

On Overfitting and Asymptotic Bias in Batch Reinforcement Learning with Partial Observability

Vincent François-Lavet
Guillaume Rabusseau
Joelle Pineau

*School of Computer Science, McGill University
University Street 3480, Montreal, QC, H3A 2A7, Canada*

VINCENT.FRANCOIS-LAVET@MAIL.MCGILL.CA
GUILLAUME.RABUSSEAU@MAIL.MCGILL.CA
JPINEAU@CS.MCGILL.CA

Damien Ernst

Raphael Fonteneau

*Montefiore Institute, University of Liege
Allée de la découverte 10, 4000 Liège, Belgium*

DERNST@ULIEGE.BE

RAPHAEL.FONTENEAU@ULIEGE.BE

Abstract

This paper provides an analysis of the tradeoff between asymptotic bias (suboptimality with unlimited data) and overfitting (additional suboptimality due to limited data) in the context of reinforcement learning with partial observability. Our theoretical analysis formally characterizes that while potentially increasing the asymptotic bias, a smaller state representation decreases the risk of overfitting. This analysis relies on expressing the quality of a state representation by bounding L_1 error terms of the associated belief states. Theoretical results are empirically illustrated when the state representation is a truncated history of observations, both on synthetic POMDPs and on a large-scale POMDP in the context of smartgrids, with real-world data. Finally, similarly to known results in the fully observable setting, we also briefly discuss and empirically illustrate how using function approximators and adapting the discount factor may enhance the tradeoff between asymptotic bias and overfitting in the partially observable context.

1. Introduction

This paper studies sequential decision-making problems that may be modeled as Markov Decision Processes (MDP) but for which the state is partially observable. This class of problems is called Partially Observable Markov Decision Processes (POMDPs) (Sondik, 1978). Within this setting, we focus on decision-making strategies computed using Reinforcement Learning (RL). When the model of the environment is not available, RL approaches rely on observations gathered through interactions with the (PO)MDP, and, although some RL approaches have strong convergence guarantees, classic RL approaches are challenged by data scarcity. When acquisition of new observations is possible (the “online” case), data scarcity is gradually phased out using strategies balancing the exploration / exploitation (E/E) tradeoff. The scientific literature related to this topic is vast; in particular, Bayesian RL techniques (Ross, Pineau, Chaib-draa, & Kreitmann, 2011; Ghavamzadeh, Mannor, Pineau, & Tamar, 2015) offer an elegant way of formalizing the E/E tradeoff.

However, such E/E strategies are not applicable when the acquisition of new observations is not possible anymore. In the pure “batch” setting, the task is to learn the best possible policy from a fixed set of transition samples (Farahmand, 2011; Lange, Gabel, & Riedmiller,

2012). Within this context, we propose to revisit RL as a learning paradigm that faces, similarly to supervised learning, a tradeoff between simultaneously minimizing two sources of error: an asymptotic bias and an overfitting error. The asymptotic bias (also simply called bias in the following) directly relates to the choice of the RL algorithm (and its parameterization). Any RL algorithm defines a policy class as well as a procedure to search within this class, and the bias may be defined as the performance gap between actual optimal policies and the best policies within the policy class considered. This bias does not depend on the set of observations. On the other hand, overfitting is an error term induced by the fact that only a limited amount of data is available to the algorithm. This overfitting error vanishes as the size and the quality of the dataset increase.

In this paper, we focus on studying the interactions between these two sources of error, in a setting where the state is partially observable. Due to this particular setting, one needs to build a state representation from the sequence of observations, actions and rewards in a trajectory (Singh, Jaakkola, & Jordan, 1994; Aberdeen, 2003). By increasing the cardinality of the state representation, the algorithm may be provided with a more informative representation of the POMDP, but at the price of simultaneously increasing the size of the set of candidate policies, thus also increasing the risk of overfitting. We analyze this tradeoff in the case where the RL algorithm provides an optimal solution to the frequentist-based MDP associated with the state representation (independently of the method used by the learning algorithm to converge towards that solution). Our novel analysis relies on expressing the quality of a state representation by bounding L_1 error terms of the associated belief states, thus introducing the concept of ϵ -sufficient statistics in the hidden state dynamics.

Experimental results illustrate the theoretical findings on a distribution of synthetic POMDPs as well as a large-scale POMDP with real-world data. In addition, we illustrate the link between the variance observed when dealing with different datasets (directly linked to the size of the dataset) and overfitting, where the link is that variance leads to overfitting if we have a (too) large feature space.

By extending known results for MDPs, we also briefly discuss and illustrate how using function approximators and adapting the discount factor play a role in the tradeoff between bias and overfitting when the state is partially observable. This has the advantage of providing the reader with an overview of key elements involved in the bias-overfitting tradeoff, specifically for the POMDP case.

The remainder of the paper is organized as follows. Section 2 formalizes POMDPs, (limited) sets of observations and state representations. Section 3 details the main contribution of this paper: an analysis of the bias-overfitting tradeoff in learning POMDPs in the batch setting. Section 4 empirically illustrates the main theoretical results, while Section 5 concludes with a discussion of the findings

2. Formalization

We consider a discrete-time POMDP (Sondik, 1978) model M described by the tuple $(\mathcal{S}, \mathcal{A}, T, R, \Omega, O, \gamma)$ where

- \mathcal{S} is a finite set of states $\{1, \dots, N_S\}$,

- \mathcal{A} is a finite set of actions $\{1, \dots, N_{\mathcal{A}}\}$,
- $T : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$ is the transition function (set of conditional transition probabilities between states),
- $R : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathcal{R}$ is the reward function, where \mathcal{R} is a continuous set of possible rewards in a range $R_{max} \in \mathbb{R}^+$ (e.g., $[0, R_{max}]$ without loss of generality),
- Ω is a finite set of observations $\{1, \dots, N_{\Omega}\}$,
- $O : \mathcal{S} \times \Omega \rightarrow [0, 1]$ is a set of conditional observation probabilities, and
- $\gamma \in [0, 1)$ is the discount factor.

The initial state is drawn from an initial distribution $b(s_0)$. At each time step $t \in \mathbb{N}_0$, the environment is in a state $s_t \in \mathcal{S}$. At the same time, the agent receives an observation $\omega_t \in \Omega$ which depends on the state of the environment with probability $O(s_t, \omega_t)$ and the agent has to take an action $a_t \in \mathcal{A}$. Then, the environment transitions to state $s_{t+1} \in \mathcal{S}$ with probability $T(s_t, a_t, s_{t+1})$ and the agent receives a reward $r_t \in \mathcal{R}$ equal to $R(s_t, a_t, s_{t+1})$. In this paper, the conditional transition probabilities T , the reward function R and the conditional observation probabilities O are unknown. The only information available to the agent is the past experience it gathered while interacting with the POMDP. A POMDP is illustrated in Fig. 1.




Figure 1: Graphical model of a POMDP.

2.1 Processing a History of Data

Policies considered in this paper are mappings from (an ordered set of) observation(s) into actions. A simple approach to build a space of candidate policies is to consider the set of mappings taking only the very last observation(s) as input (Whitehead & Ballard, 1990). However, in a POMDP setting, this leads to candidate policies that are likely not rich enough to capture the system dynamics, thus suboptimal (Singh et al., 1994; Wolfe, 2006). The alternative is to use a history of previously observed features to better estimate the hidden state dynamics (McCallum, 1996; Littman & Sutton, 2002; Singh, James, & Rudary, 2004).

We denote by $\mathcal{H}_t = \Omega \times (\mathcal{A} \times \mathcal{R} \times \Omega)^t$ the set of histories observed up to time t for $t \in \mathbb{N}_0$, and by $\mathcal{H} = \bigcup_{t=0}^{\infty} \mathcal{H}_t$ the space of all possible observable histories.

A straightforward approach is to take the whole history $H_t \in \mathcal{H}$ as input of candidate policies. However, taking a too long history may have several drawbacks. Indeed, increasing the size of the set of candidate optimal policies generally implies: (i) more computation to search within this set (Littman, 1994; McCallum, 1996) and (ii) an increased risk of including candidate policies suffering overfitting (see Section 3). In this paper, we are specifically interested in minimizing the latter overfitting drawback while keeping an informative state representation.

In this paper, we consider a mapping $\phi : \mathcal{H} \rightarrow \phi(\mathcal{H})$, where $\phi(\mathcal{H}) = \{\phi(H) | H \in \mathcal{H}\}$ is of finite cardinality $|\phi(\mathcal{H})|$. On the one hand, we will show that when ϕ discards information from the whole history, the state representation $\phi(H)$ that the agent uses to take decision might depart from sufficient statistics, which can hurt performance. On the other hand, we will show that it is beneficial to use a mapping ϕ that has a low cardinality $|\phi(\mathcal{H})|$ to avoid overfitting. This can be intuitively understood since a mapping $\phi(\cdot)$ induces an upper bound on the number of candidate policies: $|\Pi_{\phi(\mathcal{H})}| \leq |\mathcal{A}|^{|\phi(\mathcal{H})|}$. In the following, we discuss this tradeoff formally.

Let us first introduce a notion of information on the latent hidden state s through the notion of belief state (Cassandra, Kaelbling, & Littman, 1994).

Definition 2.1. *The belief state $b(s|H)$ is defined as the vector of probabilities where the i^{th} component ($i \in \{1, \dots, N_S\}$) is given by $\mathbb{P}(s = i | H)$, for any history $H \in \mathcal{H}$.*

Definition 2.2. *The belief state $b_{\phi}(s|\phi(H))$ is defined as the vector of probabilities where the i^{th} component ($i \in \{1, \dots, N_S\}$) is given by $\mathbb{P}(s = i | \phi(H))$, for any history $H \in \mathcal{H}$.¹*

Among all possible mappings ϕ , we are particularly interested in the ones that extract enough information from the history to accurately capture the corresponding belief state, with the notion of sufficient statistics (Kaelbling, Littman, & Cassandra, 1998; Aberdeen, Buffet, & Thomas, 2007). We thus define this notion in the context of the mapping ϕ .

Definition 2.3. *In a POMDP M , a statistic $\phi(H)$ is a sufficient statistic at the condition that $\forall s \in \mathcal{S}$:*

$$\mathbb{P}(s | H) = \mathbb{P}(s | \phi(H)),$$

for $H \in \mathcal{H}$. A mapping ϕ which provides sufficient statistics for all histories $H \in \mathcal{H}$ is called a sufficient mapping and is denoted as ϕ_0 .

However, as mentioned previously, it may happen that a mapping ϕ does not capture enough information from a history to contain the same information than the belief state. One of the key notions on which our analysis relies is the one of ‘‘approximately sufficient mappings’’, i.e. mappings whose corresponding belief state lies in an L_1 -ball of radius ϵ centered on $b(\cdot|H)$:

1. Note that s and H are random variables and their exact distribution will depend on the context that is considered. For any given probability distribution \mathcal{D}_H over histories: $H \sim \mathcal{D}_H$, the probability $P(s|\phi(H))$ is the expectation of the state when $\phi(H)$ is observed: $b_{\phi}(s | \varphi) = \mathbb{E}_{H \sim \mathcal{D}_H, \varphi = \phi(H)} b(s | H)$.

Definition 2.4. In a POMDP M , a statistic $\phi(H)$ is an ϵ -sufficient statistic at the condition that it meets the following condition with $\epsilon \geq 0$ and with the L_1 norm:

$$\|b_\phi(\cdot|\phi(H)) - b(\cdot|H)\|_1 \leq \epsilon,$$

for $H \in \mathcal{H}$. A mapping ϕ that provides ϵ -sufficient statistics for all histories $H \in \mathcal{H}$ is called an ϵ -sufficient mapping and is denoted as $\phi_\epsilon(H)$.

2.2 Working with a Limited Dataset

Let $\mathcal{M}(\mathcal{S}, \mathcal{A}, \Omega, \gamma)$ be a set of POMDPs with fixed \mathcal{S} , \mathcal{A} , Ω , and γ . For any POMDP $M(T, R, O) \in \mathcal{M}$, we denote by $D_{M, \pi_s, N_{tr}, N_l}$ a random dataset generated according to a probability distribution $\mathcal{D}_{M, \pi_s, N_{tr}, N_l}$ over the set of N_{tr} trajectories of length N_l . One such trajectory is defined as the observable history $H_{N_l} \in \mathcal{H}_{N_l}$ obtained in M when starting from s_0 and following a stochastic sampling policy π_s that ensures a non-zero probability of taking any action given an observable history $H \in \mathcal{H}$. For simplicity we denote $D_{M, \pi_s, N_{tr}, N_l}$, simply as $D \sim \mathcal{D}_M$. For the purpose of the analysis, we also introduce the asymptotic dataset $D_\infty = D_{M, \pi_s, N_{tr} \rightarrow \infty, N_l \rightarrow \infty}$ that would be theoretically obtained in the case where one could generate an infinite number of observations ($N_{tr} \rightarrow \infty$ and $N_l \rightarrow \infty$).

In this paper, the algorithm cannot generate additional data. The challenge is to determine a high-performance policy (in the actual environment) while having only access to a fixed dataset D . We formalize this hereafter.

2.3 Assessing the Performance of a Policy

Let us consider stationary and deterministic control policies $\pi : \phi(\mathcal{H}) \rightarrow \mathcal{A}$ with $\pi \in \Pi$. Any particular choice of ϕ induces a particular definition of the policy space Π . We introduce $V_M^\pi(\phi(H))$ with $H \in \mathcal{H}$ as the expected return obtained over an infinite time horizon when the system is controlled using policy π in the POMDP M . For any given distribution \mathcal{D}_H over histories, this is defined as:

$$V_M^\pi(\phi(H)) = \mathbb{E}_{\substack{H' \sim \mathcal{D}_H \\ \phi(H') = \phi(H)}} V_M^\pi(H' | \phi),$$

with $V_M^\pi(H | \phi)$ given by

$$V_M^\pi(H | \phi) = \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t r_t | s_0 \sim b(\cdot|H), \pi \right],$$

where we have $\mathbb{P}(\omega_t | s_t) = O(s_t, \omega_t)$, $a_t = \pi(\phi(H_t))$, $\mathbb{P}(s_{t+1} | s_t, a_t) = T(s_t, a_t, s_{t+1})$ and $r_t = R(s_t, a_t, s_{t+1})$.

We also define π^* as an optimal policy in M :

$$\pi^* \in \operatorname{argmax}_{\pi: \phi_0(\mathcal{H}) \rightarrow \mathcal{A}} V_M^\pi(\phi_0(H_0)),$$

where H_0 is taken out of the distribution of initial observations (compatible with the distribution $b(s_0)$ of initial states through the conditional observation probabilities).

3. Bias-overfitting in RL with Partial Observability

In this section, we study the performance difference (or gap) between the expected return that can be obtained following the policy built from limited data and the highest possible expected return that we would obtain if the algorithm had access to the POMDP parameters. In particular, we analyze how this performance gap can be decomposed into the sum of two terms: a term related to an asymptotic bias (suboptimality with unlimited data) and a term due to overfitting (additional suboptimality due to limited data).

3.1 Importance of the Feature Space

To study the importance of the feature space, let us assume that the policies built from limited data are optimal according to frequentist statistics, which allows removing from the analysis how the RL algorithm converges. In order to define the optimal policy according to frequentist statistics, let us first introduce a frequentist-based (augmented) MDP from the dataset D :

Definition 3.1. *With M defined by $(S, A, T, R, \Omega, O, \gamma)$ and the dataset D built from interactions with M , the frequentist-based augmented MDP $\hat{M}_{D,\phi}$, also denoted for simplicity $\hat{M}_D = (\Sigma, A, \hat{T}, \hat{R}, \Gamma)$, is defined with*

- *the state space: $\Sigma = \phi(\mathcal{H})$,*
- *the action space: $A = \mathcal{A}$,*
- *the estimated transition function: for $\sigma, \sigma' \in \Sigma$ and $a \in A$, $\hat{T}(\sigma, a, \sigma')$ is the number of times we observe the transition $(\sigma, a) \rightarrow \sigma'$ divided by the number of times we observe (σ, a) ²,*
- *the estimated reward function: for $\sigma, \sigma' \in \Sigma$ and $a \in A$, $\hat{R}(\sigma, a, \sigma')$ is the mean of the rewards observed for the tuple (σ, a, σ') ³, and*
- *the discount factor $\Gamma \leq \gamma$.*

As long as the mapping ϕ is a sufficient mapping (thus denoted ϕ_0), the asymptotic frequentist-based MDP (when unlimited data is available) actually gathers the relevant information from the actual POMDP. Indeed, when the POMDP is known (i.e. T, R, O are known), the knowledge of H_t allows one to obtain the belief state $b(s_t|H_t)$, calculated recursively thanks to the Bayes rule based on $b(s_t|H_t) = \mathbb{P}(s_t|\omega_t, a_t, b(s_{t-1}|H_{t-1}))$. It is then possible to define, from the history $H \in \mathcal{H}$ and for any action $a \in \mathcal{A}$, the expected

2. if any (σ, a) has never been encountered in a dataset, we arbitrarily set $\hat{T}(\sigma, a, \sigma') = 1/|\Sigma|, \forall \sigma'$. The theoretical results that follow are independent of how this case is treated.

3. if any (σ, a, σ') has never been encountered in a dataset, we arbitrarily set $\hat{R}(\sigma, a, \sigma')$ to the average of rewards observed over the whole dataset D . The theoretical results that follow are independent of how this case is treated.

immediate reward as well as a transition function into the next observation ω' :

$$R_{model-based}(H, a) = \sum_{s' \in \mathcal{S}} \sum_{s \in \mathcal{S}} b(s|H)T(s, a, s')R(s, a, s'), \text{ and}$$

$$T_{model-based}(H, a, \omega') = \underbrace{\sum_{s' \in \mathcal{S}} \sum_{s \in \mathcal{S}} b(s|H)T(s, a, s')O(s', \omega')}_{\text{next belief state}}.$$

In the frequentist approach, this information is estimated directly from interactions with the POMDP in \hat{R} and \hat{T} without any explicit knowledge of the POMDP parameters.

We introduce $\mathcal{V}_{\hat{M}_D}^\pi(\sigma)$ with $\sigma \in \Sigma$ as the expected return obtained over an infinite time horizon when the system is controlled using a policy $\pi : \Sigma \rightarrow \mathbf{A}$ in the augmented decision process \hat{M}_D :

$$\mathcal{V}_{\hat{M}_D}^\pi(\sigma) = \mathbb{E} \left[\sum_{t=0}^{\infty} \Gamma^k \hat{r}_t | \sigma_0 = \sigma, \pi \right],$$

where \hat{r}_t is the reward s.t. $\hat{r}_t = \hat{R}(\sigma_t, a_t, \sigma_{t+1})$ and the transition is defined by $\mathbb{P}(\sigma_{t+1} | \sigma_t, a_t) = \hat{T}(\sigma_t, a_t, \sigma_{t+1})$.

A policy π is defined to be better than or equal to a policy π' if its expected return is greater than or equal to that of π' for all states. In an MDP, there is always at least one policy that is better than or equal to all other policies and this is an optimal policy (Sutton & Barto, 1998). In the augmented MDP \hat{M}_D , we denote the optimal policy as $\pi_{D,\phi}$ and we also call it the frequentist-based policy. Let us now decompose the error of using a frequentist-based policy $\pi_{D,\phi}$ in the actual POMDP:

$$\begin{aligned} \mathbb{E}_{D \sim \mathcal{D}_M} \left[V_M^{\pi^*}(\phi_0(H)) - V_M^{\pi_{D,\phi}}(\phi(H)) \right] &= \underbrace{\left(V_M^{\pi^*}(\phi_0(H)) - V_M^{\pi_{D_\infty,\phi}}(\phi(H)) \right)}_{\substack{\text{bias function of dataset } D_\infty \text{ (function of } \pi_s) \\ \text{and frequentist-based policy } \pi_{D_\infty,\phi} \text{ (function of } \phi \text{ and } \Gamma)}} \\ &+ \underbrace{\mathbb{E}_{D \sim \mathcal{D}_M} \left[V_M^{\pi_{D_\infty,\phi}}(\phi(H)) - V_M^{\pi_{D,\phi}}(\phi(H)) \right]}_{\substack{\text{overfitting due to finite dataset } D \text{ (function of } \pi_s, N_l, N_{tr}) \\ \text{in the context of frequentist-based policy } \pi_{D,\phi} \\ \text{(function of } \phi \text{ and } \Gamma)}}}. \end{aligned} \tag{1}$$

The term *bias* actually refers to an asymptotic bias when the size of the dataset tends to infinity, while the term *overfitting* refers to the expected suboptimality due to the finite size of the dataset (and thus due to the variance in the estimated transition function and reward function).

Selecting the feature space $\phi(\mathcal{H})$ carefully allows building a class of policies that have the potential to accurately capture information from data (low bias), but also generalize well (low overfitting). On the one hand, using too many non-informative features will increase overfitting, as stated in Theorem 3 below. On the other hand, a mapping $\phi(\mathcal{H})$ that discards useful available information will suffer an asymptotic bias, as stated in Theorem 1 below (arbitrarily large depending on the POMDP and on the features discarded).

We start by providing a bound on the bias, which is an original result based on the belief states via the ϵ -sufficient statistic.

Theorem 1. "Bound on the bias": Let M be a POMDP described by the γ -tuple $(\mathcal{S}, \mathcal{A}, T, R, \Omega, O, \gamma)$. Let \hat{M}_{D_∞} be an augmented MDP $(\Sigma, \mathcal{A}, \hat{T}, \hat{R}, \Gamma = \gamma)$ estimated, according to Definition 3.1, from a dataset D_∞ . Then, for any ϵ -sufficient mapping $\phi = \phi_\epsilon$, the asymptotic bias can be bounded as follows:

$$\max_{H \in \mathcal{H}} (V_M^{\pi^*}(\phi_0(H)) - V_M^{\pi_{D_\infty, \phi}}(\phi_0(H))) \leq \frac{2\epsilon R_{max}}{(1-\gamma)^3}. \quad (2)$$

Proof. We consider the frequentist-based MDP $\hat{M}_{D_\infty, \phi_0}(\Sigma_0, \mathcal{A}, \hat{T}, \hat{R}, \Gamma = \gamma)$, for $H \in \mathcal{H}$ and $a \in \mathcal{A}$, let us define

$$\mathcal{Q}_{\hat{M}_{D_\infty, \phi_0}}^{\pi_{D_\infty, \phi_0}}(\phi_0(H), a) = \hat{R}'(\phi_0(H), a) + \gamma \sum_{\varphi \in \phi_0(\mathcal{H})} \hat{T}(\phi_0(H), a, \varphi) \mathcal{V}_{\hat{M}_{D_\infty, \phi_0}}^{\pi_{D_\infty, \phi_0}}(\varphi),$$

where the reward

$$\hat{R}'(\phi_0(H), a) = \sum_{\varphi \in \phi_0(\mathcal{H})} \hat{T}(\phi_0(H), a, \varphi) \hat{R}(\phi_0(H), a, \varphi).$$

Then the main part of the proof is to demonstrate Proposition 2 below. From there, by applying Lemma 1 by Abel, Hershkowitz, and Littman (2016), we have:

$$\left\| \mathcal{V}_{M_{D_\infty, \phi_0}}^{\pi_{D_\infty, \phi_0}} - \mathcal{V}_{M_{D_\infty, \phi_0}}^{\pi_{D_\infty, \phi_\epsilon}} \right\|_\infty \leq \frac{2\epsilon R_{max}}{(1-\gamma)^2} = \frac{2\epsilon R_{max}}{(1-\gamma)^3}.$$

By further noticing that, when starting in s_0 , $\hat{M}_{D_\infty, \phi_0}$ and M provide an identical value function for a given policy $\pi_{D, \phi}$ and that $\pi_{D_\infty, \phi_0} \sim \pi^*$, i.e. $V_M^{\pi^*} = V_M^{\pi_{D_\infty, \phi_0}}$, the theorem follows. \square

Remark. As compared to Hutter (2014) and Abel et al. (2016), this bound relates directly to the capacity of the mapping $\phi(H)$ to retrieve sufficient information on the latent hidden state. As compared to PBVI (Pineau, Gordon, & Thrun, 2003) and similar approaches, we do not make the assumption that T , R and O are known and, as such, they need to be estimated from data.

We now provide Proposition 2, which is the key result required in the proof of Theorem 1.

Proposition 2. Let ϕ_ϵ be an ϵ -sufficient mapping, and let ϕ_0 be a sufficient mapping. Then, for any $H^{(1)}, H^{(2)}$ such that $\phi_\epsilon(H^{(1)}) = \phi_\epsilon(H^{(2)})$, we have

$$\max_a \left| \mathcal{Q}_{\hat{M}_{D_\infty, \phi_0}}^{\pi_{D_\infty, \phi_0}}(\phi_0(H^{(1)}), a) - \mathcal{Q}_{\hat{M}_{D_\infty, \phi_0}}^{\pi_{D_\infty, \phi_0}}(\phi_0(H^{(2)}), a) \right| \leq \epsilon \frac{R_{max}}{(1-\gamma)}.$$

Proof. For this proposition, we rely on the fact that since $\phi_\epsilon(H^{(1)}) = \phi_\epsilon(H^{(2)})$, we are able to bound the L1 error terms of the associated belief states of $H^{(1)}$ and $H^{(2)}$. This is illustrated in Figure 2. From that bound, we then present two different ways of independent interest to prove Proposition 2. The details of the proofs are given in Appendix A.1.

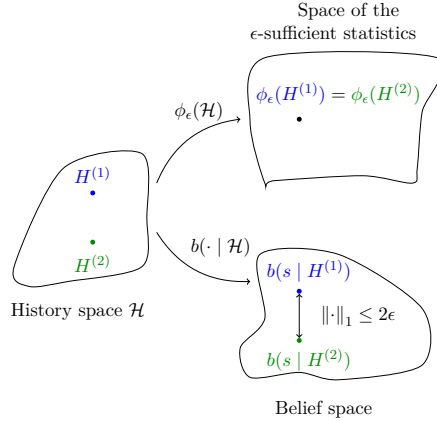


Figure 2: Illustration of the ϕ_ϵ mapping and the belief for $H^{(1)}, H^{(2)} \in \mathcal{H}$: $\phi_\epsilon(H^{(1)}) = \phi_\epsilon(H^{(2)})$.

The first proof makes use of a tree of possible future observations, rewards and corresponding actions given a policy π , and we show that when starting from $H^{(1)}, H^{(2)}$ such that $\phi_\epsilon(H^{(1)}) = \phi_\epsilon(H^{(2)})$, the bound holds.

We also provide an alternative proof that makes use of the formalism of the bisimulation metric (Ferns, Panangaden, & Precup, 2004) along with the data processing inequality. Intuitively, the proof relies on the fact that the histories $H^{(1)}$ and $H^{(2)}$ are close according to a bisimulation metric that measures the behavioral similarity (future rewards and observations). The data processing inequality is used for finding that distance as a function of the L1 error terms of the associated belief states. \square

We now provide a bound on the overfitting error that monotonically grows with $|\phi(\mathcal{H})|$. Theorem 3 shows that using a large set of features allows a larger policy class, hence potentially leading to a stronger drop in performance when the available dataset D is limited (the bound decreases proportionally to $\frac{1}{\sqrt{n}}$). A theoretical analysis in the context of MDPs with a finite dataset was performed by Jiang, Kulesza, and Singh (2015a).

Theorem 3. "Bound on the overfitting": *Let M be a POMDP described by the 7-tuple $(\mathcal{S}, \mathcal{A}, T, R, \Omega, O, \gamma)$. Let \hat{M}_D be an augmented MDP $(\Sigma, \mathcal{A}, \hat{T}, \hat{R}, \Gamma = \gamma)$ estimated, according to Definition 3.1, from a dataset D with the assumption that D has n transitions from any possible pair $(\phi(H), a) \in \phi(\mathcal{H}) \times \mathcal{A}$ (sampled i.i.d according to \mathcal{D}_M). Then the overfitting due to using the frequentist-based policy $\pi_{D, \phi}$ instead of $\pi_{D_\infty, \phi}$ in the POMDP M can be bounded as follows:*

$$\max_{H \in \mathcal{H}} (V_M^{\pi_{D_\infty, \phi}}(\phi(H)) - V_M^{\pi_{D, \phi}}(\phi(H))) \leq \frac{2R_{max}}{(1-\gamma)^2} \sqrt{\frac{1}{2n} \ln \left(\frac{2|\phi(\mathcal{H})||\mathcal{A}|^{1+|\phi(\mathcal{H})|}}{\delta} \right)}, \quad (3)$$

with probability at least $1 - \delta$.

Proof. The proof of Theorem 3 is deferred to Appendix A.2. \square

Overall, Theorems 1 and 3 can help choose a good state representation for POMDPs as they provide bounds on the two terms that appear in the bias-overfitting decomposition of Equation 1. For example, an additional feature in the mapping ϕ has an overall positive effect only if it provides a significant increase of information on the belief state (i.e. if it allows one to obtain a more accurate knowledge of the underlying state of the MDP defined by T and R when given $\phi(H)$). This increase of information must be significant enough to compensate for the additional risk of overfitting when choosing a large cardinality of $\phi(\mathcal{H})$. Note that one could combine the two bounds to theoretically define an optimal choice of the state representation with lower bound guarantees regarding the bias-overfitting tradeoff. In practice, as the two bounds are loose, other techniques described in Section 3.3 are usually more useful.

Related work Partial observability is very common in real world domains and though there has been many works in state abstraction and homomorphisms in the MDP setting (Ravindran & Barto, 2004; Ferns et al., 2004), there has been relatively little in the POMDP setting. One related work in POMDPs by Castro, Panangaden, and Precup (2009) has discussed the parallel between the notion of bisimulation and a notion of trace equivalence, under which states are considered equivalent if they generate the same conditional probability distributions over observation sequences (where the conditioning is on action sequences).

In this paper, we introduce the definition of ϵ -sufficient statistics in POMDPs and derive bounds for the bias based on this property. An insightful part of the proof of Proposition 2 is to use the bisimulation metric (Ferns et al., 2004) and the data processing inequality. Note that the bisimulation metrics may also be used to take into account how the errors on the belief states may have less of an impact at the condition that the hidden states affected by these errors are close according to the bisimulation metric. In case there is some knowledge on the underlying dynamics, one could also extend the notion of the bisimulation metric to allow certain distinct actions to be essentially equivalent (Arun-Kumar, 2006). In that context, Taylor, Precup, and Panagaden (2009) have generalized the notion of the bisimulation metric to a lax bisimulation metric, which may take into account some symmetries and special structures.

3.2 Discussion on Function Approximators

As described earlier, a straightforward mapping $\phi(\cdot)$ may be obtained by discarding some features from the observable history. In addition, the theoretical work developed in this paper can be useful to understand how a specific design of deep neural networks may work well (or not). Indeed, deep neural networks can be seen as a composition of many learnable mappings (constrained by the design of the neural network) and we provide bounds that depend on the property of a given mapping ϕ of the inputs (e.g., the first layers of a deep Q-network). If there is, for instance, an attention mechanism in the first layers of a deep neural network, the mapping made up of those first layers can be analyzed through the bounds developed in our work (e.g., it should not discard important observations that prevent the mapping to be close to some sufficient statistics). As another example, our work helps explain why basic recurrence in a neural network may not be well-suited for POMDPs that require a long history of observations for approaching some sufficient statistics. Indeed, basic recurrent cells are known to have difficulties to convey the long-term dependencies in

sequences (i.e. features from the first time steps in a long time series) and LSTMs (Hochreiter & Schmidhuber, 1997) or other variants should be preferred.

Note that we could add a theorem showing that using, for instance, Rademacher complexity would allow providing other bounds (potentially tighter) than Theorem 3. However, such a bound is usually of little interest in practice, specifically concerning deep learning, because the bounds based on complexity measures fail to provide insights on the generalization capabilities of neural networks (Zhang, Bengio, Hardt, Recht, & Vinyals, 2016). Indeed, it has been empirically demonstrated that deep neural networks have strong generalization capabilities even with a high number of parameters (hence a potentially large complexity). In other words, a very loose bound on the overfitting error due to a large number of parameters may not lead to a performance drop in practice.

3.3 Selection of the Parameters with Validation or Cross-Validation to Balance the Bias-Overfitting Tradeoff

In the batch setting case, the selection of the policy parameters to effectively balance the bias-overfitting tradeoff can be done similarly to that in supervised learning (e.g., cross-validation) as long as the performance criterion can be estimated from a subset of the trajectories from the dataset D not used during training (validation set). One possibility is to fit an MDP model from data via the frequentist approach (or regression), and evaluate the policy against the model (with Monte-Carlo simulations). Another approach is to use the idea of importance sampling (Precup, 2000). A mix of the two approaches has also been developed (Jiang & Li, 2016; Thomas & Brunskill, 2016).

Importance of the discount factor Γ used in the training phase: Besides the elements related to feature selection and function approximators, artificially lowering the discount factor can also be used to improve the performance of the policy when solving MDPs with limited data (Petrik & Scherrer, 2009; Jiang, Kulesza, Singh, & Lewis, 2015b). In the partially observable setting, these results can be transferred to the frequentist-based MDP $(\Sigma, A, \hat{T}, \hat{R}, \Gamma)$ by choosing $\Gamma < \gamma$, which introduces a bias but reduces the risk of overfitting.

4. Experiments

This section provides empirical illustrations of the theoretical results both on synthetic POMDPs and on a large-scale POMDP in the context of smartgrids, with real-world data. On the synthetic POMDPs, we first illustrate the main theoretical results of this paper related to the state representation of POMDPs and we also illustrate the use of function approximators and the training discount factor Γ . On the large-scale POMDP, we illustrate that an efficient feature selection process can be useful, even when used in addition to function approximators and a biased discount factor.

4.1 Synthetic POMDPs

In order to be representative of a diversity of environments, we provide results on a distribution of POMDPs.

4.1.1 PROTOCOL

We randomly sample N_P POMDPs such that $N_S = 5$, $N_A = 2$ and $N_\Omega = 5$ (except when stated otherwise) from a distribution \mathcal{P} that we refer to as Random POMDP. The distribution \mathcal{P} is fully determined by specifying a distribution over the set of possible transition functions $T(\cdot, \cdot, \cdot)$, a distribution over the set of reward functions $R(\cdot, \cdot, \cdot)$, and a distribution over the set of possible conditional observation probabilities $O(\cdot, \cdot)$.

Random transition functions $T(\cdot, \cdot, \cdot)$ are drawn by assigning, for each entry (s, a, s') , a zero value with probability $3/4$, and, with probability $1/4$, a non-zero entry with a value drawn uniformly in $[0, 1]$. For all (s, a) , if all $T(s, a, s')$ are zeros, we enforce one non-zero value for a random $s' \in \mathcal{S}$. Values are normalized. Random reward functions are generated by associating to all possible (s, a, s') a reward sampled uniformly and independently from $[-1, 1]$. Random conditional observation probabilities $O(\cdot, \cdot)$ are generated in the following way: the probability of observing $o^{(i)}$ when in state $s^{(i)}$ is equal to 0.5, while all other values are chosen uniformly randomly so that it is normalized for any s . For all POMDPs, we have $\gamma = 1$ and $\Gamma = 0.95$ if not stated otherwise and we truncate the trajectories to a length of $N_l = 100$ time steps.

For each generated POMDP $P \sim \mathcal{P}$, we generate 20 datasets $D \in \mathcal{D}_P$ where \mathcal{D}_P is a probability distribution over all possible sets of n trajectories ($n \in [2, 5000]$); where each trajectory is made up of a history H_{100} of 100 time steps, when starting from an initial state $s_0 \in \mathcal{S}$ and taking uniform random decisions. Each dataset D induces a policy $\pi_{D, \phi}$, and we want to evaluate the expected return of this policy while discarding the variance related to the stochasticity of the transitions, observations and rewards. To do so, policies are tested with 1000 rollouts of the policy. For each POMDP P , we are then able to get an estimate of the average score μ_P which is defined as:

$$\mu_P = \mathbb{E}_{D \sim \mathcal{D}_P} \mathbb{E} \left[\sum_{t=0}^{N_l} \gamma^t r_t | s_0, \pi_{D, \phi} \right].$$

We are also able to get an estimate of a parametric variance σ_P^2 defined as:

$$\sigma_P^2 = \text{Var}_{D \sim \mathcal{D}_P} \mathbb{E} \left[\sum_{t=0}^{N_l} \gamma^t r_t | s_0, \pi_{D, \phi} \right].$$

4.1.2 HISTORY PROCESSING

We show experimentally that any additional feature from the history H_t is likely to reduce the asymptotic bias, but may also increase the overfitting. For any history length h , we consider the mapping ϕ_h that extracts the current observation and the last $h-1$ (observation, action) tuples. In the experiments, we compare the policies π_{D, ϕ_h} for $h = 1, 2, 3$.

The values $\mathbb{E}_{P \sim \mathcal{P}} \mu_P$ and $\mathbb{E}_{P \sim \mathcal{P}} \sigma_P$ are displayed in Figure 3. One can observe that a small set of features (small history) appears to be a better choice (in terms of total bias) when the dataset is small (only a few trajectories). With an increasing number of trajectories, the optimal choice in terms of number of features ($h = 1, 2$ or 3) also increases. In addition, one can also observe that the expected variance of the score decreases as the number of samples




Figure 3: Evolution (as a function of the size of the dataset) of estimated values of $\mathbb{E}_{P \sim \mathcal{P}} \mu_P \pm \mathbb{E}_{P \sim \mathcal{P}} \sigma_P$ computed from a sample of $N_P = 50$ POMDPs drawn from \mathcal{P} . The bars are used to represent the variance observed when dealing with different datasets drawn from a distribution; note that this is not a usual error bar.

increases. As the variance decreases, the risk of overfitting also decreases, and it becomes possible to target a larger policy class (using a larger feature set).

The overfitting error is also linked to the variance of the value function estimates in the frequentist-based MDP. When these estimates have a large variance, an overfitting term appears because of a higher chance of picking one of the suboptimal policies, as illustrated in Figure 4.

4.1.3 FUNCTION APPROXIMATOR AND DISCOUNT FACTOR

We also illustrate the effect of using function approximators on the bias-overfitting tradeoff. To do so, we process the output of the state representation $\phi(\cdot)$ into a deep Q-learning scheme (technical details are given in Appendix A.3). We can see in Figure 5 that deep Q-learning policies suffer less overfitting as compared to the frequentist-based approach (lower degradation of performance in the low-data regime) even though using a large set of features still leads to more overfitting than a small set of features. We also see that deep Q-learning policies do not introduce an important asymptotic bias (identical performance when a lot of data is available) because the neural network architecture is rich enough. Note that the variance is slightly larger than in Figure 3, and does not vanish to 0 with additional data. This is due to the additional stochasticity induced when building the Q-value function with neural networks (note that when performing the same experiments while taking the average recommendation of several Q-value functions, this variance decreases with the number of Q-value functions).

Finally, we empirically illustrate in Figure 6 the effect of modifying the discount factor Γ . When the training discount factor is lower than the one used in the actual POMDP ($\Gamma < \gamma$), there is an additional bias term, while when a high discount factor is used (Γ close to 1)




Figure 4: Illustrations of the return estimates in the augmented MDP \hat{M}_D ($h = 2$) for three different policies. The policies are selected specifically for illustration purposes based on the criterion $V_M^{\pi D}$; the best performing (in green), the worst performing (in red) and the median performing were selected in a set of 50 policies built when $D \sim \mathcal{D}$ has 5 trajectories of data from the actual POMDP. On the actual POMDP, the expected returns are $V_M^{\pi_1} = 28$, $V_M^{\pi_2} = 33$, $V_M^{\pi_3} = 42$ (in general, these values need not be the same as the expected value of the probability distribution in the two graphs).

with a limited amount of data, overfitting increases. In our experiments, the influence of the discount factor is more subtle as compared to the impact of the state representation and the function approximator. The influence is nonetheless clear: it is better to have a low discount factor when only a few data points are available, and it is better to have a high discount factor when a lot of data is available, which is in line with previous analyses for MDPs (Jiang et al., 2015b).

4.2 Real-world Application in Smartgrids

The microgrid model considered here was introduced by François-Lavet, Taralla, Ernst, and Fonteneau (2016) and we next describe the main elements ⁴.

4.2.1 MICROGRID BENCHMARK

The microgrid is powered by photovoltaic (PV) panels combined with both long-term storage (hydrogen-based) and short-term storage (such as, for instance, LiFePO₄ batteries). These two types of storage aim at fulfilling, at best, the demand by addressing the seasonal and daily fluctuations of solar irradiance. Distinguishing short- from long-term storage is mainly a question of cost: batteries are too expensive to be used for addressing seasonal variations.

4. The source code is available at <https://github.com/VinF/deer/>.




Figure 5: Evolution (as a function of the size of the dataset) of estimated values of $\mathbb{E}_{P \sim \mathcal{P}} \mu_P \pm \mathbb{E}_{P \sim \mathcal{P}} \sigma_P$ computed from a sample of $N_P = 50$ POMDPs drawn from \mathcal{P} (same as Figure 3) with neural network as a function approximator. The bars are used to represent the variance observed when dealing with different datasets drawn from a distribution; this is not a usual error bar.

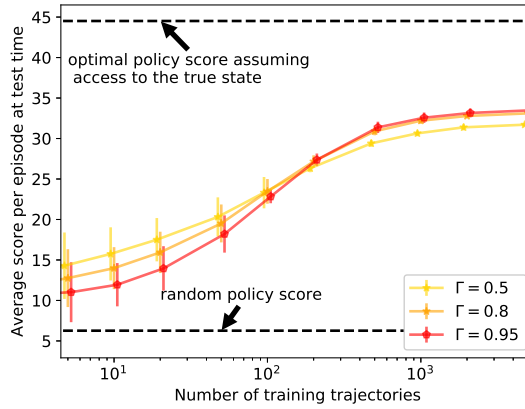


Figure 6: Evolution in the frequentist-based case (as a function of the size of the dataset) of estimated values of $\mathbb{E}_{P \sim \mathcal{P}} \mu_P \pm \mathbb{E}_{P \sim \mathcal{P}} \sigma_P$ computed from a sample of $N_P = 10$ POMDPs drawn from \mathcal{P} with $N_S = 8$ and $N_\Omega = 8$ ($h = 3$). The bars are used to represent the variance observed when dealing with different datasets drawn from a distribution; this is not a usual error bar.

Operating the microgrid is formalized as a POMDP over discrete time steps of one hour. The observations at each time step are made up of the consumption (c_t), the production

(ψ_t) and the level of energy in the battery (s_t^B) . The mapping $\phi(\cdot)$ is defined such that

$$\phi(H_t) = \left[[c_{t-h^c}, \dots, c_{t-1}], [\psi_{t-h^p}, \dots, \psi_{t-1}], s_t^B \right],$$

where $h^c = h^p$ are the lengths of the time series considered for the consumption and production, respectively.

The instantaneous reward signal r_t is obtained by adding the revenues generated by the hydrogen production r^{H_2} and the penalties r^- due to the value of loss load:

$$r_t = r(a_t, d_t) = r^{H_2}(a_t) + r^-(a_t, d_t), \quad (4)$$

where d_t denotes the net electricity demand, which is the difference between the local consumption c_t and the local production of electricity ψ_t . The penalty r^- is proportional to the total amount of energy that was not supplied to meet the demand, with the cost incurred per kilowatt-hour (kWh) not supplied within the microgrid set to 2 €/kWh (corresponding to a value of loss load). The revenues (or cost) generated by the hydrogen production r^{H_2} is proportional to the total amount of energy transformed (or consumed) in the form of hydrogen, where the revenue (or cost) per kWh of hydrogen produced (or used) is set to 0.1 €/kWh .

A synthetic residential consumption profile is considered with a daily consumption of $18kWh$. The PV production profile comes from actual data of a residential customer located in Belgium (average production changes by a factor of about 1:5 between summer and winter). A fixed sizing of the microgrid is considered. The size of the battery is $x^B = 15kWh$, the instantaneous power of the hydrogen storage is $x^{H_2} = 1.1kW$ and the peak power generation of the PV installation is $x^{PV} = 12kW_p$ (kW_p stands for kilowatt-peak, which is the maximum electric power that can be supplied by the PV panels).

The action a_t is an element of the action space made up of three discrete actions: $\mathcal{A} = \{a^{(0)}, a^{(1)}, a^{(2)}\}$. The possible actions relate to whether the microgrid creates hydrogen from electricity, creates electricity from hydrogen or leaves the long-term storage unused. The battery adapts to store excess electrical power (except when the battery is full) while avoiding any loss load (except when the battery is empty, which leads to negative rewards).

Note that in this specific application, using a history of observations from the previous time steps provides information on the latent features (of the state of the POMDP) such as the time of the day, the weather, the season (PV production and consumption time series are conditionally dependent on these latent features).

Also note that the dynamics of the system depend on exogenous time series (production and consumption time series) for which the agent only has finite data. In order to evaluate the capacity of the agent to generalize, we can break down the exogenous time series into a part that will be used for training and one part that will be used for validation (François-Lavet et al., 2016). The final performance is then evaluated on the environment with unseen time series.

4.2.2 SPLITTING THE TIMES SERIES TO AVOID OVERFITTING

In this microgrid domain, the agent is provided with up to two years of actual past realizations of the consumption (c_t) and the production (ψ_t) . These past realizations are split into a training environment (one year) and a validation environment (one year). The training

environment is used to train the policy, while the validation environment is used at each epoch ⁵ to estimate how well the policy performs on the undiscounted sum of rewards, and it selects the best (approximated) Q-network denoted \tilde{Q}^* before overfitting (by selecting a discount factor lower than the maximum and by using early stopping). It also has the advantage of picking up the Q-network at an epoch less affected by instabilities. The selected trained Q-network is then used in a test environment ($y = 3$) to provide an independent estimation of how well the resulting policy performs. Technical details relative to the deep Q-network algorithm used are given in Appendix B.1.

To empirically demonstrate the effect of the bias-overfitting tradeoff, we artificially reduce the diversity of the training and validation time-series by a factor $\kappa = \{1, 2, 4, 8, 16\}$. Overall, the time series for training and validation should still be as close to 365 days as possible (the epochs are run on approximately the same given number of time steps for all cases). In addition, the true underlying processes should be respected as closely as possible by (i) guaranteeing that the succession of the seasons are not corrupted and (ii) guaranteeing that consecutive days in the original time series should be kept consecutive whenever possible.

In order to do so, the time series of one year are divided into four seasons. Each season is then split into κ blocs of the same number of days. For each season, data reduction is done by replicating one of the blocs instead on all remaining blocs of the same season. This artificial reduction on the available data is illustrated on Figure 7.




Figure 7: Illustration of how the quantity of data in a time series of one year is artificially reduced by a factor $\kappa = 2$ while guaranteeing that the specificity of the seasons are essentially preserved. This schema applies for both the consumption and the production time-series and for both the training time series and validation time series.

4.2.3 RESULTS OF THE EXPERIMENTS

The purpose of the experiments is to illustrate the theoretical results, and are given in Figure 8. On the one hand, using a long history may pose problems relative to overfitting

5. An epoch is defined as the set of all iterations required to go through the whole year of the exogenous time series c_t and ψ_t . Each iteration is made up of a transition of one time-step in the environment as well as a gradient step of all parameters θ of the Q-network.

when the amount of training data is limited. On the other hand, using a long history of observations allows achieving a good operation when enough data is available (information on the latent features are preserved).

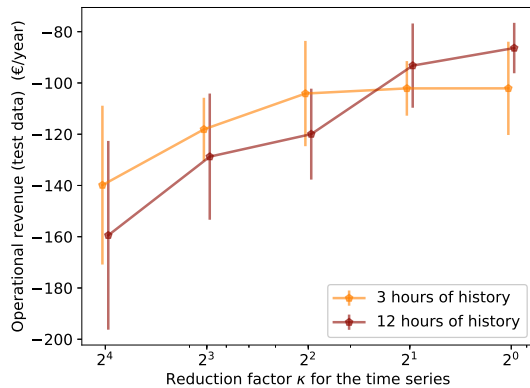


Figure 8: Evolution of the performance of the policy in the test setting (one year of unseen past time series for both the consumption and production). Using a relatively short history of observations is beneficial when the amount of data is limited while a longer history of observations leads to improved results when the amount of data increases (lower κ means that more data is available). For each point, the mean and the standard deviation are calculated over 2^3 different runs of the Deep Q-network algorithm (over different ways to split the time series defined in appendix and for different random seeds).

As already noted, there is a close link between (i) choosing a function approximator and (ii) selecting features, since the function approximator characterizes how the features will be treated into higher levels of abstraction (*a fortiori* it can thus give more or less weight to some features). Notwithstanding, this example shows that, in general, the use of function approximators may still benefit from an efficient feature selection process. Indeed, we still observe a tradeoff between bias and overfitting in this microgrid application, even with the convolutional architecture of the deep Q-network.

5. Conclusion and Future Works

This paper discusses the bias-overfitting tradeoff of batch RL algorithms in the context of POMDPs. Most current RL works compare algorithms in an online setting, where the performance depends on many factors such as improved exploration, better generalization, or better numerical efficiency of the gradient step (in the case of deep RL). In contrast, this paper is focused on one source of sub-optimality related to generalization from limited data in the case of POMDPs. In that context, we propose an analysis showing that, similarly to supervised learning techniques, RL may face a bias-overfitting dilemma in situations where the policy class is too large compared to the batch of data. In such situations, we show

that it may be preferable to concede an asymptotic bias in order to reduce overfitting. This (favorable) asymptotic bias may be introduced through different manners: (i) downsizing the state representation, (ii) using specific types of function approximators and (iii) lowering the discount factor.

The main theoretical results of this paper relate to the state representation; the originality of the setting proposed in this paper compared to Maillard, Ryabko, and Munos (2011), Ortner, Maillard, and Ryabko (2014) and the related work is mainly to formalize the problem in a batch setting (limited set of tuples) instead of the online setting, where they investigate the E/E dilemma. As compared to Jiang et al. (2015b), the originality is to consider a partially observable setting. We introduce the notion of the ϵ -sufficient statistics and showed that it allows us to formalize the intuition of the bias-overfitting trade-off in a rigorous way and provides new insights compared to the MDP case. In particular, the bound of Theorem 1 is a new result based on L1 error terms of the associated belief states. There are also interesting insights from the techniques used to derive that bound, where we use the formalism of bisimulation metrics and the data processing inequality.

The work proposed in this paper may also be of interest in online settings because, at each stage, obtaining a performant policy from given data is part of the solution to an efficient exploration/exploitation tradeoff. For instance, optimizing the bias-overfitting tradeoff suggests that it can be beneficial to dynamically adapt the feature space and the function approximator. This can be done through ad hoc regularization or by adapting the neural network architecture, using for instance the NET2NET transformation (Chen, Goodfellow, & Shlens, 2015).

Acknowledgments

We thank the anonymous reviewers for many helpful comments. The authors thank the Walloon Region (Belgium) that funded this research in the context of the BATWAL project. We also acknowledge financial support by the Natural Science and Engineering Research Council Canada (NSERC) and Samsung Electronics Co Ltd. Guillaume Rabusseau acknowledges support of an IVADO postdoctoral fellowship.

Appendix A. Proofs

A.1 Proof of Proposition 2

Proposition 2 Let ϕ_ϵ be an ϵ -sufficient mapping, and let ϕ_0 be a sufficient mapping. Then, for any $H^{(1)}, H^{(2)}$ such that $\phi_\epsilon(H^{(1)}) = \phi_\epsilon(H^{(2)})$, we have

$$\max_a \left| \mathcal{Q}_{\hat{M}_{D_\infty, \phi_0}}^{\pi_{D_\infty, \phi_0}}(\phi_0(H^{(1)}), a) - \mathcal{Q}_{\hat{M}_{D_\infty, \phi_0}}^{\pi_{D_\infty, \phi_0}}(\phi_0(H^{(2)}), a) \right| \leq \epsilon \frac{R_{max}}{(1 - \gamma)}.$$

Proof. We first define the notion of a policy conditioned on a previous history. Let $\pi : \phi_0(\mathcal{H}) \rightarrow \mathcal{A}$. Given a history $H \in \mathcal{H}$ we define $\pi_H : (\mathcal{A} \times \Omega)^* \rightarrow \mathcal{A}$ by $\pi_H(a_1\omega_1 \cdots a_k\omega_k) = \pi(\phi_0(Ha_1\omega_1 \cdots a_k\omega_k))$, i.e. π_H is the policy π conditioned on having previously observed the history H ⁶.

6. To simplify notations, one can consider that the rewards are included in the observations ω .

Given a history $H \in \mathcal{H}$ we can then define the Q-value of a state $s \in \mathcal{S}$ and an action $a \in \mathcal{A}$ in the underlying MDP following π and assuming that H has been observed:

$$\begin{aligned} \mathcal{Q}^{\pi H}(s, a) &= R'(s, a) \\ &\quad + \gamma \mathbb{E}_{s_1|s, a} \mathbb{E}_{\omega_1|s_1} R'(s_1, a_1 = \pi_H(a\omega_1)) \\ &\quad + \gamma^2 \mathbb{E}_{s_1|s, a} \mathbb{E}_{\omega_1|s_1} \mathbb{E}_{s_2|s_1, a_1} \mathbb{E}_{\omega_2|s_2} R'(s_2, \pi_H(a\omega_1 a_1 \omega_2)) \\ &\quad + \dots \end{aligned}$$

where $R'(s, a) = \mathbb{E}_{s'|s, a} R(s, a, s')$. Note that $\mathcal{Q}_{\hat{M}_{D_\infty, \phi_0}}^\pi(\phi_0(H), a) = \mathbb{E}_{s \sim b_{\phi_0}(\cdot | \phi_0(H))} \mathcal{Q}^{\pi H}(s, a)$ for any history $H \in \mathcal{H}$ and any action $a \in \mathcal{A}$ since ϕ_0 is a sufficient mapping.

We can now prove the proposition. Let $a \in \mathcal{A}$. Let $\pi^* = \pi_{D_\infty, \phi_0}$ be the optimal policy in \hat{M} and let $\pi_i = \pi_{H_i}^*$ be the optimal policy conditioned on having observed H_i for $i = 1, 2$. Since π^* is optimal we have

$$\begin{aligned} \mathcal{Q}^{\pi^*}(\phi_0(H^{(1)}), a) &= \mathbb{E}_{s \sim b_{\phi_0}(\cdot | \phi_0(H_1))} \mathcal{Q}^{\pi_1}(s, a) \\ &\geq \mathbb{E}_{s \sim b_{\phi_0}(\cdot | \phi_0(H_1))} \mathcal{Q}^{\pi_2}(s, a) \end{aligned}$$

and

$$\begin{aligned} \mathcal{Q}^{\pi^*}(\phi_0(H^{(2)}), a) &= \mathbb{E}_{s \sim b_{\phi_0}(\cdot | \phi_0(H_2))} \mathcal{Q}^{\pi_2}(s, a) \\ &\geq \mathbb{E}_{s \sim b_{\phi_0}(\cdot | \phi_0(H_2))} \mathcal{Q}^{\pi_1}(s, a). \end{aligned}$$

We then have

$$\begin{aligned} \mathcal{Q}_{\hat{M}_{D_\infty, \phi_0}}^{\pi^*}(\phi_0(H^{(1)}), a) - \mathcal{Q}_{\hat{M}_{D_\infty, \phi_0}}^{\pi^*}(\phi_0(H^{(2)}), a) &= \mathbb{E}_{s \sim b_{\phi_0}(\cdot | \phi_0(H_1))} \mathcal{Q}^{\pi_1}(s, a) - \mathbb{E}_{s \sim b_{\phi_0}(\cdot | \phi_0(H_2))} \mathcal{Q}^{\pi_2}(s, a) \\ &\leq \mathbb{E}_{s \sim b_{\phi_0}(\cdot | \phi_0(H_1))} \mathcal{Q}^{\pi_1}(s, a) - \mathbb{E}_{s \sim b_{\phi_0}(\cdot | \phi_0(H_2))} \mathcal{Q}^{\pi_1}(s, a) \\ &= \sum_{s \in \mathcal{S}} \left(b(s|H^{(1)}) - b(s|H^{(2)}) \right) \mathcal{Q}^{\pi_1}(s, a) \\ &\leq \frac{\epsilon}{1-\gamma} R_{max}, \end{aligned}$$

where we used the facts that $b(\cdot | H^{(1)}) - b(\cdot | H^{(2)})$ is a vector of L_1 -norm less than 2ϵ (by Lemma 4) whose components sum to zero and $|\mathcal{Q}^{\pi_1}(s, a) - \mathcal{Q}^{\pi_1}(s', a)| \leq \frac{R_{max}}{1-\gamma}$ for any $s, s' \in \mathcal{S}$.

Applying the same argument *mutatis mutandis* we obtain $\mathcal{Q}_{\hat{M}_{D_\infty, \phi_0}}^{\pi^*}(\phi_0(H^{(2)}), a) - \mathcal{Q}_{\hat{M}_{D_\infty, \phi_0}}^{\pi^*}(\phi_0(H^{(1)}), a) \leq \frac{\epsilon}{1-\gamma} R_{max}$ from which the result follows. \square

Alternative proof. An alternative proof of independent interest makes use of the formalism of the bisimulation metric (Ferns et al., 2004) along with the data processing inequality. Let us consider $d \in \mathcal{M}$, where \mathcal{M} is the set of all semi-metrics on Σ_0 with distance at most 1. We fix a particular d as follows: $\forall \sigma^{(1)}, \sigma^{(2)} \in \Sigma_0$:

$$\begin{aligned} d(\sigma^{(1)}, \sigma^{(2)}) &= \max_a \left(c_R \left| \hat{R}'(\sigma^{(1)}, a) - \hat{R}'(\sigma^{(2)}, a) \right| \right. \\ &\quad \left. + c_T d_P(\hat{T}(\sigma^{(1)}, a, \cdot), \hat{T}(\sigma^{(2)}, a, \cdot)) \right), \end{aligned}$$

where d_P is some probability metric, and where $c_R, c_T \geq 0$ are such that $c_R + c_T \leq 1$. We define $F : \mathcal{M} \rightarrow \mathcal{M}$ by

$$F(d)(\sigma^{(1)}, \sigma^{(2)}) = \max_a \left(c_R \left| \hat{R}'(\sigma^{(1)}, a) - \hat{R}'(\sigma^{(2)}, a) \right| + c_T T_K(d)(\hat{T}(\sigma^{(1)}, a, \cdot), \hat{T}(\sigma^{(2)}, a, \cdot)) \right).$$

where $T_K(d)$ is the Kantorovich metric induced by d . From Lemma 5, F has a least fixed-point d_{fix} , and it is a bisimulation metric. From Lemmas 6, 7 (using the data processing inequality), we have $\forall \sigma^{(1)}, \sigma^{(2)} \in \Sigma_0$ that

$$d_{fix}(\sigma^{(1)}, \sigma^{(2)}) \leq \frac{1}{2} \|b(\cdot | \sigma^{(1)}) - b(\cdot | \sigma^{(2)})\|_1 R_{max}.$$

Using Lemma 8, it follows that

$$\max_a \left| \mathcal{Q}_{\hat{M}_{D_\infty, \phi_0}}^{\pi_{D_\infty, \phi_0}}(\phi_0(H^{(1)}), a) - \mathcal{Q}_{\hat{M}_{D_\infty, \phi_0}}^{\pi_{D_\infty, \phi_0}}(\phi_0(H^{(2)}), a) \right| \leq \epsilon \frac{R_{max}}{(1 - \gamma)}.$$

□

Lemma 4. *Let $H^{(1)}, H^{(2)} \in H$ be two histories such that $\phi_\epsilon(H^{(1)}) = \phi_\epsilon(H^{(2)})$. Then,*

$$\|b(\cdot | H^{(1)}) - b(\cdot | H^{(2)})\|_1 \leq 2\epsilon.$$

Proof.

$$\begin{aligned} \|b(\cdot | H^{(1)}) - b(\cdot | H^{(2)})\|_1 &= \|b(\cdot | H^{(1)}) - b_{\phi_\epsilon}(\cdot | \phi_\epsilon(H^{(1)})) \\ &\quad + b_{\phi_\epsilon}(\cdot | \phi_\epsilon(H^{(2)})) - b(\cdot | H^{(2)})\|_1 \\ &\leq \|b(\cdot | H^{(1)}) - b_{\phi_\epsilon}(\cdot | \phi_\epsilon(H^{(1)}))\|_1 \\ &\quad + \|b_{\phi_\epsilon}(\cdot | \phi_\epsilon(H^{(2)})) - b(\cdot | H^{(2)})\|_1 \\ &\leq 2\epsilon, \end{aligned}$$

□

Lemma 5. *Let $c_R, c_T \geq 0$ and $c_R + c_T \leq 1$. Define $F : \mathcal{M} \rightarrow \mathcal{M}$ by*

$$F(d)(\sigma^{(1)}, \sigma^{(2)}) = \max_a \left(c_R \left| \hat{R}'(\sigma^{(1)}, a) - \hat{R}'(\sigma^{(2)}, a) \right| + c_T T_K(d)(\hat{T}(\sigma^{(1)}, a, \cdot), \hat{T}(\sigma^{(2)}, a, \cdot)) \right).$$

Then F has a least fixed-point, d_{fix} , and d_{fix} is a bisimulation metric.

Proof. See proof of Theorem 4.5 by Ferns et al. (2004). □

Lemma 6. *Let $H^{(1)}, H^{(2)} \in \mathcal{H}$ such that*

$$\|b(\cdot | \phi_0(H^{(1)})) - b(\cdot | \phi_0(H^{(2)}))\|_1 \leq 2\epsilon.$$

Then

$$\max_a \left| \hat{R}'(\phi_0(H^{(1)}), a) - \hat{R}'(\phi_0(H^{(2)}), a) \right| \leq \epsilon R_{max}.$$

Proof. This follows directly from the fact that $\forall H$:

$$\hat{R}'(\phi_0(H), a) = \sum_{s \in \mathcal{S}} \mathbb{P}(s | \phi_0(H)) R'(s, a),$$

where $R'(s, a) = \mathbb{E}_{s' | s, a} R(s, a, s')$. □

Lemma 7. *Let $H^{(1)}, H^{(2)} \in \mathcal{H}$ such that*

$$\left\| b(\cdot | \phi_0(H^{(1)})) - b(\cdot | \phi_0(H^{(2)})) \right\|_1 \leq 2\epsilon.$$

Then

$$\sum_{\omega \in \Omega} \left\| \mathbb{P}(\omega | \phi_0(H^{(1)}), a) b(\cdot | \phi_0(H^{(1)}), a, \omega) - \mathbb{P}(\omega | \phi_0(H^{(2)}), a) b(\cdot | \phi_0(H^{(2)}), a, \omega) \right\|_1 \leq 2\epsilon.$$

Proof. For any $a \in \mathcal{A}$ by direct application of the data processing inequality (DPI) in the special case of the total variation:

$$\left\| b(\cdot | \phi_0(H^{(1)}), a) - b(\cdot | \phi_0(H^{(2)}), a) \right\|_1 \leq 2\epsilon.$$

Moreover, as illustrated on Figure 9 and once more by application of the DPI in the context of the special case of the total variation, we have that

$$\frac{1}{2} \mathbb{E}_{Q_Y} \left| \frac{P_Y}{Q_Y} - 1 \right| \leq \frac{1}{2} \mathbb{E}_{Q_X} \left| \frac{P_X}{Q_X} - 1 \right|.$$

with the distributions

- $P_X = \mathbb{P}(s | \phi_0(H^{(1)}), a)$,
- $Q_X = \mathbb{P}(s | \phi_0(H^{(2)}), a)$,
- $P_Y = \mathbb{P}(\omega, s | \phi_0(H^{(1)}), a)$, and
- $Q_Y = \mathbb{P}(\omega, s | \phi_0(H^{(2)}), a)$.

Note that the output distributions is on the tuple (ω, s) .

Since the right hand-side in the aforementioned DPI can be bounded by ϵ :

$$\frac{1}{2} \sum_{s \in \mathcal{S}} \mathbb{P}(s | \phi_0(H^{(2)}), a) \left| \frac{\mathbb{P}(s | \phi_0(H^{(1)}), a)}{\mathbb{P}(s | \phi_0(H^{(2)}), a)} - 1 \right| \leq \epsilon,$$

we have

$$\frac{1}{2} \sum_{\omega \in \Omega} \sum_{s \in \mathcal{S}} \left| \mathbb{P}(\omega | \phi_0(H^{(1)}), a) \mathbb{P}(s | \phi_0(H^{(1)}), a, \omega) - \mathbb{P}(\omega | \phi_0(H^{(2)}), a) \mathbb{P}(s | \phi_0(H^{(2)}), a, \omega) \right| \leq \epsilon.$$




Figure 9: Illustration of the two histories $H^{(1)}$ and $H^{(2)}$. Here we illustrate only for one given action a and two possible observations ω_1 and ω_2 . The dashed line illustrates the key relation between $H^{(1)}$ and $H^{(2)}$, which is $\left\| b(\cdot | H^{(1)}) - b(\cdot | H^{(2)}) \right\|_1 \leq 2\epsilon$. In expectation, this relation is preserved through the action and observations.

This can be rewritten as:

$$\sum_{\omega \in \Omega} \left\| \mathbb{P}(\omega | \phi_0(H^{(1)}), a) b(\cdot | \phi_0(H^{(1)}), a, \omega) - \mathbb{P}(\omega | \phi_0(H^{(2)}), a) b(\cdot | \phi_0(H^{(2)}), a, \omega) \right\|_1 \leq 2\epsilon.$$

□

Lemma 8. *Suppose $\gamma \leq c_T$. Then $\forall \sigma^{(1)}, \sigma^{(2)} \in \Sigma_0$:*

$$c_R \max_a \left| \mathcal{Q}^*(\sigma^{(1)}, a) - \mathcal{Q}^*(\sigma^{(2)}, a) \right| \leq d_{fix}(\sigma^{(1)}, \sigma^{(2)})$$

where \mathcal{Q}^* stands for $\mathcal{Q}_{\hat{M}_{D_\infty, \phi_0}}^{\pi_{D_\infty, \phi_0}}$.

Proof. The proof is done using the same techniques than the proof of Theorem 5.1 by Ferns et al. (2004). Let us first prove that

$$\max_a c_R \left| \mathcal{Q}_n(\sigma^{(1)}, a) - \mathcal{Q}_n(\sigma^{(2)}, a) \right| \leq d_n(\sigma^{(1)}, \sigma^{(2)})$$

where $F(d_n(\sigma^{(1)}, \sigma^{(2)})) = d_{n+1}(\sigma^{(1)}, \sigma^{(2)})$, and

$$\mathcal{Q}_{n+1}(\phi_0(H), a) = \hat{R}'(\phi_0(H), a) + \gamma \sum_{\varphi \in \phi_0(\mathcal{H})} \hat{T}(\phi_0(H), a, \varphi) \mathcal{V}_n(\varphi),$$

with $\mathcal{V}_n(\varphi) = \max_{a'} \mathcal{Q}_n(\varphi, a')$. The result would then follow by taking the limit.

$\forall H^{(1)}, H^{(2)} \in \mathcal{H}$: $\phi_\epsilon(H^{(1)}) = \phi_\epsilon(H^{(2)})$, it follows that

$$\begin{aligned}
 & c_R \max_a \left| \mathcal{Q}_{n+1} \left(\phi_0(H^{(1)}), a \right) - \mathcal{Q}_{n+1} \left(\phi_0(H^{(2)}), a \right) \right| \\
 & \leq c_R \max_a \left(\left| \hat{R}' \left(\phi_0(H^{(1)}), a \right) - \hat{R}' \left(\phi_0(H^{(2)}), a \right) \right| \right. \\
 & \quad \left. + \gamma \left| \hat{T}(\phi_0(H^{(1)}), a, \cdot) \mathcal{V}_n(\cdot) - \hat{T}(\phi_0(H^{(2)}), a, \cdot) \mathcal{V}_n(\cdot) \right| \right) \\
 & \leq \max_a \left(c_R \left| \hat{R}' \left(\phi_0(H^{(1)}), a \right) - \hat{R}' \left(\phi_0(H^{(2)}), a \right) \right| \right. \\
 & \quad \left. + c_T \left| \left(\hat{T}(\phi_0(H^{(1)}), a, \cdot) - \hat{T}(\phi_0(H^{(2)}), a, \cdot) \right) \frac{c_R \gamma}{c_T} \mathcal{V}_n(\cdot) \right| \right) \\
 & \leq \max_a \left(c_R \left| \hat{R}' \left(\phi_0(H^{(1)}), a \right) - \hat{R}' \left(\phi_0(H^{(2)}), a \right) \right| \right. \\
 & \quad \left. + c_T T_K(d_n) \left(\hat{T}(\phi_0(H^{(1)}), a, \cdot) - \hat{T}(\phi_0(H^{(2)}), a, \cdot) \right) \right) \\
 & = F \left(d_n(\phi_0(H^{(1)}), \phi_0(H^{(2)})) \right) \\
 & = d_{n+1} \left(\phi_0(H^{(1)}), \phi_0(H^{(2)}) \right),
 \end{aligned}$$

where we have used the fact that $\left\{ \frac{c_R \gamma}{c_T} \mathcal{V}_n(u) : u \in \Sigma \right\}$ is a feasible solution to the primal LP defined by the Kantorovich distance. \square

A.2 Proof of Theorem 3

Sketch of the proof The idea of the proof is first to bound the difference between value functions following different policies by a bound between value functions estimated in different environments but following the same policy (more precisely a max over a set of policies of such a bound). Once that is done, a bound in probability using Hoeffding's inequality can be obtained.

Let us denote $\varphi \in \phi(\mathcal{H})$:

$$\begin{aligned}
 V_M^{\pi_{D_\infty}, \phi}(\varphi) - V_M^{\pi_D, \phi}(\varphi) &= (V_M^{\pi_{D_\infty}, \phi}(\varphi) - \mathcal{V}_{\hat{M}_D}^{\pi_{D_\infty}, \phi}(\varphi)) - (V_M^{\pi_D, \phi}(\varphi) - \mathcal{V}_{\hat{M}_D}^{\pi_D, \phi}(\varphi)) \\
 & \quad + (\mathcal{V}_{\hat{M}_D}^{\pi_{D_\infty}, \phi}(\varphi) - \mathcal{V}_{\hat{M}_D}^{\pi_D, \phi}(\varphi)) \\
 & \leq (V_M^{\pi_{D_\infty}, \phi}(\varphi) - \mathcal{V}_{\hat{M}_D}^{\pi_{D_\infty}, \phi}(\varphi)) - (V_M^{\pi_D, \phi}(\varphi) - \mathcal{V}_{\hat{M}_D}^{\pi_D, \phi}(\varphi)) \\
 & \leq 2 \max_{\pi \in \{\pi_{D_\infty, \phi}, \pi_{D, \phi}\}} \left| V_M^\pi(\varphi) - \mathcal{V}_{\hat{M}_D}^\pi(\varphi) \right|.
 \end{aligned}$$

It follows that $\forall H$

$$\begin{aligned}
 V_M^{\pi_{D_\infty}, \phi}(\varphi) - V_M^{\pi_D, \phi}(\varphi) & \leq 2 \max_{\pi \in \{\pi_{D_\infty, \phi}, \pi_{D, \phi}\}} \max_{\varphi \in \phi(\mathcal{H})} \left| Q_M^\pi(\varphi, \pi(\varphi)) - \mathcal{Q}_{\hat{M}_D}^\pi(\varphi, \pi(\varphi)) \right| \\
 & \leq 2 \max_{\pi \in \{\pi_{D_\infty, \phi}, \pi_{D, \phi}\}} \max_{\varphi \in \phi(\mathcal{H}), a \in A} \left| Q_M^\pi(\varphi, a) - \mathcal{Q}_{\hat{M}_D}^\pi(\varphi, a) \right|,
 \end{aligned}$$

where $Q_M^\pi(\varphi, a)$ is the action-value function for policy π in M with $a \in \mathcal{A}$. Similarly $Q_{\hat{M}}^\pi(\varphi, a)$ is the action-value function for policy π in \hat{M} .

By Lemma 9, we have:

$$V_M^{\pi_{D_\infty, \phi}}(\varphi) - V_M^{\pi_{D, \phi}}(\varphi) \leq \frac{2}{1 - \gamma} \max_{\pi \in \{\pi_{D_\infty, \phi}, \pi_{D, \phi}\}} \max_{\varphi \in \phi(\mathcal{H}), a \in A} \left| \hat{R}'(\varphi, a) + \gamma \sum_{\varphi' \in \phi(\mathcal{H})} \hat{T}(\varphi, a, \varphi') V_M^\pi(\varphi') - Q_M^\pi(\varphi, a) \right|,$$

where $\hat{R}'(\phi_0(H), a) = \sum_{\varphi \in \phi_0(\mathcal{H})} \hat{T}(\phi_0(H), a, \varphi) \hat{R}(\phi_0(H), a, \varphi)$.

With $\mathcal{R} \in [0, R_{max}]$, we notice that $\hat{R}'(\varphi, a) + \gamma \sum_{\varphi' \in \phi(\mathcal{H})} \hat{T}(\varphi, a, \varphi') V_M^\pi(\varphi')$ is the mean of i.i.d. variables bounded in the interval $[0, \frac{R_{max}}{1-\gamma}]$ and with mean $Q_M^\pi(\varphi, a)$ for any policy $\pi : \phi(\mathcal{H}) \rightarrow A$. Therefore, according to Hoeffding's inequality (Hoeffding, 1963), we have with n the number of tuples for every pair (φ, a) :

$$\mathbb{P} \left\{ \left| \hat{R}'(\varphi, a) + \gamma \sum_{\varphi' \in \phi(\mathcal{H})} \hat{T}(\varphi, a, \varphi') V_M^\pi(\varphi') - Q_M^\pi(\varphi, a) \right| > t \right\} \leq 2 \exp \left(\frac{-2nt^2}{(R_{max}/(1-\gamma))^2} \right). \quad (5)$$

As we want to obtain a bound over all pairs (φ, a) and a union bound on all policies $\pi \in \Pi$ s.t. $|\Pi| = |A|^{|\phi(\mathcal{H})|}$ (indeed Equation 5 does not hold for $\pi_{D, \phi}$ alone because that policy is not chosen randomly), we want the right-hand side of equation 5 to be $\frac{\delta}{|\Pi| |\phi(\mathcal{H})| |A|}$.

This gives $t(\delta) = \left(\frac{R_{max}}{1-\gamma} \right) \sqrt{\left(\frac{1}{2n} \ln \left(\frac{2|\Pi| |\phi(\mathcal{H})| |A|}{\delta} \right) \right)}$ and we conclude that:

$$\max_{\varphi \in \phi(\mathcal{H})} \left(V_M^{\pi_{D_\infty, \phi}}(\varphi) - V_M^{\pi_{D, \phi}}(\varphi) \right) \leq \frac{2R_{max}}{(1-\gamma)^2} \sqrt{\frac{1}{2n} \ln \left(\frac{2|\phi(\mathcal{H})| |A|^{1+|\phi(\mathcal{H})|}}{\delta} \right)}$$

with probability at least $1 - \delta$.

Lemma 9. For any $M = (S, A, T, R, \Omega, O, \gamma)$ and the frequentist-based augmented MDP $\hat{M} = (\Sigma, A, \hat{T}, \hat{R}, \Gamma)$ defined from M according to definition 3.1, we have $\forall \pi : \phi(\mathcal{H}) \rightarrow A$:

$$\max_{\varphi \in \phi(\mathcal{H}), a \in A} \left| Q_M^\pi(\varphi, a) - Q_{\hat{M}_D}^\pi(\varphi, a) \right| \leq \frac{1}{1 - \gamma} \max_{\varphi \in \phi(\mathcal{H}), a \in A} \left| \hat{R}'(\varphi, a) + \gamma \sum_{\varphi' \in \phi(\mathcal{H})} \hat{T}(\varphi, a, \varphi') V_M^\pi(\varphi') - Q_M^\pi(\varphi, a) \right|.$$

Proof. Given any policy π , let us define $\mathcal{Q}_0, \mathcal{Q}_1, \dots, \mathcal{Q}_m$ s.t. $\mathcal{Q}_0(\varphi, a) = Q_M^\pi(\varphi, a)$ and $\mathcal{Q}_m(\varphi, a) = \hat{R}'(\varphi, a) + \gamma \sum_{\varphi' \in \phi(\mathcal{H})} \hat{T}(\varphi, a, \varphi') \mathcal{V}_{m-1}(\varphi')$, where $\mathcal{V}_{m-1}(\varphi) = \mathcal{Q}_{m-1}(\varphi, \pi(\varphi))$.

We have

$$\begin{aligned} \|\mathcal{Q}_m - \mathcal{Q}_{m-1}\|_\infty &\leq \gamma \max_{\varphi \in \phi(\mathcal{H}), a \in A} \left| \sum_{\varphi' \in \phi(\mathcal{H})} \hat{T}(\varphi, a, \varphi') (\mathcal{V}_{m-1} - \mathcal{V}_{m-2})(\varphi') \right| \\ &\leq \gamma \|(\mathcal{V}_{m-1} - \mathcal{V}_{m-2})(\varphi)\|_\infty \\ &\leq \gamma \|(\mathcal{Q}_{m-1} - \mathcal{Q}_{m-2})(\varphi, a)\|_\infty. \end{aligned}$$

Taking the limit of $m \rightarrow \infty$, $\mathcal{Q}_m \rightarrow \mathcal{Q}_M^\pi$, we have

$$\|\mathcal{Q}_{\hat{M}_D}^\pi - \mathcal{Q}_M^\pi\|_\infty \leq \frac{1}{1-\gamma} \|Q_1 - Q_0\|_\infty$$

which completes the proof. \square

A.3 Q-learning with Neural Network as a Function Approximator: Technical Details of Figure 5

The neural network is made up of three intermediate fully connected layers with 20, 50 and 20 neurons with ReLU activation function and is trained with Q-learning. Weights are initialized with a Glorot uniform initializer (Glorot & Bengio, 2010). It is trained using a target Q-network with a freeze interval (see (Mnih, Kavukcuoglu, Silver, et al., 2015)) of 100 mini-batch gradient descent steps. It uses an RMSprop update rule (learning rate of 0.005, $\rho = 0.9$), mini-batches of size 32 and 20000 mini-batch gradient descent steps.

Appendix B. Algorithmic Details

B.1 Microgrid Benchmark

B.1.1 NEURAL NETWORK ARCHITECTURE

The inputs of the neural network architecture are provided by $\phi(H_t)$ (normalized into $[0,1]$), and the outputs represent the Q-values for each discretized action.

The neural network processes the time series (when $h_c > 2$ and $h_p > 2$) thanks to a set of convolutions with 16 filters of 2×1 with stride 1, followed by a convolution with 16 filters of 2×2 with stride 1. The combination of the output of the convolutions and the non-time series inputs is then followed by two fully-connected layers with 50 and 20 neurons. The activation function used is the Rectified Linear Unit (ReLU) except for the output layer where no activation function is used. A sketch of the structure of the neural network is provided in Figure 10.

B.1.2 HYPERPARAMETERS USED IN DQN

The different hyperparameters used for the DQN algorithm (Mnih et al., 2015) are provided in Table 1.




Figure 10: Sketch of the structure of the neural network architecture. The neural network processes the time series using a set of convolutional layers. The output of the convolutions and the other inputs are followed by fully-connected layers and the output layer.

Table 1: Main parameters of the DQN algorithm. The integer k refers to the iteration number from the beginning of the training and the integer l refers to the epoch number.

Parameter	
Update rule	RMSprop ($\rho_{RMS} = 0.9$) (Tieleman, 2012)
Initial learning rate	$\alpha_0 = 0.0002$
Learning rate update rule	$\alpha_{l+1} = 0.99\alpha_l$
RMS decay	0.9
Initial discount factor	$\gamma_0 = 0.9$
Discount factor update rule	$\gamma_{l+1} = \min(0.98, 1 - 0.99(1 - \gamma_l))$
Parameter of the ϵ -greedy	$\epsilon = \max(0.3, 1 - \frac{k}{500,000})$
Replay memory size	1000000
Batch size	32
Freeze interval	1000 iterations
Initialization of the NN layers	Glorot uniform (Glorot & Bengio, 2010)

References

- Abel, D., Hershkowitz, D., & Littman, M. (2016). Near optimal behavior via approximate state abstraction. In *Proceedings of The 33rd International Conference on Machine Learning*, pp. 2915–2923.
- Aberdeen, D. (2003). A (revised) survey of approximate methods for solving partially observable Markov decision processes. *National ICT Australia, Canberra, Australia*, 1.
- Aberdeen, D., Buffet, O., & Thomas, O. (2007). Policy-gradients for PSRs and POMDPs. In *Artificial Intelligence and Statistics*, pp. 3–10.
- Arun-Kumar, S. (2006). On bisimilarities induced by relations on actions. In *Software Engineering and Formal Methods, 2006. SEFM 2006. Fourth IEEE International Conference on*, pp. 41–49. IEEE.
- Cassandra, A. R., Kaelbling, L. P., & Littman, M. L. (1994). Acting optimally in partially

- observable stochastic domains. In *Proceedings of the Twelfth AAAI National Conference on Artificial Intelligence*, Vol. 94, pp. 1023–1028.
- Castro, P. S., Panangaden, P., & Precup, D. (2009). Equivalence relations in fully and partially observable markov decision processes.. In *Twenty-First International Joint Conference on Artificial Intelligence*, Vol. 9, pp. 1653–1658.
- Chen, T., Goodfellow, I., & Shlens, J. (2015). Net2net: Accelerating learning via knowledge transfer. *arXiv preprint arXiv:1511.05641*.
- Farahmand, A.-m. (2011). *Regularization in reinforcement learning*. Ph.D. thesis, University of Alberta.
- Ferns, N., Panangaden, P., & Precup, D. (2004). Metrics for finite Markov decision processes. In *Proceedings of the 20th conference on Uncertainty in artificial intelligence*, pp. 162–169. AUAI Press.
- François-Lavet, V., Fonteneau, R., & Ernst, D. (2015). How to discount deep reinforcement learning: Towards new dynamic strategies. *arXiv preprint arXiv:1512.02011*.
- François-Lavet, V., Taralla, D., Ernst, D., & Fonteneau, R. (2016). Deep reinforcement learning solutions for energy microgrids management. In *European Workshop on Reinforcement Learning*.
- Ghavamzadeh, M., Mannor, S., Pineau, J., & Tamar, A. (2015). Bayesian reinforcement learning: a survey. *Foundations and Trends® in Machine Learning*, 8(5-6), 359–483.
- Glorot, X., & Bengio, Y. (2010). Understanding the difficulty of training deep feedforward neural networks.. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, Vol. 9, pp. 249–256.
- Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural computation*, 9(8), 1735–1780.
- Hoeffding, W. (1963). Probability inequalities for sums of bounded random variables. *Journal of the American statistical association*, 58(301), 13–30.
- Hutter, M. (2014). Extreme state aggregation beyond mdps. In *International Conference on Algorithmic Learning Theory*, pp. 185–199. Springer.
- Jiang, N., Kulesza, A., & Singh, S. (2015a). Abstraction selection in model-based reinforcement learning. In *Proceedings of the 32nd International Conference on Machine Learning (ICML-15)*, pp. 179–188.
- Jiang, N., Kulesza, A., Singh, S., & Lewis, R. (2015b). The dependence of effective planning horizon on model accuracy. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, pp. 1181–1189. International Foundation for Autonomous Agents and Multiagent Systems.
- Jiang, N., & Li, L. (2016). Doubly robust off-policy value evaluation for reinforcement learning. In *Proceedings of The 33rd International Conference on Machine Learning*, pp. 652–661.
- Kaelbling, L. P., Littman, M. L., & Cassandra, A. R. (1998). Planning and acting in partially observable stochastic domains. *Artificial intelligence*, 101(1), 99–134.

- Lange, S., Gabel, T., & Riedmiller, M. (2012). Batch reinforcement learning. In *Reinforcement learning*, pp. 45–73. Springer.
- Littman, M. L. (1994). Memoryless policies: Theoretical limitations and practical results. In *From Animals to Animats 3: Proceedings of the third international conference on simulation of adaptive behavior*, Vol. 3, p. 238. MIT Press.
- Littman, M. L., & Sutton, R. S. (2002). Predictive representations of state. In *Advances in neural information processing systems*, pp. 1555–1561.
- Maillard, O.-A., Ryabko, D., & Munos, R. (2011). Selecting the state-representation in reinforcement learning. In *Advances in Neural Information Processing Systems*, pp. 2627–2635.
- McCallum, A. K. (1996). *Reinforcement learning with selective perception and hidden state*. Ph.D. thesis, University of Rochester.
- Mnih, V., Kavukcuoglu, K., Silver, D., et al. (2015). Human-level control through deep reinforcement learning. *Nature*, 518(7540), 529–533.
- Ortner, R., Maillard, O.-A., & Ryabko, D. (2014). Selecting near-optimal approximate state representations in reinforcement learning. In *International Conference on Algorithmic Learning Theory*, pp. 140–154. Springer.
- Petrik, M., & Scherrer, B. (2009). Biasing approximate dynamic programming with a lower discount factor. In *Advances in neural information processing systems*, pp. 1265–1272.
- Pineau, J., Gordon, G., & Thrun, S. (2003). Point-based value iteration: An anytime algorithm for POMDPs. In *Proceedings of the 18th international joint conference on Artificial intelligence*, Vol. 3, pp. 1025–1032.
- Precup, D. (2000). Eligibility traces for off-policy policy evaluation. *Computer Science Department Faculty Publication Series*, 80.
- Ravindran, B., & Barto, A. G. (2004). Approximate homomorphisms: A framework for non-exact minimization in markov decision processes..
- Ross, S., Pineau, J., Chaib-draa, B., & Kreitmann, P. (2011). A Bayesian approach for learning and planning in partially observable Markov decision processes. *The Journal of Machine Learning Research*, 12, 1729–1770.
- Singh, S., James, M. R., & Rudary, M. R. (2004). Predictive state representations: A new theory for modeling dynamical systems. In *Proceedings of the 20th conference on Uncertainty in artificial intelligence*, pp. 512–519. AUAI Press.
- Singh, S. P., Jaakkola, T. S., & Jordan, M. I. (1994). Learning without state-estimation in partially observable Markovian decision processes.. In *Proceedings of the Eleventh International Conference on Machine Learning*, pp. 284–292.
- Sondik, E. J. (1978). The optimal control of partially observable Markov processes over the infinite horizon: Discounted costs. *Operations research*, 26(2), 282–304.
- Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction*, Vol. 1. MIT press Cambridge.

- Taylor, J., Precup, D., & Panagaden, P. (2009). Bounding performance loss in approximate mdp homomorphisms. In *Advances in Neural Information Processing Systems*, pp. 1649–1656.
- Thomas, P. S., & Brunskill, E. (2016). Data-efficient off-policy policy evaluation for reinforcement learning. In *International Conference on Machine Learning*.
- Tieleman, H. (2012). Lecture 6.5-rmsprop: Divide the gradient by a running average of its recent magnitude. *COURSERA: Neural Networks for Machine Learning*.
- Whitehead, S. D., & Ballard, D. H. (1990). Active perception and reinforcement learning. In *Machine Learning Proceedings 1990*, pp. 179–188. Elsevier.
- Wolfe, A. P. (2006). POMDP homomorphisms. In *NIPS RL Workshop*.
- Zhang, C., Bengio, S., Hardt, M., Recht, B., & Vinyals, O. (2016). Understanding deep learning requires rethinking generalization. *arXiv preprint arXiv:1611.03530*.