursion A

In Section 2, we introduced the recursive generation of rewards by iterating over *states* S. In this section, we extend this framework to iterate over state-action pairs $S \times A$, which is crucial for defining the state-action value function $q_{\pi} : S \times A \to R$.

A.1 State-action transition

First, note that both *pre-composing* and *post-composing* the pairing function $(\text{id}_S, \pi) : S \to S \times A$ with the transition function $p: S \times A \rightarrow S$ yield transition functions:

- state transition p^S_π: S → S := p ∘ (id_S, π) = [s ↦ p(s, π(s))] and
 state-action transition p^{S × A}_π: S × A → S × A := (id_S, π) ∘ p = [s, a ↦ p(s, a), π(p(s, a))].

We use the superscripts S and $S \times A$ to indicate the domains/codomains of these transition functions.

A.2 State-action step function and generation function

Then, following the definitions of the state step function $\operatorname{step}_{\pi,\mathbf{p},\mathbf{r},\omega}^{S}: S \to \{*\} + R \times S$ in Eq. (3) and generation function $\operatorname{gen}_{\pi, p, r, \omega}^{S} : S \to [R]$ in Eq. (4), we can define the *state-action step/generation functions* using the state-action transition $p_{\pi}^{S \times A}$ and the reward function r:

$$\operatorname{step}_{\pi,\mathrm{p},\mathrm{r},\omega}^{S \times A} : S \times A \to \{*\} + R \times (S \times A) \coloneqq \left[s, a \mapsto \begin{cases} * & s \in S_{\omega} \\ (\mathrm{r}(s,a), \mathrm{p}_{\pi}^{S \times A}(s,a)) & s \notin S_{\omega} \end{cases} \right],$$
(15)

$$\operatorname{gen}_{\pi,\mathrm{p,r},\omega}^{S \times A} : S \times A \to [R] := \left[s, a \mapsto \begin{cases} [] & s \in S_{\omega} \\ \operatorname{cons}(\mathbf{r}(s,a), \operatorname{gen}_{\pi,\mathrm{p,r},\omega}^{S \times A}(\mathbf{p}_{\pi}^{S \times A}(s,a))) & s \notin S_{\omega} \end{cases} \right].$$
(16)

A.3 State-action statistic function and value function

Applying the same algebraic fusion technique (Hinze et al., 2010) used for the state statistic function $\tau_{\pi}^{S}: S \to T$ in Theorem 3.2, we can define the *state-action statistic function* $\tau_{\pi}^{S \times A}: S \times A \to T$ and derive its corresponding Bellman equation as follows:

Theorem A.1 (Bellman equation for the state-action statistic function). Given a recursive reward generation function $gen_{\pi,p,r,\omega}^{S \times A}$ and a recursive statistic aggregation function $agg_{init, \triangleright}$ (Definition 3.1), their composition, called the state-action statistic function $\tau_{\pi}^{S \times A} : S \to T$, satisfies

$$\begin{aligned}
\tau_{\pi}^{S \times A} : S \times A \to T &:= \operatorname{agg}_{\operatorname{init}, \triangleright} \circ \operatorname{gen}_{\pi, p, r, \omega}^{S \times A} \\
&= \left[s, a \mapsto \begin{cases} \operatorname{init} & s \in S_{\omega} \\ \operatorname{r}(s, a) \triangleright \tau_{\pi}^{S \times A}(\operatorname{p}_{\pi}^{S \times A}(s, a)) & s \notin S_{\omega} \end{cases} \right].
\end{aligned} \tag{17}$$

Similarly, the *state-action value function* $q_{\pi}: S \times A \to R := \text{post} \circ \tau_{\pi}^{S \times A}$ is the composition of the state-action statistic function $\tau_{\pi}^{S \times A}: S \times A \to T$ with the post-processing function $\text{post}: T \to R$.

A.4 Relationship between state and state-action statistic functions

We can now state the theorem that relates the state and state-action statistic functions:

Theorem A.2 (Relationship between state and state-action statistic functions). *Given a recursive* reward generation function $gen_{\pi,p,r,\omega}$ (Definition 2.1) and a recursive statistic aggregation function $\operatorname{agg}_{\operatorname{init},\triangleright}$ (Definition 3.1), the state statistic function $\tau_{\pi}^{S}: S \to T$ in Eq. (10) and the state-action statistic function $\tau_{\pi}^{S \times A} : S \times A \to T$ in Eq. (17) satisfy

$$\tau_{\pi}^{S} = \tau_{\pi}^{S \times A} \circ \langle \mathrm{id}_{S}, \pi \rangle : S \to T \qquad (for all states), \tag{18}$$

$$\tau_{\pi}^{S \times A} = \mathbf{r} \triangleright (\tau_{\pi}^{S} \circ \mathbf{p}) : S \times A \to T \qquad (for all non-terminal states). \tag{19}$$

Corollary A.3 (Relationship between state and state-action value functions). The state value function $v_{\pi}: S \to R$ and the state-action value function $q_{\pi}: S \times A \to R$ satisfy

$$\mathbf{v}_{\pi} = \mathbf{q}_{\pi} \circ \langle \mathrm{id}_S, \pi \rangle : S \to R.$$
⁽²⁰⁾

In summary, the relationships between the state/state-action step, generation, statistic, and value functions are shown in the following diagram:



A.5 Advantage function

The advantage function (Baird, 1994),

÷

$$\mathbf{u}_{\pi}: S \times A \to R \coloneqq \mathbf{q}_{\pi} - \mathbf{v}_{\pi} \circ p_1 = [s, a \mapsto \mathbf{q}_{\pi}(s, a) - \mathbf{v}_{\pi}(s)], \tag{22}$$

is defined as the difference between the state-action value function $q_{\pi} : S \times A \to R$ and the state value function $v_{\pi} : S \to R$, where $p_1 : S \times A \to S$ is the projection function that extracts the state from a state-action pair. The advantage function measures the advantage of taking an action a in a state s over the average value of all actions in that state following the policy π , which is widely used in RL algorithms.

For a general recursive reward aggregation function post $\circ \operatorname{agg}_{\operatorname{init},\triangleright}$, the advantage function can be expressed using the state-action statistic function $\tau_{\pi}^{S \times A} : S \times A \to T$ and the state statistic function $\tau_{\pi}^{S} : S \to T$ as follows:

$$\alpha_{\pi}: S \times A \to R = \left[s, a \mapsto \text{post}(\tau_{\pi}^{S \times A}(s, a)) - \text{post}(\tau_{\pi}^{S}(s))\right]$$
(23)

$$= \left[s, a \mapsto \begin{cases} 0 & s \in S_{\omega} \\ \operatorname{post}(\mathbf{r}(s, a) \triangleright \tau_{\pi}^{S}(\mathbf{p}(s, a))) - \operatorname{post}(\tau_{\pi}^{S}(s)) & s \notin S_{\omega} \end{cases} \right].$$
(24)

Because the statistic function can be computed recursively, given a sequence of states, rewards, and statistics, we can obtain a sequence of advantage estimators:

$$\hat{\alpha}_t^{(1)} = \text{post}(r_t \triangleright \tau_{t+1}) - \text{post}(\tau_t), \tag{25}$$

$$\hat{\alpha}_t^{(2)} = \text{post}(r_t \triangleright r_{t+1} \triangleright \tau_{t+2}) - \text{post}(\tau_t), \tag{26}$$

$$\hat{\alpha}_t^{(3)} = \text{post}(r_t \triangleright r_{t+1} \triangleright r_{t+2} \triangleright \tau_{t+3}) - \text{post}(\tau_t), \tag{27}$$

$$\hat{\alpha}_{t+1}^{(1)} = \text{post}(r_{t+1} \triangleright \tau_{t+2}) - \text{post}(\tau_{t+1}), \tag{28}$$

$$\hat{\alpha}_{t+1}^{(2)} = \text{post}(r_{t+1} \triangleright r_{t+2} \triangleright \tau_{t+3}) - \text{post}(\tau_{t+1}),$$
:
(29)

The generalized advantage estimator (GAE) proposed by Schulman et al. (2016) combines these advantage estimators with a discount factor $\lambda \in [0, 1]$:

$$\hat{\alpha}_t := \hat{\alpha}_t^{(1)} + \lambda \hat{\alpha}_t^{(2)} + \lambda^2 \hat{\alpha}_t^{(3)} + \cdots$$
(30)

$$= 1 \quad (\operatorname{post}(r_t \qquad \rhd \ \tau_{t+1}) - \operatorname{post}(\tau_t)) \tag{31}$$
$$+ \lambda \quad (\operatorname{post}(r_t \rhd \ r_{t+1} \qquad \rhd \ \tau_{t+2}) - \operatorname{post}(\tau_t))$$
$$+ \lambda^2 \quad (\operatorname{post}(r_t \rhd \ r_{t+1} \rhd \ r_{t+2} \rhd \ \tau_{t+3}) - \operatorname{post}(\tau_t))$$
$$+ \cdots$$

The original GAE formulation (Schulman et al., 2016) considered only the discounted sum and an infinite horizon. For a finite horizon Ω , the advantage estimator can be expressed as follows:

$$\begin{aligned}
\alpha_{t} &= 1 & (r_{t} + \gamma v_{t+1} - v_{t}) & (32) \\
&+ \lambda & (r_{t} + \gamma r_{t+1} + \gamma^{2} r_{t+2} + \gamma^{2} v_{t+2} - v_{t}) \\
&+ \lambda^{2} & (r_{t} + \gamma r_{t+1} + \gamma^{2} r_{t+2} + \gamma^{2} v_{t+3} - v_{t}) \\
&+ \cdots \\
&+ \lambda^{\Omega - t - 1} & (r_{t} + \gamma r_{t+1} + \gamma^{2} r_{t+2} + \cdots + \gamma^{\Omega - t - 1} r_{\Omega - 1} + \gamma^{\Omega - t} v_{\Omega} - v_{t}) \\
&= \sum_{i=0}^{\Omega - t - 1} \lambda^{i} \gamma^{i} \left(\frac{1 - \lambda^{\Omega - t - i}}{1 - \lambda} r_{t+i} + \gamma v_{t+i+1} \right) - \frac{1 - \lambda^{\Omega - t}}{1 - \lambda} v_{t},
\end{aligned}$$
(32)

which has a recursive form:

$$\alpha_t = \frac{1 - \lambda^{\Omega - t}}{1 - \lambda} (r_t + \gamma v_{t+1} - v_t) + \lambda \gamma \alpha_{t+1}.$$
(34)

However, when considering a general recursive reward aggregation function $post \circ agg_{init,\triangleright}$, a recursive expression for the advantage estimator is not always available. Therefore, the advantage estimator may need to be computed directly using its original definition in Eq. (30).



Figure 7: State statistic bidirectional process $\tau_\pi^S:S\to T$



Figure 8: State statistic bidirectional process (with different behavior and target policies)



Figure 9: State statistic bidirectional process (with state as the residual)



Figure 10: State-action statistic bidirectional process $\tau_{\pi}^{S\times A}:S\times A\rightarrow T$



Figure 11: State-action statistic bidirectional process (with different behavior and target policies)



Figure 12: State-action statistic bidirectional process (with state-action as the residual)

B Algebraic structures in Markov decision process

In this section, we briefly discuss the algebraic structures used in this work. For a tutorial on algebraic programming, we refer the reader to Hutton (1999). For a theoretical treatment of algebra fusion, see Hinze et al. (2010). For an accessible and illustrative introduction to bidirectional processes, we recommend Gavranović (2022).

B.1 Algebra fusion

In this work, we mainly considered algebras and coalgebras of signature $\{*\} + R \times (-)$, i.e., lists of rewards. An *algebra* is a pair (A, f) consisting of a carrier set A and a function $f : \{*\} + R \times A \to A$. A *coalgebra* is a pair (C, g) consisting of a carrier set C and a function $g : C \to \{*\} + R \times C$. For example, the list construction [nil, cons] : $\{*\} + R \times [R] \to [R]$ is an algebra on the set [R] of lists of rewards, while the step function step $_{x,p,r,\omega}^S : S \to \{*\} + R \times S$ is a coalgebra on the set S of states.

Note that the list construction [nil, cons] is the *initial algebra*, the discounted sum function sum_{γ} is defined as the *catamorphism* from the initial algebra to the algebra $[0, +_{\gamma}]$, while the recursive reward generation function gen_{π,p,r,ω} is defined as the *hylomorphism* from the coalgebra step_{π,p,r,ω} to the initial algebra. In the field of functional programming, such operations are also known as fold and unfold (Meijer et al., 1991; Bird & de Moor, 1997; Hutton, 1999; Yang & Wu, 2022).

Due to the recursive nature of the generation and aggregation functions, we can derive the recursive structure of their composition using the algebra fusion technique (Hinze et al., 2010), which leads to the Bellman equations for the state statistic function $\tau_{\pi}^{S} : S \to T$ in Theorem 3.2 and the state-action statistic function $\tau_{\pi}^{S \times A} : S \times A \to T$ in Theorem A.1.

B.2 Bidirectional process

In Fig. 2, we illustrate the bidirectional processes for the state statistic function and state value function. In algebra, such bidirectional processes are called *lenses* and *optics* (Riley, 2018). Such bidirectional processes (Riley, 2018) have been applied to study supervised learning (Fong & Johnson, 2019), Bayesian inference (Smithe, 2020), gradient-based learning (Cruttwell et al., 2022), and reinforcement learning (Hedges & Sakamoto, 2022).

Note that there is a slight difference between the definitions of step/generation/statistic functions in Eqs. (3), (4) and (10) and the bidirectional process in Fig. 2 (reproduced in Fig. 7). In Eq. (3), a state s is duplicated and passed separately to the transition function p_{π} and the reward function r_{π} , requiring the policy π to compute the action a twice. In contrast, in Fig. 7, the state s is passed to the policy π only once, and the action a is computed only once and then copied to the transition function p and the reward function r. These two approaches are equivalent only when the following equation holds:



For functions, copying an input and then passing the copies to two identical functions is equivalent to passing the input to the function once and then copying the output. However, for stochastic functions, these two approaches are not equivalent, which requires additional care when defining bidirectional processes for stochastic functions (see also Fritz, 2020, Definition 10.1).

Strictly speaking, the definitions in Eqs. (3), (4) and (10) correspond to a bidirectional process illustrated in Fig. 8, where different behavior and target policies can be considered. In this setting, the target policy π_{θ} , parameterized by θ , is used to compute the reward and is optimized, while the potentially unknown behavior policy π is passed to the transition function. Further, the *internal state* between the forward and backward processes — also known as the *residual* (Gavranović, 2022) — can be the state itself rather than the reward, as shown in Fig. 9. Similar considerations extend to the state-action statistic function, as illustrated in Figs. 10 to 12.

We believe that such bidirectional processes offer a clearer framework for reinforcement learning, including offline reinforcement learning, inverse reinforcement learning, and imitation learning (Hussein et al., 2017; Arora & Doshi, 2021; Hedges & Sakamoto, 2022; Murphy, 2024). Further research is needed to explore the full potential of bidirectional processes in reinforcement learning.

B.3 Non-uniqueness of update function and post-processing function

It is important to note that for a given aggregation function, the corresponding update function $\triangleright: R \times T \to T$ and post-processing function post $: T \to R$ are not necessarily unique.

Mean For example, the mean function can be computed recursively in different ways: one approach updates the sum and the length, while another updates the mean and the length. Each approach has its own advantages and disadvantages. Updating the sum allows for a straightforward implementation, but when both the sum and the length are large, numerical instability may arise. In contrast, updating the mean may require additional computation, but if the rewards are bounded, the mean remains bounded as well, which can improve numerical stability.

Variance Similarly, the variance can also be computed recursively through multiple formulations. A common method maintains the sum of squares, the mean, and the length, while a more numerically stable alternative, Welford's algorithm (Welford, 1962), updates the variance directly using incremental differences. Specifically, the update rule is given by:

$$\sigma_{t+1}^2 = \sigma_t^2 + \frac{t(r_{t+1} - \mu_t)^2 - (t+1)\sigma_t^2}{(t+1)^2},$$
(36)

where r_t , μ_t , and σ_t^2 denote the reward observed at time step t, the mean of the rewards up to time t, and the variance of the rewards up to time t, respectively. To compute the variance iteratively using this formulation, it is sufficient to maintain and update the length, the mean, and the variance at each step. This formulation improves numerical stability by preventing catastrophic cancellation (Goldberg, 1991; Muller et al., 2018), which occurs when subtracting two large and nearly identical values, leading to significant precision loss in floating-point arithmetic.

	Premetric	Strict premetric	Metric
Indiscernibility of identities $(a_1 = a_2) \rightarrow (d_A(a_1, a_2) = 0)$	\checkmark	\checkmark	\checkmark
Identity of indiscernibles $(d_A(a_1, a_2) = 0) \rightarrow (a_1 = a_2)$		\checkmark	\checkmark
Symmetry $d_A(a_1, a_2) = d_A(a_2, a_1)$			\checkmark
Triangle inequality $d_A(a_1, a_3) \le d_A(a_1, a_2) + d_A(a_2, a_3)$			\checkmark

Table 3: Properties of metrics

C Metrics and Bellman operators

In this section, we discuss the *metrics* on the statistics T and rewards R and the *Bellman operators* for the state/state-action statistic functions.

C.1 Preliminaries

Recall the definitions of metrics, as summarized in Table 3:

Definition C.1 (Premetric). A *premetric* on a set A is a function $d_A : A \times A \rightarrow [0, \infty]$ such that $\forall a \in A. \ d_A(a, a) = 0.$

Definition C.2 (Strict premetric). A *strict premetric* on a set A is a function $d_A : A \times A \rightarrow [0, \infty]$ such that $\forall a_1, a_2 \in A$. $(d_A(a_1, a_2) = 0) \leftrightarrow (a_1 = a_2)$.

Given a function to a premetric space, we can define a premetric on the domain by pullback:

Lemma C.3 (Pullback premetric). Let $d_B : B \times B \to [0, \infty]$ be a premetric on a set B, and let $f : A \to B$ be a function. The pullback premetric $d_A : A \times A \to [0, \infty]$ is defined by $\forall a_1, a_2 \in A. \ d_A(a_1, a_2) := d_B(f(a_1), f(a_2)).$ (37)

If d_B is a strict premetric, then d_A is also a strict premetric if and only if the function f is injective.

C.2 Metrics on statistics and rewards

By Lemma C.3, we can define a premetric d_T on statistics T by pulling back a premetric d_R on rewards R through a post-processing function post : $T \to R$:

$$\forall \tau_1, \tau_2 \in T. \ d_T(\tau_1, \tau_2) \coloneqq d_R(\operatorname{post}(\tau_1), \operatorname{post}(\tau_2)).$$
(38)

However, when rewards R are real-valued while statistics T are multi-dimensional, the pullback premetric d_T may not be a strict premetric, as different statistics may map to the same reward value.

For example, consider the range of rewards, where the statistics $T = \mathbb{R}^2$ are the maximum and minimum of rewards. We can directly define a metric on statistics by

$$d_T \left(\begin{bmatrix} m_1 \\ n_1 \end{bmatrix}, \begin{bmatrix} m_2 \\ n_2 \end{bmatrix} \right) := \sqrt{(m_1 - m_2)^2 + (n_1 - n_2)^2}.$$
(39)

If we use the pullback premetric, we have

$$d_T\left(\begin{bmatrix}m_1\\n_1\end{bmatrix},\begin{bmatrix}m_2\\n_2\end{bmatrix}\right) \coloneqq d_R\left(\operatorname{post}\left(\begin{bmatrix}m_1\\n_1\end{bmatrix}\right),\operatorname{post}\left(\begin{bmatrix}m_2\\n_2\end{bmatrix}\right)\right) \tag{40}$$

$$= d_R(m_1 - n_1, m_2 - n_2) \tag{41}$$

$$= |(m_1 - n_1) - (m_2 - n_2)|.$$
(42)

C.3 Bellman operators

Recall the definition of the Bellman operator for a state statistic function $\tau^S: S \to T$:

Definition 3.4 (Bellman operator). Given a policy π , a transition function p, a reward function r, a terminal condition ω , and a recursive statistic aggregation function $\underset{i}{\operatorname{agg}}$ (Definition 3.1), the *Bellman operator* $\mathcal{B}_{\pi} : [S,T] \to [S,T]$ for a function $\tau : S \to T$ is defined by

$$\mathcal{B}_{\pi}\tau: S \to T := \begin{bmatrix} s \mapsto \begin{cases} \text{init} & s \in S_{\omega} \\ r_{\pi}(s) \triangleright \tau(p_{\pi}(s)) & s \notin S_{\omega} \end{bmatrix}.$$
(11)

We can define a Bellman operator for a state-action statistic function $\tau^{S \times A} : S \times A \to T$ similarly:

Definition C.4 (Bellman operator). Given a policy π , a transition function p, a reward function r, a terminal condition ω , and a recursive statistic aggregation function $\underset{\pi}{\operatorname{agg}}_{\operatorname{init},\triangleright}$ (Definition 3.1), the *Bellman operator* $\mathcal{B}_{\pi}^{S \times A} : [S \times A, T] \to [S \times A, T]$ for a function $\tau^{S \times A} : S \times A \to T$ is defined by

$$\mathcal{B}_{\pi}^{S \times A} \tau^{S \times A} : S \times A \to T := \left[s, a \mapsto \begin{cases} \text{init} & s \in S_{\omega} \\ \mathbf{r}(s, a) \triangleright \tau^{S \times A}(\mathbf{p}_{\pi}^{S \times A}(s, a)) & s \notin S_{\omega} \end{cases} \right].$$
(43)

C.4 Existence of fixed points of Bellman operators

The existence of fixed points of the Bellman operators \mathcal{B}_{π}^{S} and $\mathcal{B}_{\pi}^{S \times A}$ is established by the Bellman equations for the state statistic function $\tau_{\pi}^{S} : S \to T$ in Theorem 3.2 and the state-action statistic function $\tau_{\pi}^{S \times A} : S \times A \to T$ in Theorem A.1.

Remark 4 (Banach fixed point theorem). Note that the classical fixed point theorem for Bellman operators typically relies on the *Banach fixed point theorem*, which requires the underlying space to be a *complete metric space*. This is not an issue in the standard discounted sum setting, as the space \mathbb{R} of real numbers has a complete metric structure. However, in our setting, the space T of statistics may lack such a complete metric structure, posing potential challenges for establishing fixed point guarantees. That said, the triangle inequality of the metric and the completeness of the space are only necessary for ensuring the *existence* of fixed points: the triangle inequality guarantees that the iterative sequence is a Cauchy sequence, while completeness ensures that the sequence has a limit within the space. Since the existence of fixed points follows directly from the Bellman equations, our focus shifts to the *uniqueness* of fixed points, which only requires the space to be a premetric space.

C.5 Uniqueness of fixed points of Bellman operators

Recall that Theorem 3.6 establishes the uniqueness of fixed points of the Bellman operator \mathcal{B}^S_{π} for state statistic functions $\tau^S : S \to T$:

Theorem 3.6 (Uniqueness of fixed points of Bellman operator). Let $\tau_1, \tau_2 : S \to T$ be fixed points of the Bellman operator \mathcal{B}_{π} (Definition 3.4). If the update function \triangleright is contractive with respect to a premetric d_T on statistics T (Definition 3.5), then $d_T(\tau_1(s), \tau_2(s)) = 0$ for all states $s \in S$. If d_T is a strict premetric, then $\tau_1 = \tau_2 = \tau_{\pi}$.

Similarly, we can extend this result to the Bellman operator $\mathcal{B}_{\pi}^{S \times A}$ for state-action statistic functions $\tau^{S \times A} : S \times A \to T$:

Theorem C.5 (Uniqueness of fixed points of the Bellman operator). Let $\tau_1^{S \times A}, \tau_2^{S \times A} : S \times A \to T$ be fixed points of the Bellman operator $\mathcal{B}_{\pi}^{S \times A}$ (Definition C.4). If the update function \triangleright is contractive with respect to a premetric d_T on statistics T (Definition 3.5), then $d_T(\tau_1^{S \times A}(s, a), \tau_2^{S \times A}(s, a)) = 0$ for all states $s \in S$ and actions $a \in A$. If d_T is a strict premetric, then $\tau_1^{S \times A} = \tau_2^{S \times A} = \tau_{\pi}^{S \times A}$.

	Preorder	Partial order	Total preorder	Total order
$\begin{array}{c} \text{Reflexivity} \\ a \leq_A a \end{array}$	\checkmark	\checkmark	\checkmark	\checkmark
$\begin{array}{c} \text{Transitivity} \\ (a_1 \leq_A a_2) \land (a_2 \leq_A a_3) \to (a_1 \leq_A a_3) \end{array}$	\checkmark	\checkmark	\checkmark	\checkmark
Antisymmetry $(a_1 \leq_A a_2) \land (a_2 \leq_A a_1) \rightarrow (a_1 = a_2)$		\checkmark		\checkmark
Totality $(a_1 \leq_A a_2) \lor (a_2 \leq_A a_1)$			\checkmark	\checkmark

 Table 4: Properties of orders

D Orders and Bellman optimality operators

In this section, we discuss the *orders* on the statistics T and rewards R and the *Bellman optimality operators* for the state/state-action statistic functions.

D.1 Preliminaries

Recall the definitions of orders, as summarized in Table 4:

Definition D.1 (Preorder). A *preorder* on a set A is a relation \leq_A that is reflexive $\forall a \in A$. $a \leq_A a$ and transitive $\forall a_1, a_2, a_3 \in A$. $(a_1 \leq_A a_2) \land (a_2 \leq_A a_3) \rightarrow (a_1 \leq_A a_3)$.

Definition D.2 (Partial order). A *partial order* on a set A is a relation \leq_A that is reflexive, transitive, and antisymmetric $\forall a_1, a_2 \in A$. $(a_1 \leq_A a_2) \land (a_2 \leq_A a_1) \rightarrow (a_1 = a_2)$.

Definition D.3 (Total preorder). A *total preorder* on a set A is a relation \leq_A that is reflexive, transitive, and total $\forall a_1, a_2 \in A$. $(a_1 \leq_A a_2) \lor (a_2 \leq_A a_1)$.

Definition D.4 (Total order). A *total order* on a set A is a relation \leq_A that is reflexive, transitive, antisymmetric, and total.

Given a function to a preorder space, we can define a preorder on the domain by pullback:

Lemma D.5 (Pullback preorder). Let \leq_B be a preorder on a set B, and let $f : A \to B$ be a function. The pullback preorder \leq_A on a set A is defined by

$$\forall a_1, a_2 \in A. \ (a_1 \leq_A a_2) := (f(a_1) \leq_B f(a_2)). \tag{44}$$

If \leq_B is total, then \leq_A is also total. If \leq_B is antisymmetric, then \leq_A is also antisymmetric if and only if f is injective.

Given a preorder and a premetric, wen can consider how the premetric preserves the preorder:

Definition D.6 (Preorder-preserving premetric). A *premetric* $d_B : B \times B \to [0, \infty]$ on a set B preserves a preorder \leq_B on the set B if

$$\forall b_1, b_2, b_3 \in B. \ (b_1 \leq_B b_2 \leq_B b_3) \rightarrow (d_B(b_1, b_2) \leq d_B(b_1, b_3)) \land (d_B(b_3, b_2) \leq d_B(b_3, b_1)). \ (45)$$

Note that since a premetric is not required to be symmetric, there are in total eight possible inequalities that we can consider for the preorder preservation of a premetric, which are omitted here for brevity.

Given a preorder-preserving premetric, we can consider an inequality for the supremum of functions:

Lemma D.7 (Preorder-preserving premetric's supremum inequality). Let $d_B : B \times B \to [0, \infty]$ be a premetric that preserves a premetric \leq_B on a set B. Then, for functions $f_1, f_2 : A \to B$ whose suprema are attained in B, we have

$$d_B(\sup_{a \in A} f_1(a), \sup_{a \in A} f_2(a)) \le \sup_{a \in A} d_B(f_1(a), f_2(a)).$$
(46)

This lemma is useful for proving the contraction property of the Bellman optimality operator, as we will see later.

D.2 Orders on statistics and rewards

By Lemma D.5, we can define a preorder \leq_T on statistics T by pulling back a preorder \leq_R on rewards R through a post-processing function post : $T \rightarrow R$:

$$\tau_1, \tau_2 \in T. \ (\tau_1 \leq_T \tau_2) := (\text{post}(\tau_1) \leq_R \text{post}(\tau_2)).$$
 (47)

Since the (pre)order \leq_R on rewards R is usually the total order of real numbers, we can guarantee that the preorder \leq_T on statistics T is also total.

For example, consider the arithmetic mean of rewards, where the statistics $T = \mathbb{N} \times \mathbb{R}$ are the length and the sum of rewards. We can compare two statistics (n_1, s_1) and (n_2, s_2) by comparing the means $\frac{s_1}{n_1}$ and $\frac{s_2}{n_2}$. This is a total preorder on the statistics T.

D.3 Bellman optimality operators

We can define the Bellman optimality operators as follows:

Definition D.8 (Bellman optimality operator). Given a policy π , a transition function p, a reward function r, a terminal condition ω , a recursive statistic aggregation function $\underset{t}{\operatorname{agg}}_{\operatorname{init}, \triangleright}$ (Definition 3.1), and a preorder \leq_T on statistics T, the *Bellman optimality operator* $\mathcal{B}^S_* : [S,T] \to [S,T]$ for a function $\tau^S : S \to T$ is defined by

$$\mathcal{B}^{S}_{*}\tau^{S}: S \to T := \left[s \mapsto \begin{cases} \text{init} & s \in S_{\omega} \\ \sup_{a \in A} \left(\mathbf{r}(s, a) \triangleright \tau^{S}(\mathbf{p}(s, a)) \right) & s \notin S_{\omega} \end{cases} \right].$$
(48)

Definition D.9 (Bellman optimality operator). Given a policy π , a transition function p, a reward function r, a terminal condition ω , a recursive statistic aggregation function $\arg_{init,\triangleright}$ (Definition 3.1), and a preorder \leq_T on statistics T, the *Bellman optimality operator* $\mathcal{B}^{S \times A}_* : [S \times A, T] \to [S \times A, T]$ for a function $\tau^{S \times A} : S \times A \to T$ is defined by

$$\mathcal{B}^{S \times A}_{*}\tau^{S \times A} : S \times A \to T := \left[s, a \mapsto \begin{cases} \text{init} & s \in S_{\omega} \\ \sup_{a' \in A} \left(\mathbf{r}(s, a) \triangleright \tau^{S \times A}(\mathbf{p}(s, a), a') \right) & s \notin S_{\omega} \end{cases} \right].$$
(49)

D.4 Existence of fixed points of Bellman optimality operators

Recall that Theorem 3.8 establishes the existence of a fixed point of the Bellman optimality operator \mathcal{B}^S_* for state statistic functions $\tau^S: S \to T$:

Theorem 3.8 (Bellman optimality equation for the state statistic function). Given a preorder \leq_T on statistics *T*, the optimal state statistic function τ_* (Definition 3.7) satisfies

$$\tau_*: S \to T := \left[s \mapsto \begin{cases} \text{init} & s \in S_{\omega} \\ \sup_{a \in A} (\mathbf{r}(s, a) \triangleright \tau_*(\mathbf{p}(s, a))) & s \notin S_{\omega} \end{cases} \right].$$
(12)

We can similarly establish the existence of a fixed point of the Bellman optimality operator $\mathcal{B}_*^{S \times A}$ for state-action statistic functions $\tau^{S \times A} : S \times A \to T$:

Theorem D.10 (Bellman optimality equation for the state-action statistic function). Given a preorder \leq_T on statistics T, the optimal state-action statistic function $\tau^{S \times A}_*$ satisfies

$$\tau_*^{S \times A} : S \times A \to T := \left[s, a \mapsto \begin{cases} \text{init} & s \in S_{\omega} \\ \sup_{a' \in A} \left(\mathbf{r}(s, a) \triangleright \tau_*^{S \times A}(\mathbf{p}(s, a), a') \right) & s \notin S_{\omega} \end{cases} \right].$$
(50)

D.5 Uniqueness of fixed points of Bellman optimality operators

Similarly to Theorem 3.6, we can guarantee the uniqueness of fixed points of the Bellman optimality operators \mathcal{B}_*^S and $\mathcal{B}_*^{S \times A}$ under certain conditions:

		Definition	Existence	Uniqueness
Bellman operator	\mathcal{B}^S_π	Definition 3.4	Theorem 3.2	Theorem 3.6
Definital operator	$\mathcal{B}^{S imes A}_{\pi}$	Definition C.4	Theorem A.1	Theorem C.5
Bellman ontimality operator	\mathcal{B}^S_*	Definition D.8	Theorem 3.8	Theorem D.11
Bennan optimanty operator	$\mathcal{B}^{S imes A}_*$	Definition D.9	Theorem D.10	Theorem D.12

Table 5: Fixed points of the Bellman operators and the Bellman optimality operators

Theorem D.11 (Uniqueness of fixed points of Bellman optimality operator). Let $\tau_1^S, \tau_2^S : S \to T$ be fixed points of the Bellman optimality operator \mathcal{B}^S_* (Definition D.8). If the update function \triangleright is contractive with respect to a premetric d_T on statistics T (Definition 3.5), and the premetric d_T preserves the preorder \leq_T on statistics T (Definition D.6), then $d_T(\tau_1^S(s), \tau_2^S(s)) = 0$ for all states $s \in S$. If d_T is a strict premetric, then $\tau_1^S = \tau_2^S = \tau_*^S$.

Theorem D.12 (Uniqueness of fixed points of Bellman optimality operator). Let $\tau_1^{S \times A}, \tau_2^{S \times A}$: $S \times A \to T$ be fixed points of the Bellman optimality operator $\mathcal{B}_*^{S \times A}$ (Definition D.9). If the update function \triangleright is contractive with respect to a premetric d_T on statistics T (Definition 3.5), and the premetric d_T preserves the preorder \leq_T on statistics T (Definition D.6), then $d_T(\tau_1^{S \times A}(s, a), \tau_2^{S \times A}(s, a)) = 0$ for all states $s \in S$ and actions $a \in A$. If d_T is a strict premetric, then $\tau_1^{S \times A} = \tau_2^{S \times A} = \tau_*^{S \times A}$.

In summary, the definitions and results on the fixed points of the Bellman operators and the Bellman optimality operators are summarized in Table 5.

E Stochastic Markov decision process

In this section, we discuss the stochastic extension of the deterministic Markov decision processes introduced in Sections 2 and 3.

E.1 Composition of stochastic functions

The composition rules of stochastic functions and deterministic functions are defined as follows:

■ Composition of two stochastic functions f : A → PB and g : B → PC by marginalizing over the intermediate variable, as described by the Chapman–Kolmogorov equation (Giry, 1982):

$$(g \circ f)(c|a) := \int_{B} g(c|b)f(b|a) \, db.$$

$$A \xrightarrow{g \circ f} B \xrightarrow{g \circ f} C$$

$$f \xrightarrow{g \circ f} g \xrightarrow{g} \mathbb{P} C$$

$$\stackrel{f}{\mathbb{P}} B \xrightarrow{\mu_{C}} P \xrightarrow{\mu_{C}} P \xrightarrow{\mu_{C}} P \mathbb{P} C$$

$$(52)$$

• Composition of a stochastic function $f : A \to \mathbb{P}B$ with a deterministic function $g : B \to C$:

$$(g \circ f)(c|a) := g_* f(b|a) = \int_B \delta_g(b) f(b|a) \, \mathrm{d}b. \tag{53}$$

$$A \xrightarrow{g \circ f} \xrightarrow{g \circ f} C$$

$$f \xrightarrow{g \circ f} \delta_g \xrightarrow{\delta_g} C$$

$$\mathbb{P}B \xrightarrow{g_*} \mathbb{P}C$$

$$\overset{\uparrow}{\mathbb{P}} \delta_g \xrightarrow{\downarrow_{L_C}} 0$$

$$\mathbb{P}B \xrightarrow{g \circ f} 0$$

• Composition of a deterministic function $f : A \to B$ with a stochastic function $g : B \to \mathbb{P}C$: $(g \circ f)(c|a) := g(c|f(a)).$ (55)



	expected maximum rewards	maximum expected rewards
definition	$\mathbb{E}_{\pi}[\max(r_1, r_2, \dots, r_{\Omega})]$	$\max(\mathbb{E}_{\pi}[r_1], \mathbb{E}_{\pi}[r_2], \dots, \mathbb{E}_{\pi}[r_{\Omega}])$
statistic T	max reward distribution $\in \mathbb{P}\overline{\mathbb{R}}$	max reward expectation $\in \overline{\mathbb{R}}$
initial value	Dirac delta measure $\delta_{-\infty} \in \mathbb{P}\overline{\mathbb{R}}$	reward value $-\infty \in \overline{\mathbb{R}}$
update function	pushforward measure update	expected value update
	$\mathbb{P}\overline{\mathbb{R}} \times \mathbb{P}\overline{\mathbb{R}} \to P(\overline{\mathbb{R}} \times \overline{\mathbb{R}}) \xrightarrow{\max_*} \mathbb{P}\overline{\mathbb{R}}$	$\mathbb{P}\overline{\mathbb{R}}\times\overline{\mathbb{R}}\xrightarrow{\mathbb{E}_{\overline{\mathbb{R}}}\times\operatorname{id}_{\overline{\mathbb{R}}}}\overline{\mathbb{R}}\times\overline{\mathbb{R}}\xrightarrow{\max}\overline{\mathbb{R}}$
post-processing	expectation $\mathbb{E}_{\overline{\mathbb{R}}} : \mathbb{P}\overline{\mathbb{R}} \to \overline{\mathbb{R}}$	identity $\operatorname{id}_{\overline{\mathbb{R}}}:\overline{\mathbb{R}}\to\overline{\mathbb{R}}$

Table 6: Expected aggregated rewards vs. aggregated expected rewards: maximum as an example

E.2 Stochastic recursion

In Section 4, we introduced the stochastic state transition and statistic functions. Similarly, we can define the stochastic state-action transition $p_{\pi}^{S \times A}$ as follows: $p_{\pi}^{S \times A} : S \times A \to \mathbb{P}(S \times A) := \langle \operatorname{id}_{S}, \pi \rangle \circ p$

$$p_{\pi} \longrightarrow S \times A \to \mathbb{P}(S \times A) := \langle \mathrm{Id}_{S}, \pi \rangle \otimes p$$
$$= \left[s, a \mapsto \left(s' \sim \mathrm{p}(s'|s, a), a' \sim \int_{S} \pi(a'|s') \mathrm{p}(s'|s, a) \, \mathrm{d}s' \right) \right].$$
(57)

The stochastic state-action statistic function $\tau_{\pi}^{S \times A}$ satisfies the following recursive equation: $\tau_{\pi}^{S \times A} : S \times A \to \mathbb{P}T$

$$= \left[s, a \mapsto \tau \sim \begin{cases} \delta_{\text{init}} & s \in S_{\omega} \\ \mathbf{r}(s, a) \triangleright \tau' \mid \tau' \sim \int_{S \times A} \tau_{\pi}^{S \times A}(\tau'|s', a') \mathbf{p}_{\pi}^{S \times A}(s', a'|s, a) \, \mathrm{d}s' \, \mathrm{d}a' & s \notin S_{\omega} \end{cases} \right].$$
(58)

Further characterizations of stochastic state/state-action statistic functions, including the (pre)metrics and (pre)orders on statistics, as well as the contractivity of stochastic Bellman (optimality) operators, are left for future work.

E.3 Relationship between stochastic state and state-action statistic functions

In the stochastic setting, the state/state-action statistic functions are related by the following equations, which are analogous to Theorem A.2:

$$\tau_{\pi}^{S}(\tau|s) = \int_{A} \tau_{\pi}^{S \times A}(\tau|s, a) \pi(a|s) \,\mathrm{d}a \qquad (\text{for all states}), \tag{59}$$

$$\tau_{\pi}^{S \times A}(\tau|s,a) = \mathbf{r}(s,a) \triangleright \int_{S} \tau_{\pi}^{S}(\tau|s') \mathbf{p}(s'|s,a) \,\mathrm{d}s' \qquad \text{(for all non-terminal states).} \tag{60}$$

E.4 Expected aggregated rewards vs. aggregated expected rewards

As discussed in Section 4, the expected discounted sum of rewards equals the discounted sum of expected rewards. However, the expected aggregated rewards and the aggregated expected rewards are not equal in general. For example, the expected maximum reward is not equal to the maximum expected reward because the expectation operator does not distribute over the maximum operator, as shown in Table 6. This issue was also raised by Gottipati et al. (2020); Cui & Yu (2023); Veviurko et al. (2024). However, we argue that even though the expected aggregated rewards and the aggregated expected rewards are not equal, they are both valid and useful learning objectives for different purposes, and the choice between them depends on the specific application. If we want to optimize the expected aggregated rewards, a more straightforward approach is to estimate the distributions of the aggregated rewards, using *distributional reinforcement learning* (Morimura et al., 2010a;a; Bellemare et al., 2017; 2023). Further theoretical and empirical investigations are left for future work.

F Proofs

In this section, we present the proofs of the theorems and lemmas introduced in the main text.

In some derivations, we use <u>underwave notation</u> to highlight the specific subterm being rewritten or replaced. This syntactic marking corresponds to substituting one path in a commutative diagram with another path sharing the same source and target:

$$\underbrace{f}_{\sim} = h \circ g \quad \text{or} \quad \underbrace{f}_{\sim} (a) = h(g(a)) \quad \text{means} \quad \underbrace{g}_{f} \stackrel{D}{\longrightarrow} h \quad (61)$$

Theorem 3.2 (Bellman equation for the state statistic function). Given a recursive reward generation function gen_{π,p,r,ω} (Definition 2.1) and a recursive statistic aggregation function $\operatorname{agg}_{init,\triangleright}$ (Definition 3.1), their composition, called the state statistic function $\tau_{\pi}: S \to T$, satisfies

$$\tau_{\pi}: S \to T := \operatorname{agg}_{\operatorname{init}, \triangleright} \circ \operatorname{gen}_{\pi, p, r, \omega} = \left[s \mapsto \begin{cases} \operatorname{init} & s \in S_{\omega} \\ r_{\pi}(s) \triangleright \tau_{\pi}(p_{\pi}(s)) & s \notin S_{\omega} \end{cases} \right].$$
(10)

Proof. Similarly to the diagram in Eq. (7), the state statistic function $\tau_{\pi} : S \to T$ can be represented using the following diagram:

 $\operatorname{id}_{\{*\}} + \operatorname{id}_R \times \tau_{\pi}$

$$\begin{cases} * \} + R \times S \xrightarrow{\operatorname{id}_{\{*\}} + \operatorname{id}_R \times \operatorname{gen}_{\pi, p, r, \omega}} } \\ \operatorname{step}_{\pi, p, r, \omega} \uparrow & [\operatorname{nil}, \operatorname{cons}] \downarrow & [\operatorname{init}, \triangleright] \downarrow \\ S \xrightarrow{\operatorname{gen}_{\pi, p, r, \omega}} & > [R] \xrightarrow{\tau_{\pi}} \end{cases}$$

which can be non-rigorously interpreted as a "combination" of the following two diagrams:

$$\begin{array}{c} \underset{\{*\}}{\overset{\operatorname{id}_{\{*\}}}{\longrightarrow}} \{*\} \xrightarrow{\operatorname{id}_{\{*\}}} \{*\} \xrightarrow{\operatorname{id}_{\{*\}}} \{*\} \xrightarrow{\operatorname{id}_{\{*\}}} \{*\} \xrightarrow{\operatorname{id}_{\{*\}}} \{*\} \xrightarrow{\operatorname{id}_{\{*\}}} \{*\} \xrightarrow{\operatorname{id}_{\{*\}}} R \times S \xrightarrow{\operatorname{id}_R \times \operatorname{gen}_{\pi, p, r, \omega}} R \times [R] \xrightarrow{\operatorname{id}_R \times \operatorname{agg}_{\operatorname{init}, p}} R \times T \xrightarrow{\operatorname{e}_S \uparrow} \operatorname{nil} \xrightarrow{\operatorname{init}} \operatorname{init} \xrightarrow{\operatorname{init}} \langle r_{\pi, p_{\pi}} \rangle \uparrow \xrightarrow{\operatorname{cons}} R \times [R] \xrightarrow{\operatorname{id}_R \times \operatorname{agg}_{\operatorname{init}, p}} R \xrightarrow{\operatorname{r}_{\pi}} \xrightarrow{\operatorname{r}_{\pi}} T \xrightarrow{\operatorname{r}_{\pi}} \xrightarrow{\operatorname{r}_{\pi}} T \xrightarrow{\operatorname{r}_{\pi}} \xrightarrow{\operatorname{r}_{\pi}} \begin{array}{c} \operatorname{id}_R \times \operatorname{gen}_{\pi, p, r, \omega} R \times [R] \xrightarrow{\operatorname{id}_R \times \operatorname{agg}_{\operatorname{init}, p}} R \times T \xrightarrow{\operatorname{r}_{\pi}} \xrightarrow{r}_{\pi}} \xrightarrow{\operatorname{r}_{\pi}} \xrightarrow{\operatorname{r}_{\pi}} \xrightarrow{\operatorname{r}_{\pi}} \xrightarrow{r}_{\pi}} \xrightarrow{\operatorname{r}_{\pi}} \xrightarrow{\operatorname{r}_{\pi}} \xrightarrow{\operatorname{r}_{\pi}} \xrightarrow{r}_{\pi}} \xrightarrow{\operatorname{r}_{\pi}} \xrightarrow{\operatorname{r}_{\pi}} \xrightarrow{r}_{\pi}} \xrightarrow{r}_{\pi} \xrightarrow{r}_{\pi$$

where $e_S : S \to \{*\}$ is the unique function from states to the singleton set, and $\langle \mathbf{r}_{\pi}, \mathbf{p}_{\pi} \rangle : S \to R \times S$ is the pairing of the reward and transition functions, which constitute the step function $\operatorname{step}_{\pi, \mathbf{p}, \mathbf{r}, \omega}$.

The left diagram shows that when a state $s \in S_{\omega}$ is terminal,

 $\tau_{\pi}(s) = \underset{\sim}{\operatorname{agg}_{\operatorname{init}, \triangleright}(\operatorname{gen}_{\pi, p, r, \omega}(s))}$ $= \underset{\sim}{\operatorname{agg}_{\operatorname{init}, \triangleright}(\operatorname{nil})}$ (by terminal condition of $\operatorname{gen}_{\pi, p, r, \omega}$) (63)

$$= \text{init}. \qquad (by \text{ initial condition of } \operatorname{agg}_{\operatorname{init},\triangleright}) (64)$$

The right diagram shows that when a state $s \notin S_{\omega}$ is non-terminal, $\tau_{\pi}(s) = \underset{\sim}{\operatorname{agg}_{\operatorname{init},\triangleright}}(\underset{\sim}{\operatorname{gen}_{\pi;\mathrm{p},\mathrm{r},\omega}}(s))$ (by definition of τ_{π}) (65)

$$= \underset{\sim}{\operatorname{agg}_{\operatorname{init},\triangleright}(\operatorname{cons}(\mathbf{r}_{\pi}(s), \operatorname{gen}_{\pi, \mathbf{p}, \mathbf{r}, \omega}(\mathbf{p}_{\pi}(s))))}{\operatorname{cons}(\mathbf{r}_{\pi}(s) \triangleright \underset{\sim}{\operatorname{agg}_{\operatorname{init}, \triangleright}(\operatorname{gen}_{\pi, \mathbf{p}, \mathbf{r}, \omega}(\mathbf{p}_{\pi}(s))))}}$$
(by recursive definition of $\operatorname{agg}_{\operatorname{init}, \triangleright}$) (66)
$$= \mathbf{r}_{\pi}(s) \triangleright \underset{\sim}{\operatorname{agg}_{\operatorname{init}, \triangleright}(\operatorname{gen}_{\pi, \mathbf{p}, \mathbf{r}, \omega}(\mathbf{p}_{\pi}(s))))}$$
(by recursive definition of $\operatorname{agg}_{\operatorname{init}, \triangleright}$) (67)

$$= \mathbf{r}_{\pi}(s) \triangleright \tau_{\pi}(\mathbf{p}_{\pi}(s)).$$
 (by definition of τ_{π}) (68)

By combining Eq. (64) and Eq. (68), we obtain the desired result in Eq. (10).

We omit the proof for Theorem A.1 as the derivation is similar to that of Theorem 3.2.

Lemma C.3 (Pullback premetric). Let $d_B : B \times B \to [0, \infty]$ be a premetric on a set B, and let $f : A \to B$ be a function. The pullback premetric $d_A : A \times A \to [0, \infty]$ is defined by

$$da_1, a_2 \in A. \ d_A(a_1, a_2) := d_B(f(a_1), f(a_2)).$$
(37)

If d_B is a strict premetric, then d_A is also a strict premetric if and only if the function f is injective. *Proof.* The pullback premetric d_A is a premetric because

$$\forall a \in A. \ d_A(a, a) := d_B(f(a), f(a)) = 0.$$

$$(69)$$

If d_B is a strict premetric, we have

$$\forall a_1, a_2 \in A. \ (d_A(a_1, a_2) := d_B(f(a_1), f(a_2)) = 0) \to (f(a_1) = f(a_2)). \tag{70}$$

For the pullback premetric d_A to be a strict premetric, we require that $\forall a_1, a_2 \in A. \ (f(a_1) = f(a_2)) \rightarrow (a_1 = a_2),$ (71)

which is equivalent to the injectivity of the function f.

Lemma D.5 (Pullback preorder). Let \leq_B be a preorder on a set B, and let $f : A \to B$ be a function. The pullback preorder \leq_A on a set A is defined by

$$\forall a_1, a_2 \in A. \ (a_1 \leq_A a_2) := (f(a_1) \leq_B f(a_2)). \tag{44}$$

If \leq_B is total, then \leq_A is also total. If \leq_B is antisymmetric, then \leq_A is also antisymmetric if and only if f is injective.

Proof. The pullback preorder \leq_A is reflexive because

$$a \in A. \ (a \leq_A a) := (f(a) \leq_B f(a)). \tag{72}$$

The pullback preorder \leq_A is transitive because

$$\forall a_1, a_2, a_3 \in A. \ (a_1 \leq_A a_2) \land (a_2 \leq_A a_3) \coloneqq (f(a_1) \leq_B f(a_2)) \land (f(a_2) \leq_B f(a_3))$$
(73)

$$\to (f(a_1) \le_B f(a_3)) =: (a_1 \le_A a_3). \tag{74}$$

If \leq_B is total, then \leq_A is also total because

$$\forall a_1, a_2 \in A. \ (a_1 \leq_A a_2) \lor (a_2 \leq_A a_1) := (f(a_1) \leq_B f(a_2)) \lor (f(a_2) \leq_B f(a_1)).$$
(75) If \leq_B is antisymmetric, we have

$$\forall a_1, a_2 \in A. \ (a_1 \leq_A a_2) \land (a_2 \leq_A a_1) := (f(a_1) \leq_B f(a_2)) \land (f(a_2) \leq_B f(a_1))$$
(76)

$$\rightarrow (f(a_1) = f(a_2)). \tag{77}$$

For the pullback preorder \leq_A to be antisymmetric, we require that $\forall a_1, a_2 \in A. \ (f(a_1) = f(a_2)) \rightarrow (a_1 = a_2),$ (78)

$$a_1, a_2 \in A. \ (f(a_1) = f(a_2)) \to (a_1 = a_2), \tag{78}$$

ectivity of the function f.

which is equivalent to the injectivity of the function f.

Lemma D.7 (Preorder-preserving premetric's supremum inequality). Let $d_B : B \times B \to [0, \infty]$ be a premetric that preserves a premetric \leq_B on a set B. Then, for functions $f_1, f_2 : A \to B$ whose suprema are attained in B, we have

$$d_B(\sup_{a \in A} f_1(a), \sup_{a \in A} f_2(a)) \le \sup_{a \in A} d_B(f_1(a), f_2(a)).$$
(46)

Proof. By assumption, the functions f_1 and f_2 have suprema in B. We denote $a_1 = \arg \sup_{a \in A} f_1(a)$ and $a_2 = \arg \sup_{a \in A} f_2(a)$. Then, $f_1(a_1) = \sup_{a \in A} f_1(a)$ and $f_2(a_2) = \sup_{a \in A} f_2(a)$.

If $f_1(a_1) \leq_B f_2(a_2)$, we have $f_1(a_2) \leq_B f_1(a_1) \leq_B f_2(a_2)$. By the preorder preservation of the premetric d_B , we have

$$d_B(f_1(a_1), f_2(a_2)) \le d_B(f_1(a_2), f_2(a_2)) \le \sup_{a \in A} d_B(f_1(a), f_2(a)).$$
(79)

Similarly, if $f_2(a_2) \leq_B f_1(a_1)$, we have $f_2(a_1) \leq_B f_2(a_2) \leq_B f_1(a_1)$. By the preorder preservation of the premetric d_B , we have

$$d_B(f_1(a_1), f_2(a_2)) \le d_B(f_1(a_1), f_2(a_1)) \le \sup_{a \in A} d_B(f_1(a), f_2(a)).$$
(80)

Therefore, we have $d_B(\sup_{a \in A} f_1(a), \sup_{a \in A} f_2(a)) \leq \sup_{a \in A} d_B(f_1(a), f_2(a)).$

We use the following lemmas to prove Theorem 3.6.

Lemma F.1 (Induced premetric on a set of functions). Let $d_B : B \times B \to [0, \infty]$ be a premetric on a set B. For functions $f, f' : A \to B$, define $d_{[A,B]} : [A,B] \times [A,B] \to [0,\infty]$ as follows:

$$d_{[A,B]}(f,f') := \sup_{a \in A} d_B(f(a), f'(a)).$$
(81)

Then, $d_{[A,B]}$ is also a premetric. Moreover, if d_B is a strict premetric, $d_{[A,B]}$ is also a strict premetric. Proof. $d_{[A,B]}$ is a premetric because $d_{[A,B]}(f,f) = \sup_{a \in A} d_B(f(a), f(a)) = 0$. For two functions $f, f' : A \to B$, $d_{[A,B]}(f,f') = \sup_{a \in A} d_B(f(a), f'(a)) = 0$ implies that $d_B(f(a), f'(a)) = 0$ for all $a \in A$. If d_B is a strict premetric, then $d_B(f(a), f'(a)) = 0$ implies f(a) = f'(a) for all $a \in A$, which means that f = f', hence if d_B is a strict premetric, $d_{[A,B]}$ is also a strict premetric. \Box

Lemma F.2 (Data processing inequality). Let $d_{[A,B]}$ be the induced premetric defined in Lemma F.1. For functions $f, f' : A \to B$ and $g : A \to A$, we have

$$d_{[A,B]}(f \circ g, f' \circ g) \le d_{[A,B]}(f, f').$$
(82)

 $\begin{array}{l} \textit{Proof.} \ \ d_{[A,B]}(f \circ g, f' \circ g) \coloneqq \sup_{a \in A} d_B(f(g(a)), f'(g(a))) = \sup_{a' \in g(A)} d_B(f(a'), f'(a')) \\ \leq \sup_{a' \in A} d_B(f(a'), f'(a')) = : d_{[A,B]}(f, f'). \end{array}$

Lemma F.3 (Uniqueness of fixed points of a premetric contraction). Let a_1 and a_2 be fixed points of a function $f : A \to A$. If the function f is contractive with respect to a premetric d_A on the set A, then $d_A(a_1, a_2) = 0$. Moreover, if d_A is a strict premetric, then $a_1 = a_2$.

Proof. Because a_1 and a_2 are fixed points of f, and f is contractive with respect to d_A , there exists a constant $k \in [0, 1)$ such that

$$d_A(a_1, a_2) = d_A(f(a_1), f(a_2)) \le k \cdot d_A(a_1, a_2).$$
(83)

Given that $d_A(a_1, a_2) \ge 0$, the only possible solution is $d_A(a_1, a_2) = 0$. If d_A is a strict premetric, then $d_A(a_1, a_2) = 0$ implies $a_1 = a_2$. In other words, a premetric contraction has unique fixed points *up to premetric indiscernibility*, while a strict premetric contraction has a unique fixed point. \Box

Lemma F.4 (Contraction of Bellman operator). *If the update function* \triangleright *is contractive with respect to a premetric* d_T *on statistics* T (*Definition 3.5*), *then the Bellman operator* \mathcal{B}^S_{π} (*Definition 3.4*) *is contractive with respect to the induced premetric* $d_{[S,T]}$ *defined in Lemma F.1*.

Proof. For any functions
$$\tau_1^S, \tau_2^S : S \to T$$
, we have

$$d_{[S,T]}(\mathcal{B}^S_{\pi}\tau_1^S, \mathcal{B}^S_{\pi}\tau_2^S) = \sup_{s \in S} d_T((\mathcal{B}^S_{\pi}\tau_1^S)(s), (\mathcal{B}^S_{\pi}\tau_2^S)(s)).$$
(84)

When a state $s \in S_{\omega}$ is terminal, for any $k \in [0, 1)$, we have

$$d_T((\mathcal{B}^{S}_{\pi}\tau_1^{S})(s), (\mathcal{B}^{S}_{\pi}\tau_2^{S})(s))$$
(85)

$$= d_T(\text{init}, \text{init})$$
 (by definition of \mathcal{B}_{π}) (86)

$$= 0 \le k \cdot d_T(\tau_1^S(\mathbf{p}_\pi^S(s)), \tau_2^S(\mathbf{p}_\pi^S(s)))$$

$$(d_T \text{ is a premetric) (87)}$$

When a state $s \notin S_{\omega}$ is non-terminal, there exists a constant $k \in [0, 1)$ such that $d_{\omega}((\mathcal{B}^{S} \sigma^{S})(a), (\mathcal{B}^{S} \sigma^{S})(a))$

$$\leq k \cdot d_T(\tau_1^S(\mathbf{p}_\pi^S(s)), \tau_2^S(\mathbf{p}_\pi^S(s)))$$
 (by contractivity of \triangleright) (90)
Then, we have

$$d_{[S,T]}(\mathcal{B}^{\mathcal{S}}_{\pi}\tau_1^S, \mathcal{B}^{\mathcal{S}}_{\pi}\tau_2^S) \tag{91}$$

$$\leq k \cdot \sup_{s \in S} d_T(\tau_1^S(\mathbf{p}_{\pi}^S(s)), \tau_2^S(\mathbf{p}_{\pi}^S(s)))$$
 (by monotonicity and homogeneity of sup) (92)
$$= k \cdot d_{[S,T]}(\tau_1^S \circ \mathbf{p}_{\pi}^S, \tau_2^S \circ \mathbf{p}_{\pi}^S)$$
 (by definition of $d_{[S,T]}$) (93)

$$\leq k \cdot d_{[S,T]}(\tau_1^S, \tau_2^S) \tag{Lemma F.2} (94)$$

Therefore, the Bellman operator \mathcal{B}_{π}^{S} is contractive with respect to the premetric $d_{[S,T]}$.

(00)

Theorem 3.6 (Uniqueness of fixed points of Bellman operator). Let $\tau_1, \tau_2 : S \to T$ be fixed points of the Bellman operator \mathcal{B}_{π} (Definition 3.4). If the update function \triangleright is contractive with respect to a premetric d_T on statistics T (Definition 3.5), then $d_T(\tau_1(s), \tau_2(s)) = 0$ for all states $s \in S$. If d_T is a strict premetric, then $\tau_1 = \tau_2 = \tau_{\pi}$.

Proof. Let $d_{[S,T]}$ be the induced premetric defined in Lemma F.1. By Lemmas F.3 and F.4, we have $d_{[S,T]}(\tau_1,\tau_2) = \sup_{s \in S} d_T(\tau_1(s),\tau_2(s)) = 0,$ (95) which means that $d_T(\tau_1(s),\tau_2(s)) = 0$ for all states $s \in S$. When d_T is a strict premetric, we have (95)

 $\tau_1 = \tau_2$, which means that τ_{π} is the unique fixed point of the Bellman operator \mathcal{B}_{π} .

We omit the proof for Theorem C.5 as the derivation is similar to that of Theorem 3.6.

Theorem 3.8 (Bellman optimality equation for the state statistic function). Given a preorder \leq_T on statistics T, the optimal state statistic function τ_* (Definition 3.7) satisfies

$$\tau_*: S \to T := \left[s \mapsto \begin{cases} \text{init} & s \in S_{\omega} \\ \sup_{a \in A} (\mathbf{r}(s, a) \triangleright \tau_*(\mathbf{p}(s, a))) & s \notin S_{\omega} \end{cases} \right].$$
(12)

Proof. When a state $s \in S_{\omega}$ is terminal, we have $\tau_*(s) = \text{init.}$ When a state $s \notin S_{\omega}$ is non-terminal, we have

$$\begin{aligned} \tau_*(s) &\coloneqq \tau_{\pi_*}(s) & \text{(by definition of } \tau_*) \ (96) \\ &= r_{\pi_*}(s) \triangleright \tau_*(p_{\pi_*}(s)) & \text{(by recursive definition of } \tau_{\pi_*}) \ (97) \\ &= r(s, \pi_*(s)) \triangleright \tau_*(p(s, \pi_*(s))) & \text{(by definitions of } r_{\pi_*} \ \text{and } p_{\pi_*}) \ (98) \\ &= \sup_{a \in A} (r(s, a) \triangleright \tau_*(p(s, a))). & \text{(pointwise maximization)} \ (99) \end{aligned}$$

Theorem D.10 (Bellman optimality equation for the state-action statistic function). Given a preorder \leq_T on statistics T, the optimal state-action statistic function $\tau^{S \times A}_*$ satisfies

$$\tau^{S \times A}_* : S \times A \to T := \left[s, a \mapsto \begin{cases} \text{init} & s \in S_{\omega} \\ \sup_{a' \in A} \left(\mathbf{r}(s, a) \triangleright \tau^{S \times A}_*(\mathbf{p}(s, a), a') \right) & s \notin S_{\omega} \end{cases} \right].$$
(50)

Proof. When a state $s \in S_{\omega}$ is terminal, we have $\tau_*^{S \times A}(s, a) = \text{init for all actions } a \in A$. When a state $s \notin S_{\omega}$ is non-terminal, we have $\tau_{\cdot}^{S \times A}(s, a) := \tau^{S \times A}(s, a)$ $\tau^{S \times A}_{*}(s, a)$ v definition of $\tau^{S \times A}$ (100)

$$\begin{aligned} \mathbf{t} &:= \tau_{\pi_*}^{S\times A}(\mathbf{s}, a) & \text{(by definition of } \tau_*^{S\times A}) \\ &= \mathbf{r}(s, a) \triangleright \tau_*^{S\times A}(\mathbf{p}_{\pi_*}^{S\times A}(\mathbf{s}, a)) & \text{(by recursive definition of } \tau_{\pi_*}^{S\times A}) & \text{(101)} \\ &= \mathbf{r}(s, a) \triangleright \tau_*^{S\times A}(\mathbf{p}(s, a), \pi_*(\mathbf{p}(s, a)))) & \text{(by definition of } \mathbf{p}_{\pi_*}^{S\times A}) & \text{(102)} \\ &= \sup_{a' \in A} \left(\mathbf{r}(s, a) \triangleright \tau_*^{S\times A}(\mathbf{p}(s, a), a') \right). & \text{(pointwise maximization)} & \text{(103)} \end{aligned}$$

Similarly to Lemma F.4 and Theorem 3.6, we use the following lemma to prove Theorem D.11.

Lemma F.5 (Contraction of Bellman optimality operator). If the update function \triangleright is contractive with respect to a premetric d_T on statistics T (Definition 3.5), and the premetric d_T preserves the preorder \leq_T on statistics T (Definition D.6), then the Bellman optimality operator \mathcal{B}^S_* (Definition D.8) is contractive with respect to the induced premetric $d_{[S,T]}$ defined in Lemma F.1.

Proof. For any functions
$$\tau_1^S, \tau_2^S : S \to T$$
, we have

$$d_{[S,T]}(\mathcal{B}^S_*\tau_1^S, \mathcal{B}^S_*\tau_2^S) = \sup_{s \in S} d_T((\mathcal{B}^S_*\tau_1^S)(s), (\mathcal{B}^S_*\tau_2^S)(s)).$$
(104)

When a state $s \in S_{\omega}$ is terminal, for any $k \in [0, 1)$, we have

$$d_T((\mathcal{B}^S_*\tau_1^S)(s), (\mathcal{B}^S_*\tau_2^S)(s))$$
(105)

$$= d_T(\text{init}, \text{init})$$
 (by definition of \mathcal{B}^S_*) (106)

$$= 0 \le k \cdot \sup_{a \in A} d_T(\tau_1^S(\mathbf{p}(s,a)), \tau_2^S(\mathbf{p}(s,a)))$$
 (*d_T* is a premetric) (107)

When a state $s \notin S_{\omega}$ is non-terminal, there exists a constant $k \in [0, 1)$ such that $d_T((\mathcal{B}^S_*\tau_1^S)(s), (\mathcal{B}^S_*\tau_2^S)(s))$

$$= d_T(\sup_{a \in A} (\mathbf{r}(s, a) \triangleright \tau_1^S(\mathbf{p}(s, a))), \sup_{a \in A} (\mathbf{r}(s, a) \triangleright \tau_2^S(\mathbf{p}(s, a))))$$
 (by definition of \mathcal{B}_*) (109)

$$\leq \sup_{a \in A} d_T(\mathbf{r}(s, a) \triangleright \tau_1^S(\mathbf{p}(s, a)), \mathbf{r}(s, a) \triangleright \tau_2^S(\mathbf{p}(s, a)))$$
 (by monotonicity of d_T) (110)

$$\leq \sup_{a \in A} k \cdot d_T(\tau_1^S(\mathbf{p}(s,a)), \tau_2^S(\mathbf{p}(s,a)))$$
 (by contractivity of \triangleright) (111)

$$= k \cdot \sup_{a \in A} d_T(\tau_1^S(\mathbf{p}(s,a)), \tau_2^S(\mathbf{p}(s,a)))$$

Then, we have

$$d_{[S,T]}(\mathcal{B}^S_*\tau_1, \mathcal{B}^S_*\tau_2) \tag{113}$$

$$\leq k \cdot \sup_{s \in S} \sup_{a \in A} d_T(\tau_1^S(\mathbf{p}(s, a)), \tau_2^S(\mathbf{p}(s, a)))$$
 (by monotonicity and homogeneity of sup) (114)
$$= k \cdot \sup_{a \in A} \sup_{s \in S} d_T(\tau_1^S(\mathbf{p}(s, a)), \tau_2^S(\mathbf{p}(s, a)))$$
 (by commutativity of sup) (115)

$$= k \cdot \sup_{a \in A} d_{[S,T]}(\tau_1^S \circ p(-,a), \tau_2^S \circ p(-,a))$$
 (by definition of $d_{[S,T]}$) (116)
$$\leq k \cdot d_{[S,T]}(\tau_1^S, \tau_2^S)$$
 (Lemma F.2) (117)

Therefore, the Bellman optimality operator \mathcal{B}^{S}_{*} is contractive with respect to the premetric $d_{[S,T]}$.

Theorem D.11 (Uniqueness of fixed points of Bellman optimality operator). Let $\tau_1^S, \tau_2^S : S \to T$ be fixed points of the Bellman optimality operator \mathcal{B}^S_* (Definition D.8). If the update function \triangleright is contractive with respect to a premetric d_T on statistics T (Definition 3.5), and the premetric d_T preserves the preorder \leq_T on statistics T (Definition D.6), then $d_T(\tau_1^S(s), \tau_2^S(s)) = 0$ for all states $s \in S$. If d_T is a strict premetric, then $\tau_1^S = \tau_2^S = \tau_*^S$.

Proof. Let $d_{[S,T]}$ be the induced premetric defined in Lemma F.1. By Lemmas F.3 and F.5, we have $d_{[S,T]}(\tau_1, \tau_2) = \sup_{s \in S} d_T(\tau_1^S(s), \tau_2^S(s)) = 0,$ (118)

which means that $d_T(\tau_1^S(s), \tau_2^S(s)) = 0$ for all states $s \in S$. When d_T is a strict premetric, we have $\tau_1^S = \tau_2^S$, which means that τ_*^S is the unique fixed point of the Bellman optimality operator \mathcal{B}_*^S . \Box We omit the proof for Theorem D.12 as the derivation is similar to that of Theorem D.11.

(108)

(by homogeneity of sup) (112)



Theorem A.2 (Relationship between state and state-action statistic functions). Given a recursive reward generation function $gen_{\pi,p,r,\omega}$ (Definition 2.1) and a recursive statistic aggregation function $agg_{init,\triangleright}$ (Definition 3.1), the state statistic function $\tau_{\pi}^{S} : S \to T$ in Eq. (10) and the state-action statistic function $\tau_{\pi}^{S \times A} : S \times A \to T$ in Eq. (17) satisfy

$$\tau_{\pi}^{S} = \tau_{\pi}^{S \times A} \circ \langle \mathrm{id}_{S}, \pi \rangle : S \to T \qquad (for all states), \tag{18}$$

$$\tau_{\pi}^{S \times A} = \mathbf{r} \triangleright (\tau_{\pi}^{S} \circ \mathbf{p}) : S \times A \to T \qquad (for all non-terminal states). \tag{19}$$

Proof. Notice the following relation:

We can show that when a state $s \in S_{\omega}$ is terminal,

$$\left(\operatorname{gen}_{\pi,\mathrm{p},\mathrm{r},\omega}^{S\times A}\circ\left\langle\operatorname{id}_{S},\pi\right\rangle\right)(s) = \operatorname{gen}_{\pi,\mathrm{p},\mathrm{r},\omega}^{S}(s) = [], \tag{120}$$

and when a state $s \notin S_\omega$ is non-terminal,

$$\operatorname{gen}_{\pi,\mathrm{p},\mathrm{r},\omega}^{S\times A} \circ \langle \operatorname{id}_S, \pi \rangle \Big)(s) = \left(\operatorname{cons} \circ \langle \operatorname{r}, \operatorname{gen}_{\pi,\mathrm{p},\mathrm{r},\omega}^{S\times A} \circ p_{\pi}^{S\times A} \rangle \circ \langle \operatorname{id}_S, \pi \rangle \right)(s)$$
(121)

$$= \left(\cos \circ \langle \mathbf{r} \circ \langle \mathbf{id}_{S}, \pi \rangle, \operatorname{gen}_{\pi, \mathbf{p}, \mathbf{r}, \omega}^{S \times A} \circ \underline{\mathbf{p}}_{\pi}^{S \times A} \circ \langle \mathbf{id}_{S}, \pi \rangle \right) (s) \quad (122)$$

$$= \left(\operatorname{cons} \circ \langle \mathbf{r}_{\pi}, \operatorname{gen}_{\pi, \mathbf{p}, \mathbf{r}, \omega}^{S \times A} \circ \langle \operatorname{id}_{S}, \pi \rangle \circ \mathbf{p}_{\pi}^{S} \rangle \right) (s),$$
(123)

which shows that $\operatorname{gen}_{\pi,\mathrm{p,r},\omega}^{S \times A} \circ \langle \operatorname{id}_S, \pi \rangle$ satisfies the same recursive equation as $\operatorname{gen}_{\pi,\mathrm{p,r},\omega}^S$ in Eq. (4). Due to the uniqueness of the recursive coalgebra (Hinze et al., 2010, Eq. (5)), we can conclude that $\operatorname{gen}_{\pi,\mathrm{p,r},\omega}^S = \operatorname{gen}_{\pi,\mathrm{p,r},\omega}^{S \times A} \circ \langle \operatorname{id}_S, \pi \rangle : S \to [R].$ (124)

Given Eq. (124), we have

$$\tau_{\pi}^{S} := \operatorname{agg}_{\operatorname{init},\triangleright} \circ \operatorname{gen}_{\pi,\mathrm{p},\mathrm{r},\omega}^{S} = \operatorname{agg}_{\operatorname{init},\triangleright} \circ \operatorname{gen}_{\pi,\mathrm{p},\mathrm{r},\omega}^{S \times A} \circ \langle \operatorname{id}_{S}, \pi \rangle = \tau_{\pi}^{S \times A} \circ \langle \operatorname{id}_{S}, \pi \rangle : S \to T.$$
(125)

Next, for a non-terminal state $s \notin S_{\omega}$ and an action $a \in A$, we have

$$\tau_{\pi}^{S \times A}(s, a) = \left(\mathbf{r} \triangleright \left(\tau_{\pi}^{S \times A} \circ \mathbf{p}_{\pi}^{S \times A} \right) \right)(s, a) \tag{126}$$

$$= \left(\mathbf{r} \triangleright \left(\tau_{\pi}^{S} \circ \mathbf{p} \right) \right) (s, a).$$
(128)

However, for a terminal state $s \in S_{\omega}$ and an action $a \in A$, the equation $\tau_{\pi}^{S \times A} = r \triangleright (\tau_{\pi}^{S} \circ p)$ may not always hold and could require additional conditions on the transition function p, the reward function r, the initial value init, and the update function \triangleright .

Intuitively, Eqs. (18) and (19) arise from the decomposition of the bidirectional process, as illustrated in Fig. 13. \Box

Remark 5. In fact, we can derive Eq. (124) directly from the relation between the state step function $\operatorname{step}_{\pi,\mathrm{p,r},\omega}^S$ and the state-action step function $\operatorname{step}_{\pi,\mathrm{p,r},\omega}^{S\times A}$.

When a state $s \in S_{\omega}$ is terminal,

$$\left(\operatorname{step}_{\pi,\mathrm{p,r},\omega}^{S \times A} \circ \langle \operatorname{id}_{S}, \pi \rangle\right)(s) = \left(\operatorname{id}_{\{*\}} \circ \operatorname{step}_{\pi,\mathrm{p,r},\omega}^{S}\right)(s) = *, \tag{129}$$

and when a state $s \notin S_\omega$ is non-terminal,

$$\left(\operatorname{step}_{\pi,\mathrm{p},\mathrm{r},\omega}^{S\times A}\circ\langle\operatorname{id}_{S},\pi\rangle\right)(s) = \left(\underbrace{\langle\mathrm{r},\mathrm{p}_{\pi}^{S\times A}\rangle\circ\langle\operatorname{id}_{S},\pi\rangle}_{(130)}\right)(s)$$

$$= \left(\langle \underline{\mathbf{r} \circ \langle \mathrm{id}_S, \pi \rangle}, \underline{\mathbf{p}}_{\pi}^{S \times A} \circ \langle \mathrm{id}_S, \pi \rangle \rangle \right) (s)$$
(131)

$$= \left(\underbrace{\langle \mathbf{r}_{\pi}, \langle \mathrm{id}_{S}, \pi \rangle \circ \mathbf{p}_{\pi}^{S} \rangle}_{\langle \mathbf{X} \rangle \langle \mathbf{X} \rangle$$

$$= \left((\mathrm{id}_R \times \langle \mathrm{id}_S, \pi \rangle) \circ \langle \mathrm{r}_\pi, \mathrm{p}_\pi^S \rangle \right) (s)$$
(133)

$$= \left((\mathrm{id}_R \times \langle \mathrm{id}_S, \pi \rangle) \circ \mathrm{step}_{\pi, \mathrm{p}, \mathrm{r}, \omega}^S \right) (s). \tag{134}$$

We can conclude that $\operatorname{step}_{\pi,\mathrm{p},\mathrm{r},\omega}^{S\times A} \circ \langle \operatorname{id}_{S}, \pi \rangle = \left(\operatorname{id}_{\{*\}} + \operatorname{id}_{R} \times \langle \operatorname{id}_{S}, \pi \rangle \right) \circ \operatorname{step}_{\pi,\mathrm{p},\mathrm{r},\omega}^{S} : S \to \{*\} + R \times (S \times A), \quad (135)$ which means that $\langle \operatorname{id}_S, \pi \rangle$ is a *coalgebra homomorphism* from the state step function $\operatorname{step}_{\pi, \mathrm{p}, \mathrm{r}, \omega}^S$ to the state-action step function $\operatorname{step}_{\pi, \mathrm{p}, \mathrm{r}, \omega}^{S \times A}$. Then, by the *coalgebra fusion law* (Hinze et al., 2010, Eq. (7)) we are set the result in Eq. (7). Eq. (7)), we can get the result in Eq. (124).

G Learning algorithms with recursive reward aggregation

In this section, we list the RL algorithms with recursive reward aggregation used in our experiments. The colored lines indicate modifications compared to the standard discounted sum version.

G.1 Q-learning

```
Algorithm 1 Q-learning (Watkins & Dayan, 1992) with recursive reward aggregation
```

Input: transition function $p: S \times A \to S$, reward function $r: S \times A \to R$, terminal condition ω , recursive reward aggregation function post $\circ \operatorname{agg_{init,p}} : [R] \to R$ Parameters: learning rate $\alpha \in (0, 1]$, exploration parameter $\epsilon \in (0, 1)$ Initialize state-action statistic function $\tau: S \times A \to T$ with initial value init $\in T$ for each episode do Initialize state swhile s is not terminal do Compute state-action value function $q(s, a) = \operatorname{post}(\tau(s, a))$ for state s and all actions aSelect action a using ϵ -greedy policy based on value function q(s, a)Execute action a, observe next state s' and reward rUpdate state-action statistic function τ : $\tau(s, a) \leftarrow \tau(s, a) + \alpha \left(\max_{a' \in A} (r \triangleright \tau(s', a')) - \tau(s, a) \right),$ where $\max_{a' \in A} (r \triangleright \tau(s', a')) = r \triangleright \tau(s', a^*)$ and $a^* = \operatorname*{arg\,max\,post}(r \triangleright \tau(s', a'))$ Update state $s \leftarrow s'$

Output: estimated optimal statistic function τ and optimal policy $\pi(s) = \arg \max_{a \in A} q(s, a)$, where $q(s, a) = post(\tau(s, a))$

G.2 PPO

Algorithm 2 PPO (Schulman et al., 2017) with recursive reward aggregation **Input:** transition function $p: S \times A \to S$, reward function $r: S \times A \to R$, terminal condition ω , recursive reward aggregation function $\text{post} \circ \text{agg}_{\text{init},\triangleright} : [R] \to R$ **Parameters:** bias-variance trade-off parameter $\lambda \in [0, 1]$, critic loss coefficient c_1 , entropy regularization coefficient c_2 **Initialize** parameterized policy function (actor) $\pi_{\theta}: S \to A$ Initialize parameterized state statistic function (critic) $\tau_{\phi}: S \to T$ for each episode do Initialize state s Collect trajectories of states and rewards following policy π_{θ} till the end of the horizon Ω **Compute** statistics $\hat{\tau}_t^{(i)} = r_t \triangleright r_{t+1} \triangleright \cdots \triangleright r_{t+i-1} \triangleright \tau_{\phi}(s_{t+i})$ for $i = 1, \dots, \Omega - t$ **Compute** state value function $v_{\phi}(s_t) = \text{post}(\tau_{\phi}(s_t))$ **Compute** advantage estimates $\hat{\alpha}_t^{(i)} = \text{post}(\hat{\tau}_t^{(i)}) - v_{\phi}(s_t)$ for $i = 1, \dots, \Omega - t$ Use one of the following as advantage $\hat{\alpha}_t$: $\hat{\alpha}_t^{(1)} = \operatorname{post}(r_t \triangleright \tau_{\phi}(s_{t+1})) - \mathbf{v}_{\phi}(s_t)$ $\hat{\alpha}_t^{(\Omega-t)} = \operatorname{post}(r_t \triangleright r_{t+1} \triangleright \cdots \triangleright \tau_{\phi}(s_{\Omega})) - \mathbf{v}_{\phi}(s_t)$ • generalized advantage estimates (GAE) (Schulman et al., 2016) $(1 - \lambda) \sum_{k=1}^{M-t} \lambda^{i-1} \hat{\alpha}_{t}^{(i)}$ **Compute** critic loss: $L_c(\phi) = \sum_{t=1}^{\Omega} \left(\mathbf{v}_{\phi}(s_t) - \operatorname{post}(\hat{\tau}_t^{(\Omega-t)}) \right)^2$ **Compute** actor loss $L_a(\theta)$ with clipping or penalty using advantage $\hat{\alpha}_t$ (Schulman et al., 2017) **Compute** entropy regularization $H(\theta)$ **Optimize** $L_a(\theta) - c_1 L_c(\phi) + c_2 H(\theta)$

Output: estimated optimal statistic function τ_{ϕ} and optimal policy π_{θ}

G.3 TD3

Algorithm 3 TD3 (Fujimoto et al., 2018) with recursive reward aggregation **Input:** transition function $p: S \times A \to S$, reward function $r: S \times A \to R$, terminal condition ω , recursive reward aggregation function post $\circ \operatorname{agg}_{\operatorname{init} \triangleright} : [R] \to R$ **Parameters:** action variance σ^2 , soft target update rate $\lambda \in (0, 1)$ **Initialize** parameterized policy function (actor) $\pi_{\theta}: S \to A$ **Initialize** two parameterized state-action statistic functions (critics) $\tau_{\phi_1}, \tau_{\phi_2}: S \times A \to T$ **Initialize** targets $\theta' \leftarrow \theta$, $\phi'_1 \leftarrow \phi_1$, $\phi'_2 \leftarrow \phi_2$, and replay buffer \mathcal{D} for each episode do Initialize state s while s is not terminal do Select action $a \sim \mathcal{N}(\pi_{\theta}(s), \sigma^2)$ (optionally with clipping) Execute action a, observe next state s' and reward rStore transition tuple (s, a, r, s') in buffer \mathcal{D} **Compute** state-action value functions $q_{\phi_i}(s, a) = \text{post}(\tau_{\phi_i}(s, a))$ for i = 1, 2if update critics then Sample a batch of transitions $B = \{(s, a, r, s')\}$ from buffer \mathcal{D} Select target action $\tilde{a} \sim \mathcal{N}(\pi_{\theta'}(s'), \sigma^2)$ (optionally with clipping) **Compute** target statistic τ_{target} : $\tau_{\text{target}} = \begin{cases} \text{init} & s \in S_{\omega} \\ \min_{i=1,2} r \triangleright \tau_{\phi'_{i}}(s', \tilde{a}) & s \notin S_{\omega} \end{cases}$ where $\min_{i=1,2} r \triangleright \tau_{\phi'_{i}}(s', \tilde{a}) = \begin{cases} r \triangleright \tau_{\phi'_{1}}(s', \tilde{a}) & \text{post}(r \triangleright \tau_{\phi'_{1}}(s', \tilde{a})) \\ r \triangleright \tau_{\phi'_{2}}(s', \tilde{a}) & \text{otherwise} \end{cases}$ Update critics $\tau_{\phi_{i}}$ by gradient descent: $\nabla_{\phi_i} \frac{1}{|B|} \sum_{(s,a,r,s') \in B} \left(\mathbf{q}_{\phi_i}(s,a) - \mathrm{post}(\tau_{\mathsf{target}}) \right)^2 \text{ for } i = 1, 2$ if update actor then Update actor by gradient ascent: $\nabla_{\theta} \frac{1}{|B|} \sum_{(s,a,r,s') \in B} \mathbf{q}_{\phi_1}(s, \pi_{\theta}(s))$ Update targets: $\phi_i' \leftarrow \lambda \phi_i + (1 - \lambda) \phi_i' \text{ for } i = 1, 2$ $\theta' \leftarrow \lambda \theta + (1 - \lambda) \theta'$ **Output:** estimated optimal statistic functions τ_{ϕ_1} and τ_{ϕ_2} and optimal policy π_{θ}



(a) min (b) $0.7 \min + 0.3 \max$ (c) $0.4 \min + 0.6 \max$ (d) $0.1 \min + 0.9 \max$ (e) max

Figure 14: $\max - \lambda \operatorname{range} = \lambda \min + (1 - \lambda) \max$

H Experiments

In this section, we provide detailed descriptions of the environments used in our experiments and the specific configurations and hyperparameters employed for each task. We also present additional results for the grid-world and continuous control environments.

H.1 Grid-world environment

Implementation We implemented the environment and the Q-learning (Watkins & Dayan, 1992) algorithm using NumPy (Harris et al., 2020).

Hyperparameters We used a fixed exploration parameter of 0.3. We trained agents for total training time steps of $10\,000$. We repeated each experiment with three different random seeds and observed that all runs consistently converged to the same solution. We therefore present the result from one representative run.

Additional results Similarly to Fig. 4, Fig. 14 shows the policy preferences for range-regularized max, which is an interpolation between min and max.

H.2 Wind-world environment

Implementation We implemented the environment and the PPO (Schulman et al., 2017) algorithm using JAX (Bradbury et al., 2018) and gymnax (Lange, 2022).

Note that because PPO uses a stochastic policy, our algorithm effectively optimizes the *aggregated expected rewards*, which is different from the *expected aggregated rewards*. However, we argue that the aggregated expected rewards is still a meaningful objective. The extension to expected aggregated rewards using distributed RL is left for future work. See also Section 4 and Appendix E for details.

Hyperparameters The PPO clipping parameter was set to 0.2. We used a critic loss coefficient of 0.5 and an entropy regularization coefficient of 0.01. Agents were trained for a total of 500 000 time steps using 64 parallel environments, executed in batch via JAX to enable efficient data collection.

H.3 Continuous control environments

The **Lunar Lander Continuous** environment, part of the Box2D physics simulation suite (Brockman et al., 2016), involves controlling a lunar lander to safely land on a designated landing pad. The agent has continuous thrust control over the main engine and two side thrusters, which it must use efficiently to achieve a stable landing while minimizing fuel consumption. The reward function is designed to encourage precise and efficient landings. The agent receives positive rewards for (i) moving closer to the landing pad, (ii) achieving a soft landing, and (iii) staying upright. Conversely, penalties are applied for (i) excessive fuel usage, (ii) high-impact landings, and (iii) drifting too far from the target. The episode terminates if the lander successfully lands within the designated zone, crashes, or drifts out of bounds. If none of these conditions occur, the episode continues until reaching the time limit.

We used the **Hopper** environment (Erez et al., 2012) simulated using MuJoCo (Todorov et al., 2012), where a 2D one-legged robot must learn to balance and move forward efficiently. The agent controls three joints (thigh, knee, and foot) to generate locomotion while maintaining stability. The reward function in Hopper consists of three key components: (i) *healthy reward*, which incentivizes the agent to remain upright; (ii) *forward reward*, which encourages the agent to move forward; and (iii) *control cost*, which penalizes excessive energy use. Then, the total reward function is given by

reward = healthy reward + forward reward - control cost. (136)

The Hopper environment terminates when the agent is deemed unhealthy or reaches the predefined episode length limit. The agent is considered unhealthy if its state variables exceed the allowed range, its height falls below a certain threshold, or its torso angle deviates beyond a specified limit, indicating a loss of stability. If none of these conditions occur, the episode continues until the maximum duration is reached.

We used the **Ant** environment (Schulman et al., 2016) simulated using MuJoCo (Todorov et al., 2012), where a four-legged quadrupedal robot must learn to efficiently balance and move forward. The agent controls eight joints (two per leg) to generate stable locomotion while adapting to dynamic interactions with the environment. The reward function in the Ant environment is designed to encourage forward movement while maintaining stability and efficiency. It consists of four key components: (i) a *healthy reward*, which provides a fixed bonus as long as the agent remains upright; (ii) a *forward reward*, which encourages movement in the positive x-direction; (iii) a *control cost*, which penalizes excessive actions to promote energy efficiency; and (iv) a *contact cost*, which discourages large external contact forces. The total reward is calculated by summing the healthy and forward rewards while subtracting the penalties for control effort and contact forces:

reward = healthy reward + forward reward - control cost - contact cost. (137) In some versions of the environment, the contact cost may be excluded from the reward calculation. The Ant environment terminates when the agent is deemed unhealthy or when the episode reaches its maximum duration of 1000 time steps. The agent is considered unhealthy if any of its state space values become non-finite or if its torso height falls outside a predefined range, indicating a loss of stability. If neither of these conditions occur, the episode continues until it reaches the time limit.

Implementation We conducted experiments using a modified version of the TD3 (Fujimoto et al., 2018) implementation from Stable-Baselines3 (Raffin et al., 2021).

Hyperparameters Our agent performed 100 gradient updates per training episode and used a learning rate of 3×10^{-4} to ensure stable learning. Apart from these, our training setup adheres to the default hyperparameters and network architecture of Stable-Baselines3.

Computational resource Training a single agent takes approximately 1 hour on an NVIDIA RTX 2080 GPU, with a single CPU core used for environment simulation.

Additional results We provide additional results for Hopper and Ant environments. To comprehensively assess the performance, we present the mean values of various evaluation metrics across four random seeds using radar charts. Additionally, we visualize the trajectory of the agent in all environments, providing an intuitive representation of how different aggregation functions influence the learned policy. Animations for all three environments (Lunar Lander Continuous, Ant, and Hopper) are also available at https://github.com/Tang-Yuting/recursive-rew ard-aggregation, offering an intuitive understanding of policy behavior.

For the **Hopper** environment, we observe distinct behavioral patterns and performance outcomes under different reward aggregation strategies. The $sum_{0.99}$ aggregation, serving as the baseline method, demonstrates strong overall performance across multiple metrics, as reflected in both the radar chart and motion sequences. In contrast, the $max_{0.99}$ aggregation focuses solely on optimizing max reward, leading to strong performance in this specific metric but suboptimal outcomes in others. The corresponding images show the agent taking overly aggressive actions to maximize max reward, which causes it to lose balance quickly as the torso angle exceeds the allowed range. The min



Figure 15: **Hopper**: Comparison of five reward aggregation methods. (Left) Radar plot showing performance across six evaluation metrics, averaged over four random seeds. (Right) Sample trajectories illustrating the qualitative behaviors induced by each aggregation method.



Figure 16: **Ant**: Comparison of five reward aggregation methods. (Left) Radar plot showing performance across six evaluation metrics, averaged over four random seeds. (Right) Sample trajectories illustrating the qualitative behaviors induced by each aggregation method.

aggregation encourages the agent to maximize the minimum reward, which leads to a conservative strategy where the agent remains completely still to avoid negative rewards. The sum_{0.99} + max_{0.99} aggregation encourages the agent to optimize both the total reward and the maximum reward within an episode, leading to more aggressive movements and higher overall rewards. While the sum_{0.99} - var aggregation prioritizes stability by minimizing the difference between the maximum and minimum rewards, resulting in more controlled and consistent behavior at the cost of slightly lower rewards. These results highlight how different reward aggregation strategies shape the behavior of the agent and its learning outcomes.

For the Ant environment, different aggregation strategies lead to varied agent behaviors and trade-offs between stability, performance, and exploration. The sum_{0.99} aggregation, serving as the baseline, achieves balanced performance across multiple metrics, effectively promoting stable and efficient locomotion. In contrast, the $\max_{0.99}$ aggregation prioritizes obtaining the highest possible reward at an individual time step, leading to highly aggressive movements. As a result, the agent exhibits excessive speed, which ultimately causes instability and results in the agent losing control and rolling over. The min aggregation prioritizes minimizing the risk of low rewards, leading to an overly conservative strategy. Instead of efficient locomotion, the agent adopts passive or static behavior, often staying close to the ground to avoid unfavorable rewards. This lack of exploration and controlled movement results in instability, ultimately causing the agent to collapse and terminate early due to height constraints. Moreover, the $sum_{0.99} + max_{0.99}$ aggregation encourages aggressive behavior by jointly optimizing cumulative and peak rewards. The agent exhibits rapid, unstable locomotion, often pushing for immediate gains. While this reduces stability, reward-related metrics remain high, indicating strong overall performance at the cost of greater energy use and inconsistency. Finally, the sum_{0.99} – var aggregation prioritizes stability by penalizing reward fluctuations, leading to more controlled and steady locomotion. The agent avoids aggressive actions and achieves longer episode durations. However, while reducing variance enhances stability, it also limits the ability of agent to explore high-reward strategies, leading to robust but suboptimal overall performance.

H.4 Portfolio environment

In our experiment, we trained agents using five different random seeds over a rolling 5-year window, with a total of 10 training periods. Specifically, for each training period, training begins on January 1 of a given year and continues for five years, ending on December 31 of the fifth year. Each training period starts one year after the previous one, resulting in overlapping but not identical training datasets. Following the training phase, we evaluate the performance of agents in the subsequent year, immediately following the training period. Finally, we assess their generalization performance in the test phase, which takes place in the year after the evaluation period. This design allows us to systematically analyze the agents' performance across different temporal contexts while leveraging historical data in a structured and overlapping manner.

Implementation We conducted experiments using a modified version of the PPO (Schulman et al., 2017) implementation from Stable-Baselines3 (Raffin et al., 2021).

Hyperparameters We trained each agent for a total of 7 500 000 time steps. Compared to the default settings of PPO in Stable-Baseline3, we made several modifications to better suit our environment. The learning rate followed a linear decay schedule, starting from 3×10^{-4} and gradually decreasing to 1×10^{-5} over the course of training. We set the discount factor to 0.9 and the GAE lambda to 0.9 to reduce reliance on long-term returns, and used a slightly wider clipping range 0.25 to allow for greater policy updates. These adjustments were empirically tuned for improved stability and performance in our setting.

Computational resource Training a single agent takes approximately 1.5 hours on an NVIDIA RTX 2080 GPU, with the environment running in parallel on 10 CPU cores to accelerate data collection.



Figure 17: A year-by-year comparison of Sharpe ratios obtained by different methods during the test phase across a rolling backtesting window from 2012 to 2021. Each data point represents the mean performance across five different random seeds, with the shaded regions indicating one standard deviation to reflect variability. The horizontal dashed lines represent the mean Sharpe ratio across all years for each method, providing a summary view of their long-term performance.

Results: Sharpe ratio comparison over time Figure 17 presents a comparison of the Sharpe ratios achieved by three methods during the test phase over the ten-year backtesting period. Each backtesting period refers to a historical test year following the training and validation phases, during which the strategy is evaluated on previously unseen data to assess its out-of-sample performance (Bailey et al., 2015). The red curve represents our method, which not only achieves the highest mean Sharpe ratio with relatively low variability, but also consistently delivers the best or highly competitive performance in 8 out of the 10 backtesting periods. This consistency across random seeds and temporal splits underscores the practical generalization ability of our proposed method and its suitability for real-world financial decision-making.

I Discussion

Reward function design vs. aggregation strategies Changing the reward function and adjusting how rewards are aggregated are two complementary approaches to shaping agent behavior. Rather than asserting the superiority of one approach over the other, we examine the trade-offs and situational advantages associated with each.

Reward function modification directly encodes task objectives into the per-step feedback signal received by the agent. This approach is expressive and flexible, allowing designers to incorporate domain-specific preferences (Ng et al., 1999; Hadfield-Menell et al., 2017), intermediate goals (Andrychowicz et al., 2017), or constraints (Achiam et al., 2017). However, designing an effective reward function often requires careful tuning, may introduce unintended incentives, and can suffer from reward misspecification, especially in environments with sparse or delayed feedback (Ng et al., 1999; Ziebart et al., 2008; Hadfield-Menell et al., 2017).

In contrast to modifying the reward function itself, reward aggregation modification keeps the underlying reward signal fixed and instead alters how rewards are aggregated over time to define the training objective. This offers a structured way to influence long-term behavior without redefining the reward signal at each time step. For instance, using the max aggregation emphasizes the highest reward achieved in a trajectory, thereby promoting risk-seeking behavior, while the min aggregation emphasizes avoiding the worst-case outcomes, encouraging risk-averse strategies. This approach is effective when the reward signal provides informative feedback, but the desired policy depends on how that feedback is interpreted over time. However, limited or ambiguous reward signals may restrict the ability of any aggregation function to align with the intended behavioral goals.

In practice, modifying the reward and adjusting the aggregation function are not mutually exclusive and can be combined effectively. The reward function provides the essential feedback for learning, while the aggregation method influences how this feedback is evaluated over time. The choice of whether to modify one, the other, or both should be guided by the nature of the task, the clarity and expressiveness of the reward, and the behavioral patterns desired in the learned policy.

Limitations of sum-based objectives While standard RL typically defines the training objective as the sum of per-step rewards, this formulation tends to be effective under certain assumptions about the reward signal and task structure. First, it is generally better suited to tasks where the overall performance can be approximated by accumulating the reward at each time step. In such cases, the total return should reflect meaningful progress over time. Second, sum aggregation assumes that the timing of rewards is not a critical factor. While discounted sums introduce a preference for earlier rewards, they still impose a fixed temporal structure. Therefore, sum is suited to tasks where the timing of rewards is relatively neutral and consistent accumulation matters more than when specific rewards happen. Finally, sum-based objectives are more likely to be effective when the reward function offers sufficient granularity, providing reliable feedback at each step to support training.

Despite its simplicity and widespread adoption, summing per-step rewards may be less effective in scenarios where its underlying assumptions are not fully satisfied. In particular, many tasks do not neatly align with the structure implied by a standard sum-based objective. As a result, the learned policy may be optimal under the sum semantics but misaligned with the intended behavioral goals. For example, in safety-critical environments, aggregating rewards via summation might obscure low-reward outliers, as occasional high rewards could mask dangerous behaviors. Similarly, in peak-oriented tasks where success depends on achieving exceptional performance at a specific moment, summation may diminish the significance of these peak events by averaging them with less important steps. These observations suggest that in such contexts, adjusting how rewards are aggregated over time may offer additional flexibility for aligning the learning objective with designer intent.