

On Shallow Planning Under Partial Observability

Randy Lefebvre¹, Audrey Durand^{1,2}

¹Université Laval

²Canada CIFAR AI Chair

randy.lefebvre.1@ulaval.ca, audrey.durand@ift.ulaval.ca

Abstract

Formulating a real-world problem under the Reinforcement Learning framework involves non-trivial design choices, such as selecting a discount factor for the learning objective (discounted cumulative rewards), which articulates the planning horizon of the agent. This work investigates the impact of the discount factor on the bias-variance trade-off given structural parameters of the underlying Markov Decision Process. Our results support the idea that a shorter planning horizon might be beneficial, especially under partial observability.

Introduction

Reinforcement Learning (RL) has had tremendous success on Atari games (Mnih et al. 2013), yet applications of RL in the real-world remain limited (Dulac-Arnold et al. 2021). This complexity is due to many challenges such as sample efficiency of RL methods (Yu 2018), risk/safety issues (Gu et al. 2023) and partial observability (Sondik 1978; Francois-Lavet et al. 2019; Kaelbling, Littman, and Cassandra 1998). Formulating a real-world problem under the RL framework also involves several non-trivial decisions such as selecting a state/action space (especially when these are continuous), formulating a reward function, and formulating the learning objective (Hare 2019; Devidze et al. 2021). The learning objective usually corresponds to the discounted cumulative rewards, which depends on a discount factor articulating the considered planning horizon when attributing values to states and actions (Sutton and Barto 2018). This objective is useful since it can reduce the search space intuitively by giving less credit to future rewards and actions. In toy and simulated environments, early practitioners tend to use large discount factor values often found in the RL literature on Atari (Kaiser et al. 2024; Mnih et al. 2013). This equates to considering very long planning horizons. On the other hand, real-world applications tend to formulate sequential decision-making problems under the contextual bandit setting (i.e. with a myopic agent w.r.t. the planning horizon) in response to low data regimes (Bastani and Bayati 2020; Ding, Li, and Liu 2019; Durand et al. 2018).

The impact of reducing the planning horizon has been studied previously, and bounds on the resulting bias-variance trade-off on the state value functions have been

proposed (Amit, Meir, and Ciosek 2020; Jiang et al. 2015). Unfortunately, these results provide loose bounds that do not consider the structure of the underlying Markov Decision Process (MDP) and thus fail to capture its impact on the *optimal planning horizon*. The optimal planning horizon can be defined using the discount factor γ which minimizes the planning loss (see Eq. 2), i.e. the one which can extract the best policy possible considering the limited amount of data. Distinct results involving the structure of MDPs (Jiang, Singh, and Tewari 2016; Gheshlaghi Azar, Munos, and Kappen 2013; Wu, He, and Gu 2023; He, Zhou, and Gu 2021) have been achieved separately, but these insights have never been brought together.

Contributions In this work, we introduce new results on the bias-variance trade-off that explicitly depend on high-level structural parameters of the underlying MDP. More importantly, our results touch on Partially Observable MDPs (POMDPs), providing the first insights supporting the advantage of considering short horizons in the learning objective for practical applications under partial observability. We support and illustrate the theory with experiments, hoping that this can open the door to choices in learning objectives that are better adapted to real-world RL applications. Finally, we provide the open-source code¹ for all our experiments to ensure reproducibility and offer a framework that practitioners can modify to better understand the impact of partial observability on their specific applications.

Fully Observable Setting

An MDP can be described as a tuple $(\mathcal{S}, \mathcal{A}, P, R)$, where \mathcal{S} is a finite state space, \mathcal{A} is a finite action space, $P : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \mapsto [0, 1]$ is a transition function, and $R : \mathcal{S} \times \mathcal{A} \mapsto [0, R_{\max}]$ is a reward function, with R_{\max} denoting the maximal reward obtainable in the MDP. On each time step $t \in \mathbb{N}_0$, the current state $S_t \in \mathcal{S}$ is observed, an action $A_t \in \mathcal{A}$ is played, the environment transitions into next state S_{t+1} (using P) and generates an observed reward R_{t+1} (using R). Given an MDP (tuple) M , the value of state $s \in \mathcal{S}$ under policy $\pi : \mathcal{S} \mapsto \mathcal{A}$ is the expected sum of discounted rewards obtained by selecting actions according to policy π

¹<https://github.com/GRAAL-Research/shallow-planning-partial-observability>

from state s :

$$V_{M,\gamma}^\pi(s) = \mathbb{E}_\pi \left[\sum_{k=0}^{\infty} \gamma^k R_{t+k+1} | S_t = s \right],$$

where the discount factor $\gamma \in [0, 1]$ controls the planning horizon $1/(1-\gamma)$, i.e. the credit assigned to action A_t for future rewards. The goal of a learning agent is to find the optimal policy $\pi_{M,\gamma}^*$ that maximizes $V_{M,\gamma}^\pi(s)$ for all states $s \in \mathcal{S}$. We consider a setting where this policy is unique. We use $V_{M,\gamma}^\pi \in \mathbb{R}^{|\mathcal{S}|}$ to denote the vector of state values. A notation table can be found in Appendix A.

Blackwell Discount Factor Practitioners often believe using a higher discount factor will result in a better policy. While this is true with an infinite amount of data, it is rarely the case when using RL in practice. It has even been shown previously that there always exists a discount factor γ_{Bw} such that for any $\gamma \geq \gamma_{Bw}$, we have $V_{M,\gamma}^{\pi_{M,\gamma}^*} = V_{M,\gamma}^{\pi_{M,\gamma_{Bw}}^*}$ when $|\mathcal{S}| < \infty$ and $|\mathcal{A}| < \infty$ (Grand-Clément and Petrik (2023), Thm. 3.2). We refer to γ_{Bw} as the *Blackwell discount factor* and we say that optimal policies with discount factor $\gamma \geq \gamma_{Bw}$ are optimal according to the Blackwell optimality criterion. This concept closely resembles the idea of *effective planning horizon* by Laidlaw, Russell, and Dragan (2023), with the effective horizon defined as $1/(1-\gamma_{Bw})$ instead of a number of look ahead steps.

Planning Loss The planning loss captures the impact of using $\gamma < \gamma_{Bw}$ given that the planning is performed in an approximate model of the environment² \hat{M} with $\hat{M} \approx M$:

$$\begin{aligned} & \|V_{M,\gamma_{Bw}}^{\pi_{M,\gamma_{Bw}}^*} - V_{\hat{M},\gamma}^{\pi_{\hat{M},\gamma}^*}\|_\infty & (1) \\ &= \|V_{M,\gamma_{Bw}}^{\pi_{M,\gamma_{Bw}}^*} - V_{M,\gamma_{Bw}}^{\pi_{\hat{M},\gamma}^*} + V_{M,\gamma_{Bw}}^{\pi_{\hat{M},\gamma}^*} - V_{\hat{M},\gamma}^{\pi_{\hat{M},\gamma}^*}\|_\infty \\ &\leq \underbrace{\|V_{M,\gamma_{Bw}}^{\pi_{M,\gamma_{Bw}}^*} - V_{M,\gamma_{Bw}}^{\pi_{\hat{M},\gamma}^*}\|_\infty}_{\text{bias}} + \underbrace{\|V_{M,\gamma_{Bw}}^{\pi_{\hat{M},\gamma}^*} - V_{\hat{M},\gamma}^{\pi_{\hat{M},\gamma}^*}\|_\infty}_{\text{variance}}. & (2) \end{aligned}$$

This decomposition offers insight into two components which can affect the quality of the policy obtained when planning on an approximate model of the environment using a shallow planning horizon ($\gamma < \gamma_{Bw}$). The *bias* denotes the loss in value function (evaluated on the true MDP M and with the Blackwell planning horizon) when using a policy that is optimal with a shallow planning horizon instead of using a policy that is optimal with γ_{Bw} . On the other hand, the *variance* captures the impact of optimizing the policy under an approximate model \hat{M} with a shallow planning horizon and will tend to 0 with more data. This decomposition is different from previous work (Jiang et al. 2015) and has the advantage of being interpretable as a bias-variance trade-off from the PAC-learning literature. We can compare the bias to the approximation error and the variance to the estimation error (Shalev-Shwartz and Ben-David 2014).

²We remain agnostic to how \hat{M} is computed and how the data is collected as literature on the topic is abundant (Gheshlaghi Azar, Munos, and Kappen 2013; Wu, He, and Gu 2023; He, Zhou, and Gu 2021).

Structural parameters have been introduced previously to characterize the difficulty of an MDP.

Definition 1 (Value-function variation (Jiang, Singh, and Tewari 2016)). *Given an MDP M and discount factor γ , the value function variation captures the largest difference between the values of two different states when following the optimal policy:*

$$\kappa_{M,\gamma} = \max_{s,s' \in \mathcal{S}} \left| V_{M,\gamma}^{\pi_{M,\gamma}^*}(s) - V_{M,\gamma}^{\pi_{M,\gamma}^*}(s') \right| \leq \frac{R_{max}}{1-\gamma}.$$

Note that the value-function is evaluated with the same discount factor as used by the policy.

The value-function variation provides insight on the impact of starting in certain states over others in the MDP. A low value indicates that the starting state is not important to consider for predicting future rewards (Jiang, Singh, and Tewari 2016).

Definition 2 (Action variation (Jiang, Singh, and Tewari 2016)). *Given an MDP M with transition probabilities P , the action variation captures how much actions can impact state transitions:*

$$\delta_M = \max_{s \in \mathcal{S}} \max_{a,a' \in \mathcal{A}} \|P(\cdot|s,a) - P(\cdot|s,a')\|_1.$$

If the action variation is equal to 0, the agent cannot influence the state transitions (and therefore future rewards). In this case, we would expect the problem to be safely (and efficiently) formulated under the contextual bandit setting, which corresponds to using a myopic agent ($\gamma = 0$). The maximal value for this L_1 distance is 2, which often happens under deterministic settings. This is validated by our numerical experiments below.

Using Definitions 1 and 2, the following bound on the bias was introduced (Jiang, Singh, and Tewari 2016):

$$\underbrace{\|V_{M,\gamma_{Bw}}^{\pi_{M,\gamma_{Bw}}^*} - V_{M,\gamma_{Bw}}^{\pi_{\hat{M},\gamma}^*}\|_\infty}_{\text{bias}} \leq \frac{\delta_M/2 \cdot \kappa_{M,\gamma}(\gamma_{Bw} - \gamma)}{(1-\gamma_{Bw})(1-\gamma_{Bw}(1-\delta_M/2))}. \quad (3)$$

Unfortunately, the action variation is not sensitive to the planning horizon of the agent compared with the Blackwell planning horizon. Moreover, unlike the bias, there is no current bounds for the variance. We address these limitations with Lem. 1 and Eq. 3, which will allow us to obtain a new bound on the planning loss (upcoming Thm. 1).

Improving the Bias Bound

In order to consider the planning horizon in the existing bias bound (Eq. 3), we introduce the following definitions:

Definition 3 (Discordant state-action pairs). *The set of state-action pairs in an MDP M where policies π and π' differ:*

$$\mathcal{Z}_M(\pi \neq \pi') = \{(s,a) \in \mathcal{S} \times \mathcal{A} : \pi(s) \neq \pi'(s), \pi'(s) = a\}.$$

This new set will be used to capture the impact of a shallow-horizon policy on the action variation:

Definition 4 (Horizon-sensitive action variation). *The most important difference in transition probabilities induced by using discount factor γ instead of γ_{Bw} on an MDP M :*

$$\delta_{M,\gamma} = \max_{(s,a) \in \mathcal{Z}_M(\pi_{M,\gamma}^* \neq \pi_{M,\gamma_{Bw}}^*)} \|P(\cdot|s, \pi_{M,\gamma}^*(s)) - P(\cdot|s, a)\|_1.$$

The implementation of the action variations in proofs is to bound the difference in transition probabilities between two different policies. The highest possible bound is given by prior results (Def. 2), but we tighten this result by simply considering states for which the policies are unequal instead of all states. This has the benefit of being 0 when the policy is evaluated with a discount factor above the Blackwell. By building on previous analysis (Jiang, Singh, and Tewari 2016), we can obtain the following result to characterize the impact of optimizing the policy with a shallow horizon on k -steps transition probabilities.

Proposition 1 (Horizon-sensitive transition probabilities distance). *Given an MDP M , let $P_{s,k}^{\pi}$ denote the vector of the transition probabilities from state $s \in \mathcal{S}$ to every possible states when following policy π for $k \geq 1$ time steps. The transition probabilities when following the policy that is optimal for a shallow planning horizon ($\gamma < \gamma_{Bw}$) instead of following the policy that is optimal for the Blackwell planning horizon is bounded by:*

$$\|P_{s,k}^{\pi_{M,\gamma}^*} - P_{s,k}^{\pi_{M,\gamma_{Bw}}^*}\|_1 \leq 2 - 2(1 - \delta_{M,\gamma}/2)^k.$$

Prop. 1 can be used to extend Eq. 3 using the horizon-sensitive structural parameters:

$$\underbrace{\|V_{M,\gamma_{Bw}}^{\pi_{M,\gamma}^*} - V_{M,\gamma_{Bw}}^{\pi_{M,\gamma}^*}\|_\infty}_{\text{bias}} \leq \frac{\delta_{M,\gamma}/2 \cdot \kappa_{M,\gamma}(\gamma_{Bw} - \gamma)}{(1 - \gamma_{Bw})(1 - \gamma_{Bw}(1 - \delta_{M,\gamma}/2))}, \quad (4)$$

where we use the term $\delta_{M,\gamma}$ (instead of δ_M) to capture the divergence between the shallow and optimal-horizon policies. Since the set offered by Def. 3 is smaller than the set of all state-action pairs, Eq. 4 is tighter than Eq. 3. See Appendix B for the complete proofs.

Controlling the Variance

We now introduce new definitions and results to bound the variance (in Eq. 2). Recall that the variance captures the impact of learning an optimal policy on an (empirical) approximation \hat{M} of a true MDP M when using a shallow planning horizon ($\gamma < \gamma_{Bw}$). To this end, it is convenient to isolate the variance that does not depend on the shallow planning.

Definition 5 (Variance due to model approximation). *The maximum difference in value-function due to the approximate model \hat{M} :*

$$\hat{\epsilon} = \|V_{M,\gamma}^{\pi_{M,\gamma}^*} - V_{\hat{M},\gamma}^{\pi_{\hat{M},\gamma}^*}\|_\infty.$$

This term can be upper-bounded into the following results by using known settings in the PAC literature (Gheshlaghi Azar, Munos, and Kappen 2013; Wu, He, and Gu 2023;

He, Zhou, and Gu 2021). We can also use the discordant state-action pairs (Def. 3) to capture the action variation resulting from having optimized the policy on an approximation \hat{M} of a true MDP M .

Definition 6 (Empirical action variation). *The most important difference in transition probabilities when following the policy optimal on an MDP M vs the policy optimal on an approximate model \hat{M} :*

$$\hat{\delta}_{M,\gamma} = \max_{(s,a) \in \mathcal{Z}_M(\pi_{M,\gamma}^* \neq \pi_{\hat{M},\gamma}^*)} \|P(\cdot|s, \pi_{M,\gamma}^*(s)) - P(\cdot|s, a)\|_1.$$

The improvement over the action variation (Def. 2) is that it will tend towards 0 as $\hat{M} \approx M$ which is desirable in a bound on the variance. As was done previously in Definition 4, we can also build on previous analysis (Jiang, Singh, and Tewari 2016) to obtain the following result to characterize the impact of optimizing the policy with an approximate model \hat{M} on k -steps transition probabilities.

Proposition 2 (Empirical transition probabilities distance). *Given an MDP M and an approximate model \hat{M} . Let $P_{s,k}^{\pi}$ denote the vector of the transition probabilities from state $s \in \mathcal{S}$ to every possible states when following policy π for $k \geq 1$ time steps. For a given planning horizon, the transition probabilities when following the optimal policy on \hat{M} instead of following the optimal policy on M is bounded by:*

$$\|P_{s,k}^{\pi_{M,\gamma}^*} - P_{s,k}^{\pi_{\hat{M},\gamma}^*}\|_1 \leq 2 - 2(1 - \hat{\delta}_{M,\gamma}/2)^k.$$

Prop. 2 can be used with Def. 3 and 6 to obtain the following bound on the variance (see Appendix C).

Lemma 1 (Variance). *Consider optimal policies computed with a shallow planning horizon ($\gamma < \gamma_{Bw}$) on an MDP M and an approximate model \hat{M} . The difference between their value-functions evaluated on M with discount factor γ_{Bw} is bounded by:*

$$\begin{aligned} & \|V_{M,\gamma_{Bw}}^{\pi_{M,\gamma}^*} - V_{M,\gamma_{Bw}}^{\pi_{\hat{M},\gamma}^*}\|_\infty \\ & \leq \hat{\epsilon} \left(\frac{1 - \gamma}{1 - \gamma_{Bw}} \right) + \frac{\hat{\delta}_{M,\gamma}/2 \cdot \kappa_{M,\gamma}(\gamma_{Bw} - \gamma)}{(1 - \gamma_{Bw})(1 - \gamma_{Bw}(1 - \hat{\delta}_{M,\gamma}/2))}. \end{aligned}$$

This bound is interesting because it becomes tighter when the empirical action variation $\hat{\delta}_{M,\gamma}$ or the value function variation $\kappa_{M,\gamma}$ decrease. We can then deduce that a problem with a low value in these structural parameters lowers both the bias (Eq. 4) and the variance. Finally, the use of the empirical action variation (Def. 6) gives rise to a bound that is coherent in convergence, as it will tend towards 0 as the amount of data increases.

A New Bound on the Planning Loss

By combining the extended bias bound (Eq. 4) with our novel bound on the variance (Lem. 1), we obtain the following bound on the planning loss (Eq. 2). See Appendix D for the complete proof.




Figure 1: Proportion of randomly sampled MDPs where Eq. 5 is true given a discount factor γ .

Theorem 1 (Planning loss). *Given an MDP M , its Blackwell discount factor γ_{Bw} , and an approximate model \hat{M} . The planning loss is bounded by:*

$$\begin{aligned} & \|V_{M,\gamma_{Bw}}^{\pi_{M,\gamma_{Bw}}^*} - V_{\hat{M},\gamma_{Bw}}^{\pi_{\hat{M},\gamma_{Bw}}^*}\|_{\infty} \\ & \leq \kappa_{M,\gamma} \left(\frac{\gamma_{Bw} - \gamma}{1 - \gamma_{Bw}} \right) \left(\frac{\delta_{M,\gamma}/2}{1 - \gamma_{Bw}(1 - \delta_{M,\gamma}/2)} \right) \\ & \quad + \kappa_{M,\gamma} \left(\frac{\gamma_{Bw} - \gamma}{1 - \gamma_{Bw}} \right) \left(\frac{\hat{\delta}_{M,\gamma}/2}{1 - \gamma_{Bw}(1 - \hat{\delta}_{M,\gamma}/2)} \right) \\ & \quad + \hat{\epsilon} \left(\frac{1 - \gamma}{1 - \gamma_{Bw}} \right). \end{aligned}$$

This result provides insight into how structural parameters affect not only the bias, but also the variance. For instance, a problem with action variation $\delta_M \approx 0$ has low variance due to the limited impact of the policy over the state value (agent actions do not impact transition probabilities). Similarly to prior work (Jiang et al. 2015), although the applicability of this result is limited by not having access to the true model M , it remains a helpful guide to design heuristics and better understand how one could decide a discount factor. For example, it justifies framing recommender systems as contextual bandits when the outcome of future recommendations do not depend on current recommendations, which translates into a low value-function variation $\kappa_{M,\gamma}$. Thm. 1 is tighter than the current existing bound (Jiang et al. 2015) under the following condition on the quality of the model approximation \hat{M} (see Appendix E):

$$\hat{\epsilon} \leq \frac{R_{max}}{1 - \gamma} - \kappa_{M,\gamma} \left(\frac{\delta_{M,\gamma}/2}{1 - \gamma_{Bw}(1 - \delta_{M,\gamma}/2)} + \frac{\hat{\delta}_{M,\gamma}/2}{1 - \gamma_{Bw}(1 - \hat{\delta}_{M,\gamma}/2)} \right). \quad (5)$$

Fig. 1 supports the idea that Thm. 1 becomes tighter than prior results when the variance due to model approximation (Def. 5) is low or when $\frac{R_{max}}{1 - \gamma}$ is dominant (γ close to 1).

Bias Under Partial Observability

We now look at how partial observability impacts the structural parameters to better understand its impact on the bias.

This is important since most practical problems suffer from a form of partial observability (Dulac-Arnold et al. 2021). We consider a discrete-time POMDP (Sondik 1978) described by the MDP tuple extended with two elements: a finite set of possible observations Ω and the probabilities of receiving each observation given a state, $\mathbb{O} : \mathcal{S} \times \Omega \mapsto [0, 1]$. On each time step $t \in \mathbb{N}_0$, the current state $S_t \in \mathcal{S}$ leads the agent to receive an observation $O_t \in \Omega$ (using \mathbb{O}), an action $A_t \in \mathcal{A}$ is played, the environment transitions into the next (unknown) state S_{t+1} (using P) and generates an observed reward R_{t+1} (using R).

When facing a partially observable setting, an effective way of approximating a solution is to use a policy defined on compressed histories (Francois-Lavet et al. 2019). Let $\mathcal{H}_t = \Omega \times (\mathcal{A} \times R \times \Omega)^t$ denote the set of histories observed up to time t and let $\mathcal{H} = \bigcup_{t=0}^{\infty} \mathcal{H}_t$ denote the space of all possible histories. The belief state $b(s|H)$ is a vector where the i -th component ($i \in \{1, \dots, |S|\}$) corresponds to $\mathbb{P}(s = i|H)$, for history $H \in \mathcal{H}$. One can define a mapping $\phi : \mathcal{H} \mapsto \phi(\mathcal{H})$, where the set compressed history $\phi(\mathcal{H}) = \{\phi(H)|H \in \mathcal{H}\}$ is finite, which can be used as input to a policy $\pi : \phi(\mathcal{H}) \mapsto \mathcal{A}$. We consider POMDPs with sufficient mappings such that $\mathbb{P}(s|H) = \mathbb{P}(s|\phi(H))$, which defines an MDP on state space $\phi(\mathcal{H})$.

Given a POMDP (extended tuple) M and any given distribution \mathcal{D}_H over histories, one can define the expected return obtained over an infinite time horizon from a given history H , with $A_t \sim \pi(\phi(H_t))$:

$$V_{M,\gamma}^{\pi}(\phi(H)) = \mathbb{E}_{\substack{H' \sim \mathcal{D}_H: \\ \phi(H') = \phi(H)}} [V_M^{\pi}(H'|\phi)]$$

With $V_{M,\gamma}^{\pi}(H'|\phi)$ given by

$$V_{M,\gamma}^{\pi}(H'|\phi) = \mathbb{E}_{\pi} \left[\sum_{k=0}^{\infty} \gamma^k R_{t+k+1} \mid S_t \sim b(\cdot|H_t = H') \right].$$

For a given sufficient mapping ϕ , the optimal policy $\pi_{M,\gamma}^*$ maximizes $V_{M,\gamma}^{\pi}(\phi(H))$ for all histories $H \in \mathcal{H}$.

Extending Structural Parameters

We can extend Definitions 1 and 2 to the POMDP setting by applying them to compressed histories rather than the actual states in the underlying MDP:

$$\kappa_{M,\gamma}^{\phi} = \max_{\sigma, \sigma' \in \phi(\mathcal{H})} \left| V_{M,\gamma}^{\pi_{M,\gamma}^*}(\sigma) - V_{M,\gamma}^{\pi_{M,\gamma}^*}(\sigma') \right| \quad (6)$$

$$\delta_M^{\phi} = \max_{\sigma \in \phi(\mathcal{H})} \max_{a, a' \in \mathcal{A}} \sum_{\sigma' \in \phi(\mathcal{H})} |\mathbb{P}(\sigma'|\sigma, a) - \mathbb{P}(\sigma'|\sigma, a')|. \quad (7)$$

We introduce the following result showing how the structural parameters in the POMDP relate to the structural parameters of the underlying MDP (see Appendix F):

Theorem 2. *Given a POMDP M , let $\kappa_{M,\gamma}^S$ and δ_M^S denote the structural parameters (Definitions 1 and 2) evaluated on the underlying state space. Let $\mathcal{H}(s) = \bigcup_{t=0}^{\infty} \{H_t : b(s|H_t) > 0, H_t \in \mathcal{H}_t\}$ denote the set of all histories which can lead to being in state s at any time t , and*

$\mathcal{D}_H(s)$ denote its probability distribution. Define Δ_ϵ^ϕ s.t. $\left|V_{M,\gamma}^{\pi_{M,\gamma}^*}(s) - V_{M,\gamma}^{\pi_\phi}(s)\right| \leq \Delta_\epsilon^\phi \forall s \in \mathcal{S}$, define $\pi_\phi(s) = \mathbb{E}_{H \sim \mathcal{D}_H(s)} \pi_{M,\gamma}^*(\phi(H))$ as the optimal policy on compression of histories when executed over the underlying state space and define $b(s|\sigma) = \mathbb{E}_{H \sim \mathcal{D}_H(s), \sigma = \phi(H)} b(s|H)$ for $\sigma \in \phi(\mathcal{H})$. We have that:

$$\delta_M^\phi \leq \delta_M^S$$

$$\kappa_{M,\gamma}^\phi \leq \max_{\sigma, \sigma' \in \phi(\mathcal{H})} \frac{\|b(\cdot|\sigma) - b(\cdot|\sigma')\|_1}{2} (\kappa_{M,\gamma}^S + \Delta_\epsilon^\phi).$$

The first inequality confirms that partial observability impacts negatively the ability for the agent to control state transitions. The second inequality implies that if the policy on the partially observable domain remains accurate on the underlying state space ($\Delta_\epsilon^\phi \approx 0$), then the state-value variation observed by the agent is lower (since the L_1 distance is bounded by 2), which could make the learning task easier and as efficient with a low discount factor. This L_1 distance often has a value of 2, which makes the bound quite loose, but they illustrate the idea that the values of structural parameters decrease when taking a convex combination over states. In fact the maximal value of the expectation of a random variable happens if all the mass is concentrated on the maximal value of the support. This is explained further in appendix B. By considering that the horizon-sensitive action variation (Def. 4) is upper-bounded by the action variation (Def. 2), we can observe from Eq. 4 that the bias in the underlying MDP upper-bounds the bias of the POMDP when the optimal policy under partial observability is accurate on the true state space. This extends the ideas from Abel, Hershkowitz, and Littman (2016) where abstractions can make a problem much easier to learn while retaining good performance and in our case, lower the bias. We can get a better understanding of this trade-off in the state-abstraction setting (Abel, Hershkowitz, and Littman 2016), as shown in our numerical experiments below.

Numerical Experiments

We now conduct experiments³ to highlight the relationships between the planning horizon, the partial observability, and the structural parameters of the underlying MDP.

Random MDPs We consider the simulated environment of Jiang, Singh, and Tewari (2016). We use 2-action MDPs, with $Fixed(|\mathcal{S}|, d)$ denoting a randomly generated MDP with $d \geq 1$ next states reachable from each state. MDPs are sampled using the following procedure: 1) each state-action pair is assigned d possible next states; 2) transition probabilities to these states are sampled uniformly in $[0, 1]$, then normalized; 3) rewards are assigned to state-action pairs by sampling uniformly in $[0, 1]$.

³Experiments are conducted on a AMD Ryzen 5 3600 CPU and a GTX 1660 Ti GPU.

Extension to Partial Observability We consider the state-abstraction setting (Abel, Hershkowitz, and Littman 2016), which corresponds to a specific case of partial observability where the history compressor $\phi(\mathcal{H})$ returns only the last observation and where \mathbb{O} is a one-hot vector on an observation from Ω . For simplicity, we make sure that each observation is connected to at least one state. Using Bayes' theorem to recover the belief that the agent is in state s given observation ω , we get a constant uniform distribution on every state s which maps onto this observation:

$$b(s|\omega) = \frac{1}{|\{s \in \mathcal{S} : \mathbb{O}(o, s) > 0\}|} \forall s \in \mathcal{S} : \mathbb{O}(\omega, s) > 0, \quad \forall \omega \in \Omega.$$

and a belief of 0 otherwise. From this special case of POMDP, we can extract an abstract MDP $M_A = \langle \mathcal{S}_A, \mathcal{A}, P_A, R_A, \gamma \rangle$ from the underlying MDP $M = \langle \mathcal{S}, \mathcal{A}, P, R, \gamma \rangle$ (Abel, Hershkowitz, and Littman 2016):

$$R_A(\omega, a) = \sum_{s \in \mathcal{S}} b(s|\omega) R(s, a), \text{ and} \quad (8)$$

$$P_A(\omega, a, \omega') = \sum_{s \in \mathcal{S}} \sum_{s' \in \mathcal{S}} P(s, a, s') b(s|\omega) \mathbb{O}(\omega', s'). \quad (9)$$

For our experiments under partial observability, we start by sampling an MDP from $Fixed(|\mathcal{S}|, d)$. Then, we can map it onto abstracted MDPs (POMDPs) with different number of observations as $|\Omega|$. The number of observations encodes the level of partial observability akin to what happens in discretization. For $|\Omega| = |\mathcal{S}|$, the problem is fully observable. For $|\Omega| = 1$ (and $|\mathcal{S}| > 1$), the agent is completely blind to the state. We sample 10^4 MDPs with $Fixed(10, 3)$ and abstract each MDP into 6 configurations: $|\Omega| \in \{10, 8, 6, 4, 2, 1\}$. The Blackwell discount factors are computed by iterating from $\gamma = 1$ to $\gamma = 0$ with step size of 0.01 until the optimal policy changes.

Fig. 2 (left) shows that the mass of the Blackwell planning horizon tends to decrease as the observability decreases. Since the bias is null when planning with a discount factor larger than γ_{Bw} , we only cumulate variance above that point. When $|\Omega| = 1$, the myopic agent ($\gamma = 0$) enjoys the optimal planning horizon, which corresponds to the bandit setting. We also evaluate the normalized bias

$$\max_{s \in \mathcal{S}} \left(V_{M,\gamma_{Bw}}^{\pi_{M,\gamma_{Bw}}^*}(s) - V_{M,\gamma_{Bw}}^{\pi_\phi}(s) \right) / V_{M,\gamma_{Bw}}^{\pi_{M,\gamma_{Bw}}^*}(s)$$

for different planning horizons, averaged over different levels of partial observability. Fig. 2 (right) shows that, although the bias decreases when the planning horizon ($1/(1-\gamma)$) increases, this effect attenuates as the observability decreases. Since many real-world problems are partially observable, this supports the need to consider shallow planning more seriously.

We normalize parameters for each abstract MDP by dividing structural parameters by those of the true $Fixed(\cdot)$ MDP, i.e., $\kappa_{M,\gamma}^\phi / \kappa_{M,\gamma}^S$ for the normalized value function variation and $\delta_M^\phi / \delta_M^S$ for the normalized action variation.




Figure 2: Left: Distribution of Blackwell discount factors over 10^4 POMDPs given the number of observations. Right: Average normalized bias given the discount factor and number of observations.

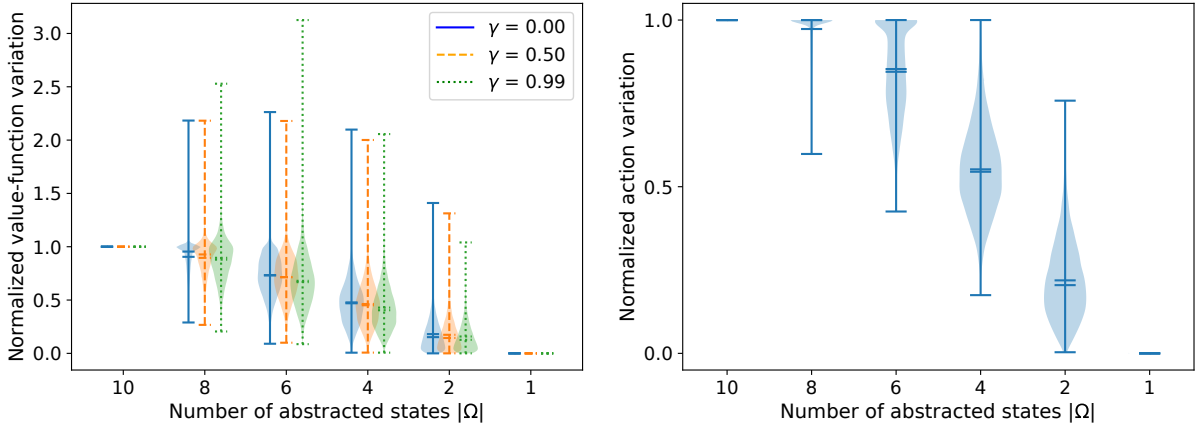


Figure 3: Left: Distribution of normalized $\kappa_{M,\gamma}^\phi$ over 10^4 POMDPs given the number of observations and the discount factor used. Right: Distribution of normalized δ_M^ϕ given the number of observations.

Fig. 3 (left) highlights the tradeoff given by Thm. 2. We see that there is a downwards trend, but for high values of gamma, the error Δ_ϵ^ϕ in some cases is so high that the value-function variation can be higher under partial observability. Fig. 3 (right) highlights the strict inequality offered by Thm. 2 on the action variation of the POMDP vs the underlying MDP. By observing Thm. 1 and Fig. 3, the reduction in structural parameters offers insight into why the bias decreases under partial observability. It also points to the fact that there might exist a bound on the Blackwell discount factor using structural parameters much tighter than Grand-Clément and Petrik (2023) or even improve upon the results of Laidlaw, Russell, and Dragan (2023) using partial observability.

Impact of Partial Observability on Deep RL

We explore the interaction of shallow planning and partial observability in the Cartpole-v1 environment (Towers et al. 2024). In Cartpole, an agent uses two discrete actions (left or right) to balance a pole on a moving cart. The agent’s

state includes the cart’s position and velocity, along with the pole’s angle and angular velocity. A reward of +1 is given for each time step before episode termination or reaching 500 time steps.

We consider the widely used PPO (Schulman et al. 2017) agent policy with the recommended hyperparameters for this task (Raffin et al. 2021). We then consider different discount factors $\gamma \in \{0, 0.24, 0.49, 0.73, 0.98\}$, with the largest $\gamma = 0.98$ from the baseline (Raffin et al. 2021). For the partially observable component, we simulate noisy sensors, which are common in real life. These are simulated by injecting noise into the state. The noise is sampled from a multivariate normal distribution $\mathcal{N}(\mathbf{0}, I\sigma^2)$ parameterized by a diagonal covariance matrix with value σ^2 on the diagonal. We consider $\sigma \in \{0, 0.1, 1\}$. We train 10 agents for each combination of (γ, σ) , resulting in 150 models, and evaluate each of these models on 100 unseen seeds.

Fig. 4 displays the average reward obtained by running the 10 models associated with each (γ, σ) configuration on




Figure 4: Average reward and standard deviation obtained by running 10 models on 100 environment seeds given the noise level and discount factor.

the 100 environment seeds. Without surprise, planning with the longest horizon (largest discount factor) maximizes rewards (as expected on this task). However, we also observe that, as the environment becomes partially observable (e.g. $\sigma = 0.1$), shallow planning can be better than long planning since the bias is lowered from the partial observability. The Blackwell discount factor could be lowered, supporting our previous results and thesis. As expected, when observability becomes too low (e.g. $\sigma = 1$), the agent cannot understand the environment anymore and is unable to learn the task no matter the discount factor considered.

Related Work

Effective Planning Horizon When it comes to using RL in practical applications, there is a trend towards a better understanding of the planning horizon on the optimality of the policy learned. Grand-Clément and Petrik (2023) explore the Blackwell optimality criterion and proves the existence of a planning horizon above which the policy learned cannot be improved. Similarly, Laidlaw, Russell, and Dragan (2023) explore the minimal horizon (in number of time steps forward) for which the optimal policy can be retrieved with high probability. Our work builds on top of these frameworks to better understand the bias-variance trade-off that arises from using a shallow planning horizon.

Planning Loss The proposed decomposition of the planning loss (Eq. 2) enables its interpretation in terms of bias and variance, unlike the previous decomposition of Jiang et al. (2015). The bias term had been studied previously (Jiang, Singh, and Tewari 2016), borrowing inspiration on the approximation term from the PAC bound literature in machine learning (Shalev-Shwartz and Ben-David 2014). However, the variance term remained unbounded. By connecting ideas from previous works Jiang et al. (2015); Jiang, Singh, and Tewari (2016), we are now able to understand not only the bias as a function of the structural parameters, but also part of the variance.

Partial Observability Most practical problem are under some form of partial observability (Dulac-Arnold et al. 2021). Different approximation schemes are explored through the literature like memoryless policies (Müller and

Montúfar 2021) or compressions of histories (Francois-Lavet et al. 2019). There is also plenty of work on properties of POMDPs such that approximations are more likely to yield good results (Liu et al. 2022). Our work extends these results to a better understanding of how these approximation schemes can impact not only the sample complexity or the performance guarantees, but the planning horizon that should be considered.

In this direction, the literature on state abstraction aims to tackle some of the challenges of practical RL through a simpler state-space (Abel, Hershkowitz, and Littman 2016; Allen et al. 2021; Abel et al. 2018). This approximation scheme is well understood, and lots of useful properties on policy performance or sample complexity have been found. These results have not yet been connected to the planning horizon that should be used under these regimes. We attempt to bridge this gap by providing insight into the behaviour of the bias under partial observability.

Conclusion

We extend existing structural parameters to consider the planning horizon (Def. 4) and the model approximation (Def. 6). This allow us to extend an existing bound on the bias (Eq. 4) and propose a new bound on the variance (Lem. 1), which result in a new bound on the planning loss (Thm. 1). We finally extend the structural parameters to POMDPs (Eq. 6 and 7) and show that these are controlled by their fully observable counterparts (Thm. 2). This complements previous results (Abel, Hershkowitz, and Littman 2016; Francois-Lavet et al. 2019) by considering the impact of the planning horizon on the bias when shallow planning under partial observability.

Limitations and Future Work Our work motivates further research on understanding when shallow planning might be beneficial by providing new theory. However, our results are limited to a finite state space. Also, the reduction in $\kappa_{M,\gamma}^S$ (Thm. 2) that comes from partial observability can be much larger than what is captured by the L_1 distance; the bound could be improved by using better mathematical tools. Even with these limitations, practitioners can use our analysis and our experiments to better understand how the nature of their problem might impact the bias-variance trade-off, and select a better planning horizon. Future work could improve upon our results by using deep-learning theory. It could also be possible to develop automatic discount factor selection algorithms, as previous research papers have indicated that even naive form of this strategy can be very impactful (François-Lavet, Fonteneau, and Ernst 2015). State-specific discount factors are also interesting. It is easy to construct an MDP such that the Blackwell discount factor is very high because of only a handful of states in the state space. A better understanding of the local structure around a state means we could create heuristics to adapt the discount factor automatically (Pitis 2019; Yoshida, Uchibe, and Doya 2013). Finally, a better understanding of the process which could lead to the downward trend in γ_{Bw} observed in Fig. 2 through our analysis could provide better insight into how to prevent over-fitting.

Acknowledgements

We acknowledge financial support from Intact and CIFAR.

References

- Abel, D.; Arumugam, D.; Lehnert, L.; and Littman, M. 2018. State Abstractions for Lifelong Reinforcement Learning. In *International Conference on Machine Learning*.
- Abel, D.; Hershkowitz, D.; and Littman, M. 2016. Near optimal behavior via approximate state abstraction. In *International Conference on Machine Learning*.
- Allen, C.; Parikh, N.; Gottesman, O.; and Konidaris, G. 2021. Learning Markov State Abstractions for Deep Reinforcement Learning. In *Advances in Neural Information Processing Systems*.
- Amit, R.; Meir, R.; and Ciosek, K. 2020. Discount Factor as a Regularizer in Reinforcement Learning. In *International Conference on Machine Learning*.
- Bastani, H.; and Bayati, M. 2020. Online decision making with high-dimensional covariates. *Operations Research*.
- Devidze, R.; Radanovic, G.; Kamalaruban, P.; and Singla, A. 2021. Explicable reward design for reinforcement learning agents. *Advances in Neural Information Processing Systems*.
- Ding, K.; Li, J.; and Liu, H. 2019. Interactive Anomaly Detection on Attributed Networks. In *ACM International Conference on Web Search and Data Mining*.
- Dulac-Arnold, G.; Levine, N.; Mankowitz, D. J.; Li, J.; Paduraru, C.; Goyal, S.; and Hester, T. 2021. An empirical investigation of the challenges of real-world reinforcement learning. *Machine Learning*.
- Durand, A.; Achilleos, C.; Iacovides, D.; Strati, K.; Mitsis, G. D.; and Pineau, J. 2018. Contextual bandits for adapting treatment in a mouse model of de novo carcinogenesis. In *Machine Learning for Healthcare Conference*.
- François-Lavet, V.; Fonteneau, R.; and Ernst, D. 2015. How to discount deep reinforcement learning: Towards new dynamic strategies. arXiv:1512.02011.
- Francois-Lavet, V.; Rabusseau, G.; Pineau, J.; Ernst, D.; and Fonteneau, R. 2019. On Overfitting and Asymptotic Bias in Batch Reinforcement Learning with Partial Observability. *Journal of Artificial Intelligence Research*.
- Gheshlaghi Azar, M.; Munos, R.; and Kappen, H. J. 2013. Minimax PAC bounds on the sample complexity of reinforcement learning with a generative model. *Machine Learning*.
- Grand-Clément, J.; and Petrik, M. 2023. Reducing Blackwell and Average Optimality to Discounted MDPs via the Blackwell Discount Factor. In *Advances in Neural Information Processing Systems*.
- Gu, S.; Yang, L.; Du, Y.; Chen, G.; Walter, F.; Wang, J.; Yang, Y.; and Knoll, A. 2023. A Review of Safe Reinforcement Learning: Methods, Theory and Applications. arXiv:2205.10330.
- Hare, J. 2019. Dealing with sparse rewards in reinforcement learning. arXiv:1910.09281.
- He, J.; Zhou, D.; and Gu, Q. 2021. Uniform-PAC bounds for reinforcement learning with linear function approximation. *Advances in Neural Information Processing Systems*.
- Jiang, N.; Kulesza, A.; Singh, S.; and Lewis, R. 2015. The dependence of effective planning horizon on model accuracy. In *International Conference on Autonomous Agents and Multiagent Systems*.
- Jiang, N.; Singh, S.; and Tewari, A. 2016. On Structural Properties of MDPs that Bound Loss Due to Shallow Planning. In *International Joint Conference on Artificial Intelligence*.
- Kaelbling, L. P.; Littman, M. L.; and Cassandra, A. R. 1998. Planning and acting in partially observable stochastic domains. *Artificial intelligence*.
- Kaiser, L.; Babaeizadeh, M.; Milos, P.; Osinski, B.; Campbell, R. H.; Czechowski, K.; Erhan, D.; Finn, C.; Kozakowski, P.; Levine, S.; Mohiuddin, A.; Sepassi, R.; Tucker, G.; and Michalewski, H. 2024. Model-Based Reinforcement Learning for Atari. arXiv:1903.00374.
- Laidlaw, C.; Russell, S. J.; and Dragan, A. 2023. Bridging RL theory and practice with the effective horizon. *Advances in Neural Information Processing Systems*.
- Liu, Q.; Chung, A.; Szepesvári, C.; and Jin, C. 2022. When Is Partially Observable Reinforcement Learning Not Scary? arXiv:2204.08967.
- Mnih, V.; Kavukcuoglu, K.; Silver, D.; Graves, A.; Antonoglou, I.; Wierstra, D.; and Riedmiller, M. 2013. arXiv:1312.5602.
- Müller, J.; and Montúfar, G. 2021. The geometry of memoryless stochastic policy optimization in infinite-horizon POMDPs. arXiv:2110.07409.
- Pitis, S. 2019. Rethinking the discount factor in reinforcement learning: A decision theoretic approach. In *AAAI conference on Artificial Intelligence*.
- Raffin, A.; Hill, A.; Gleave, A.; Kanervisto, A.; Ernestus, M.; and Dormann, N. 2021. Stable-Baselines3: Reliable Reinforcement Learning Implementations. *Journal of Machine Learning Research*, 22(268): 1–8.
- Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal policy optimization algorithms. arXiv:1707.06347.
- Shalev-Shwartz, S.; and Ben-David, S. 2014. *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press, 1st edition.
- Sondik, E. J. 1978. The optimal control of partially observable Markov processes over the infinite horizon: Discounted costs. *Operations research*.
- Sutton, R. S.; and Barto, A. G. 2018. *Reinforcement learning: an introduction*. Adaptive computation and machine learning series. Cambridge, Massachusetts: The MIT Press, second edition edition.
- Towers, M.; Kwiatkowski, A.; Terry, J. K.; Balis, J. U.; Cola, G.; Deleu, T.; Goulão, M.; Kallinteris, A.; Krimmel, M.; KG, A.; Perez-Vicente, R.; Pierré, A.; Schulhoff, S.; Tai, J. J.; Tan, H. J. S.; and Younis, O. G. 2024. Gymnasium:

A Standard Interface for Reinforcement Learning Environments. arXiv:2407.17032.

Wu, Y.; He, J.; and Gu, Q. 2023. Uniform-PAC guarantees for model-based RL with bounded eluder dimension. In *Uncertainty in Artificial Intelligence*.

Yoshida, N.; Uchibe, E.; and Doya, K. 2013. Reinforcement learning with state-dependent discount factor. In *International Conference on Development and Learning and epigenetic robotics*.

Yu, Y. 2018. Towards Sample Efficient Reinforcement Learning. In *International Joint Conference on Artificial Intelligence*.