

Article

Forward and Backward Bellman Equations Improve the Efficiency of the EM Algorithm for DEC-POMDP

Takehiro Tottori ^{1,*} and Tetsuya J. Kobayashi ^{1,2,3,4} 

¹ Department of Mathematical Informatics, Graduate School of Information Science and Technology, The University of Tokyo, Tokyo 113-8654, Japan; tetsuya@sat.t.u-tokyo.ac.jp

² Department of Electrical Engineering and Information Systems, Graduate School of Engineering, The University of Tokyo, Tokyo 113-8654, Japan

³ Institute of Industrial Science, The University of Tokyo, Tokyo 153-8505, Japan

⁴ Universal Biology Institute, The University of Tokyo, Tokyo 113-8654, Japan

* Correspondence: takehiro_tottori@sat.t.u-tokyo.ac.jp

Abstract: Decentralized partially observable Markov decision process (DEC-POMDP) models sequential decision making problems by a team of agents. Since the planning of DEC-POMDP can be interpreted as the maximum likelihood estimation for the latent variable model, DEC-POMDP can be solved by the EM algorithm. However, in EM for DEC-POMDP, the forward-backward algorithm needs to be calculated up to the infinite horizon, which impairs the computational efficiency. In this paper, we propose the Bellman EM algorithm (BEM) and the modified Bellman EM algorithm (MBEM) by introducing the forward and backward Bellman equations into EM. BEM can be more efficient than EM because BEM calculates the forward and backward Bellman equations instead of the forward-backward algorithm up to the infinite horizon. However, BEM cannot always be more efficient than EM when the size of problems is large because BEM calculates an inverse matrix. We circumvent this shortcoming in MBEM by calculating the forward and backward Bellman equations without the inverse matrix. Our numerical experiments demonstrate that the convergence of MBEM is faster than that of EM.

Keywords: decision-making; planning; multiagent; uncertainty; decentralized partially observable Markov decision process; control as inference



Citation: Tottori, T.; Kobayashi, T.J. Forward and Backward Bellman Equations Improve the Efficiency of the EM Algorithm for DEC-POMDP. *Entropy* **2021**, *23*, 551. <https://doi.org/10.3390/e23050551>

Academic Editor: Hector Zenil

Received: 19 March 2021

Accepted: 26 April 2021

Published: 29 April 2021

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2021 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Markov decision process (MDP) models sequential decision making problems and has been used for planning and reinforcement learning [1–4]. MDP consists of an environment and an agent. The agent observes the state of the environment and controls it by taking actions. The planning of MDP is to find the optimal control policy maximizing the objective function, which is typically solved by the Bellman equation-based algorithms such as value iteration and policy iteration [1–4].

Decentralized partially observable MDP (DEC-POMDP) is an extension of MDP to a multiagent and partially observable setting, which models sequential decision making problems by a team of agents [5–7]. DEC-POMDP consists of an environment and multiple agents, and the agents cannot observe the state of the environment and the actions of the other agents completely. The agents infer the environmental state and the other agents' actions from their observation histories and control them by taking actions. The planning of DEC-POMDP is to find not only the optimal control policy but also the optimal inference policy for each agent, which maximize the objective function [5–7]. Applications of DEC-POMDP include planetary exploration by a team of rovers [8], target tracking by a team of sensors [9], and information transmission by a team of devices [10]. Since the agents cannot observe the environmental state and the other agents' actions completely,

it is difficult to extend the Bellman equation-based algorithms for MDP to DEC-POMDP straightforwardly [11–15].

DEC-POMDP can be solved using control as inference [16,17]. Control as inference is a framework to interpret a control problem as an inference problem by introducing auxiliary variables [18–22]. Although control as inference has several variants, Toussaint and Storkey showed that the planning of MDP can be interpreted as the maximum likelihood estimation for a latent variable model [18]. Thus, the planning of MDP can be solved by EM algorithm, which is the typical algorithm for the maximum likelihood estimation of latent variable models [23]. Since the EM algorithm is more general than the Bellman equation-based algorithms, it can be straightforwardly extended to POMDP [24,25] and DEC-POMDP [16,17]. The computational efficiency of the EM algorithm for DEC-POMDP is comparable to that of other algorithms for DEC-POMDP [16,17,26–28], and the extensions to the average reward setting and to the reinforcement learning setting have been studied [29–31].

However, the EM algorithm for DEC-POMDP is not efficient enough to be applied to real-world problems, which often have a large number of agents or a large size of an environment. Therefore, there are several studies in which improvement of the computational efficiency of the EM algorithm for DEC-POMDP was attempted [26,27]. Because these studies achieve improvements by restricting possible interactions between agents, their applicability is limited. Therefore, it is desirable to have improvement in the efficiency for more general DEC-POMDP problems.

In order to improve the computational efficiency of EM algorithm for general DEC-POMDP problems, there are two problems that need to be resolved. The first problem is the forward–backward algorithm up to the infinite horizon. The EM algorithm for DEC-POMDP uses the forward–backward algorithm, which has also been used in EM algorithm for hidden Markov models [23]. However, in the EM algorithm for DEC-POMDP, the forward–backward algorithm needs to be calculated up to the infinite horizon, which impairs the computational efficiency [32,33]. The second problem is the Bellman equation. The EM algorithm for DEC-POMDP does not use the Bellman equation, which plays a central role in the the planning and in the reinforcement learning for MDP [1–4]. Therefore, the EM algorithm for DEC-POMDP cannot use the advanced techniques based on the Bellman equation, which makes it possible to solve large-size problems [34–36].

In some previous studies, resolution of these problems was attempted by replacing the forward–backward algorithm up to the infinite horizon with the Bellman equation [32,33]. However, in these studies, the computational efficiency could not be improved completely. For example, Song et al. replaced the forward–backward algorithm with the Bellman equation and showed that their algorithm is more efficient than EM and other DEC-POMDP algorithms by the numerical experiments [32]. However, since a parameter dependency is overlooked in [32], their algorithm may not find the optimal policy under a general situation (see Appendix D for more details). Moreover, Kumar et al. showed that the forward–backward algorithm can be replaced by linear programming with the Bellman equation as a constraint [33]. However, their algorithm may be less efficient than the EM algorithm when the size of problems is large. Therefore, previous studies have not yet completely improved the computational efficiency of EM algorithm for DEC-POMDP.

In this paper, we propose more efficient algorithms for DEC-POMDP than EM algorithm by introducing the forward and backward Bellman equations into it. The backward Bellman equation corresponds to the traditional Bellman equation, which has been used in previous studies [32,33]. In contrast, the forward Bellman equation has not yet been used for the planning of DEC-POMDP explicitly. This equation is similar to that recently proposed in the offline reinforcement learning of MDP [37–39]. In the offline reinforcement learning of MDP, the forward Bellman equation is used to correct the difference between the data sampling policy and the policy to be evaluated. In the planning of DEC-POMDP, the forward Bellman equation plays the important role in inferring the environmental state.

We propose the Bellman EM algorithm (BEM) and the modified Bellman EM algorithm (MBEM) by replacing the forward–backward algorithm with the forward and backward

Bellman equations. They are different in terms of how to solve the forward and backward Bellman equations. BEM solves the forward and backward Bellman equations by calculating an inverse matrix. BEM can be more efficient than EM because BEM does not calculate the forward–backward algorithm up to the infinite horizon. However, since BEM calculates the inverse matrix, it cannot always be more efficient than EM when the size of problems is large, which is the same problem as [33]. Actually, BEM is essentially the same as [33]. In the linear programming problem of [33], the number of variables is equal to that of constraints, which enables us to solve it only from the constraints without the optimization. Therefore, the algorithm in [33] becomes equivalent to BEM, and they suffers from the same problem as BEM.

This problem is addressed by MBEM. MBEM solves the forward and backward Bellman equations by applying the forward and backward Bellman operators to the arbitrary initial functions infinite times. Although MBEM needs to calculate the forward and backward Bellman operators infinite times, which is the same problem with EM, MBEM can evaluate approximation errors more tightly owing to the contractibility of these operators. It can also utilize the information of the previous iteration owing to the arbitrariness of the initial functions. These properties enable MBEM to be more efficient than EM. Moreover, MBEM resolves the drawback of BEM because MBEM does not calculate the inverse matrix. Therefore, MBEM can be more efficient than EM even when the size of problems is large. Our numerical experiments demonstrate that the convergence of MBEM is faster than that of EM regardless of the size of problems.

The paper is organized as follows: In Section 2, DEC-POMDP is formulated. In Section 3, the EM algorithm for DEC-POMDP, which was proposed in [16], is briefly reviewed. In Section 4, the forward and backward Bellman equations are derived, and the Bellman EM algorithm (BEM) is proposed. In Section 5, the forward and backward Bellman operators are defined, and the modified Bellman EM algorithm (MBEM) is proposed. In Section 6, EM, BEM, and MBEM are summarized and compared. In Section 7, the performances of EM, BEM, and MBEM are compared through the numerical experiment. In Section 8, this paper is concluded, and future works are discussed.

2. DEC-POMDP

DEC-POMDP consists of an environment and N agents (Figures 1 and 2a) [7,16]. $x_t \in \mathcal{X}$ is the state of the environment at time t . $y_t^i \in \mathcal{Y}^i$, $z_t^i \in \mathcal{Z}^i$, and $a_t^i \in \mathcal{A}^i$ are the observation, the memory, and the action available to the agent $i \in \{1, \dots, N\}$, respectively. \mathcal{X} , \mathcal{Y}^i , \mathcal{Z}^i , and \mathcal{A}^i are finite sets. $\mathbf{y}_t := (y_t^1, \dots, y_t^N)$, $\mathbf{z}_t := (z_t^1, \dots, z_t^N)$, and $\mathbf{a}_t := (a_t^1, \dots, a_t^N)$ are the joint observation, the joint memory, and the joint action of the N agents, respectively.

The time evolution of the environmental state x_t is given by the initial state probability $p(x_0)$ and the state transition probability $p(x_{t+1}|x_t, \mathbf{a}_t)$. Thus, agents can control the environmental state x_{t+1} by taking appropriate actions \mathbf{a}_t . The agent i cannot observe the environmental state x_t and the joint action \mathbf{a}_{t-1} completely, and obtains the observation y_t^i instead of them. Thus, the observation y_t^i obeys the observation probability $p(y_t^i|x_t, \mathbf{a}_{t-1})$. The agent i updates its memory from z_{t-1}^i to z_t^i based on the observation y_t^i . Thus, the memory z_t^i obeys the initial memory probability $v^i(z_0^i)$ and the memory transition probability $\lambda^i(z_t^i|z_{t-1}^i, y_t^i)$. The agent i takes the action a_t^i based on the memory z_t^i by following the action probability $\pi^i(a_t^i|z_t^i)$. The reward function $r(x_t, \mathbf{a}_t)$ defines the amount of reward that is obtained at each step depending on the state of the environment x_t and the joint action \mathbf{a}_t taken by the agents.

The objective function in the planning of DEC-POMDP is given by the expected return, which is the expected discounted cumulative reward:

$$J(\theta) = \mathbb{E}_{\theta} \left[\sum_{t=0}^{\infty} \gamma^t r(x_t, \mathbf{a}_t) \right]. \quad (1)$$

$\theta := (\pi, \lambda, \nu)$ is the policy, where $\pi := (\pi^1, \dots, \pi^N)$, $\lambda := (\lambda^1, \dots, \lambda^N)$, and $\nu := (\nu^1, \dots, \nu^N)$. $\gamma \in (0, 1)$ is the discount factor, which decreases the weight of the future reward. The closer γ is to 1, the closer the weight of the future reward is to that of the current reward.

The planning of DEC-POMDP is to find the policy θ that maximizes the expected return $J(\theta)$ as follows:

$$\theta^* := \arg \max_{\theta} J(\theta). \quad (2)$$

In other words, the planning of DEC-POMDP is to find how to take the action and how to update the memory for each agent to maximize the expected return.




Figure 1. Schematic diagram of DEC-POMDP. DEC-POMDP consists of an environment and N agents ($N = 2$ in this figure). x_t is the state of the environment at time t . y_t^i , z_t^i , and a_t^i are the observation, the memory, and the action available to the agent $i \in \{1, \dots, N\}$, respectively. The agents update their memories based on their observations, and take their actions based on their memories to control the environmental state. The planning of DEC-POMDP is to find their optimal memory updates and action selections that maximize the objective function.

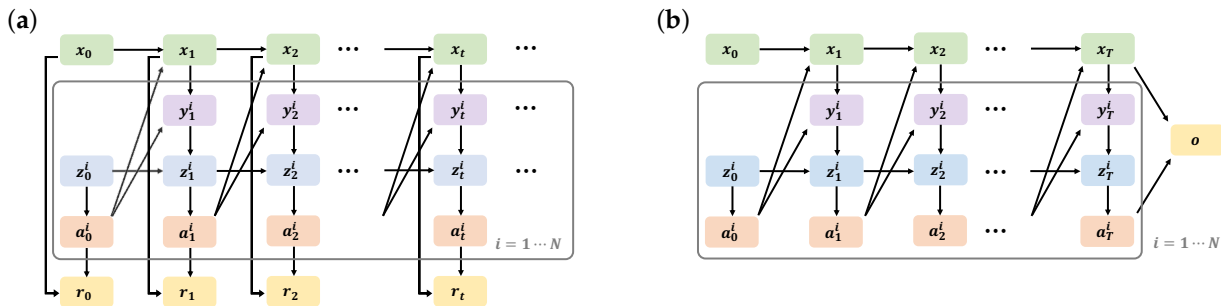


Figure 2. Dynamic Bayesian networks of DEC-POMDP (a) and the latent variable model for the time horizon $T \in \{0, 1, 2, \dots\}$ (b). x_t is the state of the environment at time t . y_t^i , z_t^i , and a_t^i are the observation, the memory, and the action available to the agent $i \in \{1, \dots, N\}$, respectively. (a) $r_t \in \mathbb{R}$ is the reward, which is generated at each time. (b) $o \in \{0, 1\}$ is the optimal variable, which is generated only at the time horizon T .

3. EM Algorithm for DEC-POMDP

In this section, we explain the EM algorithm for DEC-POMDP, which was proposed in [16].

3.1. Control as Inference

In this subsection, we show that the planning of DEC-POMDP can be interpreted as the maximum likelihood estimation for a latent variable model (Figure 2b).

We introduce two auxiliary random variables: the time horizon $T \in \{0, 1, 2, \dots\}$ and the optimal variable $o \in \{0, 1\}$. These variables obey the following probabilities:

$$p(T) = (1 - \gamma)\gamma^T, \quad (3)$$

$$p(o = 1 | x_T, \mathbf{a}_T) = \bar{r}(x_T, \mathbf{a}_T) := \frac{r(x_T, \mathbf{a}_T) - r_{\min}}{r_{\max} - r_{\min}} \quad (4)$$

where r_{\max} and r_{\min} are the maximum and the minimum value of the reward function $r(x, \mathbf{a})$, respectively. Thus, $\bar{r}(x, \mathbf{a}) \in [0, 1]$ is satisfied.

By introducing these variables, DEC-POMDP changes from Figure 2a to Figure 2b. While Figure 2a considers the infinite time horizon, Figure 2b considers the finite time horizon T , which obeys Equation (3). Moreover, while the reward $r_t := r(x_t, \mathbf{a}_t)$ is generated at each time in Figure 2a, the optimal variable o is generated only at the time horizon T in Figure 2b.

Theorem 1 ([16]). The expected return $J(\theta)$ in DEC-POMDP (Figure 2a) is linearly related to the likelihood $p(o = 1; \theta)$ in the latent variable model (Figure 2b) as follows:

$$J(\theta) = (1 - \gamma)^{-1}[(r_{\max} - r_{\min})p(o = 1; \theta) + r_{\min}]. \quad (5)$$

Note that o is the observable variable, and $x_{0:T}, \mathbf{y}_{1:T}, \mathbf{z}_{0:T}, \mathbf{a}_{0:T}, T$ are the latent variables.

Proof. See Appendix A.1. \square

Therefore, the planning of DEC-POMDP is equivalent to the maximum likelihood estimation for the latent variable model as follows:

$$\theta^* = \arg \max_{\theta} p(o = 1; \theta). \quad (6)$$

Intuitively, while the planning of DEC-POMDP is to find the policy which maximizes the reward, the maximum likelihood estimation for the latent variable model is to find the policy which maximizes the probability of the optimal variable. Since the probability of the optimal variable is proportional to the reward, the planning of DEC-POMDP is equivalent to the maximum likelihood estimation for the latent variable model.

3.2. EM Algorithm

Since the planning of DEC-POMDP can be interpreted as the maximum likelihood estimation for the latent variable model, it can be solved by the EM algorithm [16]. EM algorithm is the typical algorithm for the maximum likelihood estimation of latent variable models, which iterates two steps, E step and M step [23].

In the E step, we calculate the Q function, which is defined as follows:

$$Q(\theta; \theta_k) := \mathbb{E}_{\theta_k} [\log p(o = 1, x_{0:T}, \mathbf{y}_{1:T}, \mathbf{z}_{0:T}, \mathbf{a}_{0:T}, T; \theta) | o = 1] \quad (7)$$

where θ_k is the current estimator of the optimal policy.

In the M step, we update θ_k to θ_{k+1} by maximizing the Q function as follows:

$$\theta_{k+1} := \arg \max_{\theta} Q(\theta; \theta_k). \quad (8)$$

Since each iteration between the E step and the M step monotonically increases the likelihood $p(o = 1; \theta_k)$, we can find θ^* that locally maximizes the likelihood $p(o = 1; \theta)$.

3.3. M Step

Proposition 1 ([16]). In the EM algorithm for DEC-POMDP, Equation (8) can be calculated as follows:

$$\pi_{k+1}^i(a|z) = \frac{\sum_{\mathbf{a}^{-i}, \mathbf{z}^{-i}} \pi_k(\mathbf{a}|z) \sum_{x, x', z'} p(x', z'|x, z, \mathbf{a}; \lambda_k) F(x, z; \theta_k) (\bar{r}(x, \mathbf{a}) + \gamma V(x, z; \theta_k))}{\sum_{\mathbf{a}, \mathbf{z}} \pi_k(\mathbf{a}|z) \sum_{x, x', z'} p(x', z'|x, z, \mathbf{a}; \lambda_k) F(x, z; \theta_k) (\bar{r}(x, \mathbf{a}) + \gamma V(x, z; \theta_k))}, \quad (9)$$

$$\lambda_{k+1}^i(z'|z, y') = \frac{\sum_{\mathbf{z}^{-i'}, \mathbf{z}^{-i}, \mathbf{y}^{-i'}} \lambda_k(z'^i|z^i, y'^i) \sum_{x', x} p(x', y'|x, z; \pi_k) F(x, z; \theta_k) V(x', z'; \theta_k)}{\sum_{\mathbf{z}', \mathbf{z}, \mathbf{y}'} \lambda_k(z'^i|z^i, y'^i) \sum_{x', x} p(x', y'|x, z; \pi_k) F(x, z; \theta_k) V(x', z'; \theta_k)}, \quad (10)$$

$$v_{k+1}^i(z) = \frac{\sum_{\mathbf{z}^{-i}} v_k(z) \sum_x p_0(x) V(x, z; \theta_k)}{\sum_{\mathbf{z}} v_k(z) \sum_x p_0(x) V(x, z; \theta_k)}. \quad (11)$$

$\mathbf{a}^{-i} := (a^1, \dots, a^{i-1}, a^{i+1}, \dots, a^N)$. \mathbf{y}^{-i} and \mathbf{z}^{-i} are defined in the same way. $F(x, z; \theta)$ and $V(x, z; \theta)$ are defined as follows:

$$F(x, z; \theta) := \sum_{t=0}^{\infty} \gamma^t p_t(x, z; \theta), \quad (12)$$

$$V(x, z; \theta) := \sum_{t=0}^{\infty} \gamma^t p_0(o = 1|x, z, T = t; \theta) \quad (13)$$

where $p_t(x, z; \theta) := p(x_t = x, z_t = z; \theta)$, and $p_0(o = 1|x, z, T; \theta) := p(o = 1|x_0 = x, z_0 = z, T; \theta)$.

Proof. See Appendix A.2. \square

$F(x, z; \theta)$ quantifies the frequency of the state x and the memory z , which is called the frequency function in this paper. $V(x, z; \theta)$ quantifies the probability of $o = 1$ when the initial state and memory are x and z , respectively. Since the probability of $o = 1$ is proportional to the reward, $V(x, z; \theta)$ is called the value function in this paper. Actually, $V(x, z; \theta)$ corresponds to the value function [33].

3.4. E Step

$F(x, z; \theta_k)$ and $V(x, z; \theta_k)$ need to be obtained to calculate Equations (9)–(11). In [16], $F(x, z; \theta_k)$ and $V(x, z; \theta_k)$ are calculated by the forward-backward algorithm, which has been used in EM algorithm for the hidden Markov model [23].

In [16], the forward probability $\alpha_t(x, z)$ and the backward probability $\beta_t(x, z)$ are defined as follows:

$$\alpha_t(x, z; \theta_k) := p_t(x, z; \theta_k), \quad (14)$$

$$\beta_t(x, z; \theta_k) := p_0(o = 1|x, z, T = t; \theta_k). \quad (15)$$

It is easy to calculate $\alpha_0(x, z; \theta_k)$ and $\beta_0(x, z; \theta_k)$ as follows:

$$\alpha_0(x, z; \theta_k) = p_0(x, z; \mathbf{v}_k) := p_0(x) v_k(z), \quad (16)$$

$$\beta_0(x, z; \theta_k) = \bar{r}(x, z; \pi_k) := \sum_{\mathbf{a}} \pi_k(\mathbf{a}|z) \bar{r}(x, \mathbf{a}). \quad (17)$$

Moreover, $\alpha_{t+1}(x, z; \theta_k)$ and $\beta_{t+1}(x, z; \theta_k)$ are easily calculated from $\alpha_t(x, z; \theta_k)$ and $\beta_t(x, z; \theta_k)$:

$$\alpha_{t+1}(x, z; \theta_k) = \sum_{x', z'} p(x, z|x', z'; \theta_k) \alpha_t(x', z'; \theta_k), \quad (18)$$

$$\beta_{t+1}(x, z; \theta_k) = \sum_{x', z'} \beta_t(x', z'; \theta_k) p(x', z'|x, z; \theta_k) \quad (19)$$

where

$$p(x', z' | x, z; \theta_k) = \sum_{y', a} \lambda_k(z' | z, y') p(y' | x', a) p(x' | x, a) \pi_k(a | z). \quad (20)$$

Equations (18) and (19) are called the forward and backward equations, respectively. Using Equations (16)–(19), $\alpha_t(x, z; \theta_k)$ and $\beta_t(x, z; \theta_k)$ can be efficiently calculated from $t = 0$ to $t = \infty$, which is called the forward–backward algorithm [23].

By calculating the forward–backward algorithm from $t = 0$ to $t = \infty$, $F(x, z; \theta_k)$ and $V(x, z; \theta_k)$ can be obtained as follows [16]:

$$F(x, z; \theta_k) = \sum_{t=0}^{\infty} \gamma^t \alpha_t(x, z; \theta_k), \quad (21)$$

$$V(x, z; \theta_k) = \sum_{t=0}^{\infty} \gamma^t \beta_t(x, z; \theta_k). \quad (22)$$

However, $F(x, z; \theta_k)$ and $V(x, z; \theta_k)$ cannot be calculated exactly by this approach because it is practically impossible to calculate the forward–backward algorithm until $t = \infty$. Therefore, the forward–backward algorithm needs to be terminated at $t = T_{\max}$, where T_{\max} is finite. In this case, $F(x, z; \theta_k)$ and $V(x, z; \theta_k)$ are approximated as follows:

$$F(x, z; \theta_k) = \sum_{t=0}^{\infty} \gamma^t \alpha_t(x, z; \theta_k) \approx \sum_{t=0}^{T_{\max}} \gamma^t \alpha_t(x, z; \theta_k), \quad (23)$$

$$V(x, z; \theta_k) = \sum_{t=0}^{\infty} \gamma^t \beta_t(x, z; \theta_k) \approx \sum_{t=0}^{T_{\max}} \gamma^t \beta_t(x, z; \theta_k). \quad (24)$$

T_{\max} needs to be large enough to reduce the approximation errors. In the previous study, a heuristic termination condition was proposed as follows [16]:

$$\gamma^{T_{\max}} p(o = 1 | T_{\max}; \theta_k) \ll \sum_{T=0}^{T_{\max}-1} \gamma^T p(o = 1 | T; \theta_k). \quad (25)$$

$p(o = 1 | T; \theta_k) = \sum_{x, z} \alpha_{T'}(x, z; \theta_k) \beta_{T''}(x, z; \theta_k)$ where $T = T' + T''$. However, the relation between T_{\max} and the approximation errors is unclear in Equation (25). We propose a new termination condition to guarantee the approximation errors as follows:

Proposition 2. We set an acceptable error bound $\varepsilon > 0$. If

$$T_{\max} > \frac{\log(1 - \gamma)\varepsilon}{\log \gamma} - 1 \quad (26)$$

is satisfied, then

$$\left\| F(x, z; \theta_k) - \sum_{t=0}^{T_{\max}} \gamma^t \alpha_t(x, z; \theta_k) \right\|_{\infty} < \varepsilon, \quad (27)$$

$$\left\| V(x, z; \theta_k) - \sum_{t=0}^{T_{\max}} \gamma^t \beta_t(x, z; \theta_k) \right\|_{\infty} < \varepsilon \quad (28)$$

are satisfied.

Proof. See Appendix A.3. \square

3.5. Summary

In summary, the EM algorithm for DEC-POMDP is given by Algorithm 1. In the E step, we calculate $\alpha_t(x, z; \theta_k)$ and $\beta_t(x, z; \theta_k)$ from $t = 0$ to $t = T_{\max}$ by the forward-backward algorithm. In M step, we update θ_k to θ_{k+1} using Equations (9)–(11). The time complexities of the E step and the M step are $\mathcal{O}((|\mathcal{X}||\mathcal{Z}|)^2 T_{\max})$ and $\mathcal{O}((|\mathcal{X}||\mathcal{Z}|)^2 |\mathcal{Y}||\mathcal{A}|)$, respectively. Note that $\mathcal{A} := \otimes_{i=1}^N \mathcal{A}^i$, and \mathcal{Y} and \mathcal{Z} are defined in the same way. The EM algorithm for DEC-POMDP is less efficient when the discount factor γ is closer to 1 or the acceptable error bound ε is smaller because T_{\max} needs to be larger in these cases.

Algorithm 1 EM algorithm for DEC-POMDP

```

 $k \leftarrow 0$ , Initialize  $\theta_k$ .
 $T_{\max} \leftarrow \lceil (\log(1 - \gamma)\varepsilon) / \log \gamma - 1 \rceil$ 
while  $\theta_k$  or  $J(\theta_k)$  do not converge do
    Calculate  $p(x', z' | x, z; \theta_k)$  by Equation (20).
    //—E step—//
     $\alpha_0(x, z; \theta_k) \leftarrow p_0(x, z; \nu_k)$ 
     $\beta_0(x, z; \theta_k) \leftarrow \bar{r}(x, z; \pi_k)$ 
    for  $t = 1, 2, \dots, T_{\max}$  do
         $\alpha_t(x, z; \theta_k) \leftarrow \sum_{x', z'} p(x, z | x', z'; \theta_k) \alpha_{t-1}(x', z'; \theta_k)$ 
         $\beta_t(x, z; \theta_k) \leftarrow \sum_{x', z'} \beta_{t-1}(x', z'; \theta_k) p(x', z' | x, z; \theta_k)$ 
    end for
     $F(x, z; \theta_k) \leftarrow \sum_{t=0}^{T_{\max}} \gamma^t \alpha_t(x, z; \theta_k)$ 
     $V(x, z; \theta_k) \leftarrow \sum_{t=0}^{T_{\max}} \gamma^t \beta_t(x, z; \theta_k)$ 
    //—M step—//
    Update  $\theta_k$  to  $\theta_{k+1}$  by Equations (9)–(11).
     $k \leftarrow k + 1$ 
end while
return  $\theta_k$ 

```

4. Bellman EM Algorithm

In the EM algorithm for DEC-POMDP, $\alpha_t(x, z; \theta_k)$ and $\beta_t(x, z; \theta_k)$ are calculated from $t = 0$ to $t = T_{\max}$ to obtain $F(x, z; \theta_k)$ and $V(x, z; \theta_k)$. However, T_{\max} needs to be large to reduce the approximation errors of $F(x, z; \theta_k)$ and $V(x, z; \theta_k)$, which impairs the computational efficiency of the EM algorithm for DEC-POMDP [32,33]. In this section, we calculate $F(x, z; \theta_k)$ and $V(x, z; \theta_k)$ directly without calculating $\alpha_t(x, z; \theta_k)$ and $\beta_t(x, z; \theta_k)$ to resolve the drawback of EM.

4.1. Forward and Backward Bellman Equations

The following equations are useful to obtain $F(x, z; \theta_k)$ and $V(x, z; \theta_k)$ directly:

Theorem 2. $F(x, z; \theta)$ and $V(x, z; \theta)$ satisfy the following equations:

$$F(x, z; \theta) = p_0(x, z; \nu) + \gamma \sum_{x', z'} p(x, z | x', z'; \theta) F(x', z'; \theta), \quad (29)$$

$$V(x, z; \theta) = \bar{r}(x, z; \pi) + \gamma \sum_{x', z'} p(x', z' | x, z; \theta) V(x', z'; \theta). \quad (30)$$

Equations (29) and (30) are called the forward Bellman equation and the backward Bellman equation, respectively.

Proof. See Appendix B. \square

Note that the direction of time is different between Equations (29) and (30). In Equation (29), x' and z' are earlier state and memory than x and z , respectively. In Equation (30), x' and z' are later state and memory than x and z , respectively.

The backward Bellman Equation (30) corresponds to the traditional Bellman equation, which has been used in other algorithms for DEC-POMDP [11–14]. In contrast, the forward Bellman equation, which is introduced in this paper, is similar to that recently proposed in the offline reinforcement learning [37–39].

Since the forward and backward Bellman equations are linear equations, they can be solved exactly as follows:

$$F(\theta) = (I - \gamma P(\theta))^{-1} p(v), \quad (31)$$

$$V(\theta) = ((I - \gamma P(\theta))^{-1})^T r(\pi) \quad (32)$$

where

$$F_i(\theta) := F((x, z) = i; \theta), V_i(\theta) := V((x, z) = i; \theta)$$

$$P_{ij}(\theta) := p((x', z') = i | (x, z) = j; \theta)$$

$$p_i(v) := p_0((x, z) = i; v), r_i(\pi) := \bar{r}((x, z) = i; \pi)$$

Therefore, we can obtain $F(x, z; \theta_k)$ and $V(x, z; \theta_k)$ from the forward and backward Bellman equations.

4.2. Bellman EM Algorithm (BEM)

The forward–backward algorithm from $t = 0$ to $t = T_{\max}$ in the EM algorithm for DEC-POMDP can be replaced by the forward and backward Bellman equations. In this paper, the EM algorithm for DEC-POMDP that uses the forward and backward Bellman equations instead of the forward–backward algorithm from $t = 0$ to $t = T_{\max}$ is called the Bellman EM algorithm (BEM).

4.3. Comparison of EM and BEM

BEM is summarized as Algorithm 2. The M step in Algorithm 2 is almost the same as that in Algorithm 1—only the E step is different. While the time complexity of the E step in EM is $\mathcal{O}((|\mathcal{X}||\mathcal{Z}|)^2 T_{\max})$, that in BEM is $\mathcal{O}((|\mathcal{X}||\mathcal{Z}|)^3)$.

Algorithm 2 Bellman EM algorithm (BEM)

```

 $k \leftarrow 0$ , Initialize  $\theta_k$ .
while  $\theta_k$  or  $J(\theta_k)$  do not converge do
  Calculate  $p(x', z' | x, z; \theta_k)$  by Equation (20).
  //—E step—//
   $F(\theta_k) \leftarrow (I - \gamma P(\theta_k))^{-1} p(v_k)$ 
   $V(\theta_k) \leftarrow ((I - \gamma P(\theta_k))^{-1})^T r(\pi_k)$ 
  //—M step—//
  Update  $\theta_k$  to  $\theta_{k+1}$  by Equations (9)–(11).
   $k \leftarrow k + 1$ 
end while
return  $\theta_k$ 

```

BEM can calculate $F(x, z; \theta_k)$ and $V(x, z; \theta_k)$ exactly. Moreover, BEM can be more efficient than EM when the discount factor γ is close to 1 or the acceptable error bound ε is small because T_{\max} needs to be large enough in these cases. However, when the size of the state space $|\mathcal{X}|$ or that of the joint memory space $|\mathcal{Z}|$ is large, BEM cannot always be more efficient than EM because BEM needs to calculate the inverse matrix $(I - \gamma P(\theta_k))^{-1}$. To circumvent this shortcoming, we propose a new algorithm, the modified Bellman EM algorithm (MBEM), to obtain $F(x, z; \theta_k)$ and $V(x, z; \theta_k)$ without calculating the inverse matrix.

5. Modified Bellman EM Algorithm

5.1. Forward and Backward Bellman Operators

We define the forward and backward Bellman operators as follows:

$$A_{\theta}f(x, z) := p_0(x, z; \nu) + \gamma \sum_{x', z'} p(x, z | x', z'; \theta) f(x', z'), \quad (33)$$

$$B_{\theta}v(x, z) := \bar{r}(x, z; \pi) + \gamma \sum_{x', z'} p(x', z' | x, z; \theta) v(x', z') \quad (34)$$

where $\forall f, v : \mathcal{X} \times \mathcal{Z} \rightarrow \mathbb{R}$. From the forward and backward Bellman equations, A_{θ} and B_{θ} satisfy the following equations:

$$F(x, z; \theta) = A_{\theta}F(x, z; \theta), \quad (35)$$

$$V(x, z; \theta) = B_{\theta}V(x, z; \theta). \quad (36)$$

Thus, $F(x, z; \theta_k)$ and $V(x, z; \theta_k)$ are the fixed points of A_{θ} and B_{θ} , respectively. A_{θ} and B_{θ} have the following useful property:

Proposition 3. A_{θ} and B_{θ} are contractive operators as follows:

$$\|A_{\theta}f - A_{\theta}g\|_1 \leq \gamma \|f - g\|_1, \quad (37)$$

$$\|B_{\theta}u - B_{\theta}v\|_{\infty} \leq \gamma \|u - v\|_{\infty} \quad (38)$$

where $\forall f, g, u, v : \mathcal{X} \times \mathcal{Z} \rightarrow \mathbb{R}$.

Proof. See Appendix C.1. \square

Note that the norm is different between Equations (37) and (38). It is caused by the difference of the time direction between A_{θ} and B_{θ} . While x' and z' are earlier state and memory than x and z , respectively, in the forward Bellman operator A_{θ} , x' and z' are later state and memory than x and z , respectively, in the backward Bellman operator B_{θ} .

We obtain $F(x, z; \theta_k)$ and $V(x, z; \theta_k)$ using Equations (35)–(38), as follows:

Proposition 4.

$$\lim_{L \rightarrow \infty} A_{\theta}^L f(x, z) = F(x, z; \theta), \quad (39)$$

$$\lim_{L \rightarrow \infty} B_{\theta}^L v(x, z) = V(x, z; \theta) \quad (40)$$

where $\forall f, v : \mathcal{X} \times \mathcal{Z} \rightarrow \mathbb{R}$.

Proof. See Appendix C.2. \square

Therefore, it is shown that $F(x, z; \theta_k)$ and $V(x, z; \theta_k)$ can be calculated by applying the forward and backward Bellman operators, A_{θ} and B_{θ} , to arbitrary initial functions, $f(x, z)$ and $v(x, z)$, infinite times.

5.2. Modified Bellman EM Algorithm (MBEM)

The calculation of the forward and backward Bellman equations in BEM can be replaced by that of the forward and backward Bellman operators. In this paper, BEM that uses the forward and backward Bellman operators instead of the forward and backward Bellman equations is called modified Bellman EM algorithm (MBEM).

5.3. Comparison of EM, BEM, and MBEM

Since MBEM does not need the inverse matrix, MBEM can be more efficient than BEM when the size of the state space $|\mathcal{X}|$ and that of the joint memory space $|\mathcal{Z}|$ are large. Thus, MBEM resolves the drawback of BEM.

On the other hand, MBEM has the same problem as EM. MBEM calculates A_{θ_k} and B_{θ_k} infinite times to obtain $F(x, z; \theta_k)$ and $V(x, z; \theta_k)$. However, since it is practically impossible to calculate A_{θ_k} and B_{θ_k} infinite times, the calculation of A_{θ_k} and B_{θ_k} needs to be terminated after L_{\max} times, where L_{\max} is finite. In this case, $F(x, z; \theta_k)$ and $V(x, z; \theta_k)$ are approximated as follows:

$$F(x, z; \theta_k) = A_{\theta_k}^{\infty} f(x, z) \approx A_{\theta_k}^{L_{\max}} f(x, z), \quad (41)$$

$$V(x, z; \theta_k) = B_{\theta_k}^{\infty} v(x, z) \approx B_{\theta_k}^{L_{\max}} v(x, z). \quad (42)$$

L_{\max} needs to be large enough to reduce the approximation errors of $F(x, z; \theta_k)$ and $V(x, z; \theta_k)$, which impairs the computational efficiency of MBEM. Thus, MBEM can potentially suffer from the same problem as EM. However, we can theoretically show that MBEM is more efficient than EM by comparing T_{\max} and L_{\max} under the condition that the approximation errors of $F(x, z; \theta_k)$ and $V(x, z; \theta_k)$ are smaller than the acceptable error bound ε .

When $f(x, z) = p_0(x, z; \nu_k)$ and $v(x, z) = \bar{r}(x, z; \pi_k)$, Equations (41) and (42) can be calculated as follows:

$$A_{\theta_k}^{L_{\max}} f(x, z) = \sum_{t=0}^{L_{\max}} \gamma^t p_t(x, z; \theta), \quad (43)$$

$$B_{\theta_k}^{L_{\max}} v(x, z) = \sum_{t=0}^{L_{\max}} \gamma^t p_0(o = 1 | x, z, T = t; \theta) \quad (44)$$

which are the same with Equations (23) and (24), respectively. Thus, in this case, $L_{\max} = T_{\max}$, and the computational efficiency of MBEM is the same as that of EM. However, MBEM has two useful properties that EM does not have, and therefore, MBEM can be more efficient than EM. In the following, we explain these properties in more detail.

The first property of MBEM is the contractibility of the forward and backward Bellman operators, A_{θ_k} and B_{θ_k} . From the contractibility of the Bellman operators, L_{\max} is determined adaptively as follows:

Proposition 5. We set an acceptable error bound $\varepsilon > 0$. If

$$\|A_{\theta_k}^L f - A_{\theta_k}^{L-1} f\|_1 < \frac{1-\gamma}{\gamma} \varepsilon, \quad (45)$$

$$\|B_{\theta_k}^L v - B_{\theta_k}^{L-1} v\|_{\infty} < \frac{1-\gamma}{\gamma} \varepsilon \quad (46)$$

are satisfied, then

$$\|F(x, z; \theta_k) - A_{\theta_k}^L f(x, z; \theta_k)\|_{\infty} < \varepsilon, \quad (47)$$

$$\|V(x, z; \theta_k) - B_{\theta_k}^L v(x, z; \theta_k)\|_{\infty} < \varepsilon \quad (48)$$

are satisfied.

Proof. See Appendix C.3. \square

T_{\max} is always constant for every E step, $T_{\max} = \lceil (\log(1 - \gamma)\varepsilon) / \log \gamma - 1 \rceil$. Thus, even if the approximation errors of $F(x, z; \theta_k)$ and $V(x, z; \theta_k)$ are smaller than ε when $t \ll T_{\max}$, the forward-backward algorithm cannot be terminated until $t = T_{\max}$ because the approximation errors of $F(x, z; \theta_k)$ and $V(x, z; \theta_k)$ cannot be evaluated in the forward-backward algorithm.

L_{\max} is adaptively determined depending on $A_{\theta_k}^L f(x, z)$ and $B_{\theta_k}^L v(x, z)$. Thus, if $A_{\theta_k}^L f(x, z)$ and $B_{\theta_k}^L v(x, z)$ are close enough to $F(x, z; \theta_k)$ and $V(x, z; \theta_k)$, the E step of MBEM can be terminated because the approximation errors of $F(x, z; \theta_k)$ and $V(x, z; \theta_k)$ can be evaluated owing to the contractibility of the forward and backward Bellman operators.

Indeed, when $f(x, z) = p_0(x, z; \nu_k)$ and $v(x, z) = \bar{r}(x, z; \pi_k)$, MBEM is more efficient than EM as follows:

Proposition 6. When $f(x, z) = p_0(x, z; \nu_k)$ and $v(x, z) = \bar{r}(x, z; \pi_k)$, $L_{\max} \leq T_{\max}$ is satisfied.

Proof. See Appendix C.4. \square

The second property of MBEM is the arbitrariness of the initial functions, $f(x, z)$ and $v(x, z)$. In MBEM, the initial functions, $f(x, z)$ and $v(x, z)$, converge to the fixed points, $F(x, z; \theta_k)$ and $V(x, z; \theta_k)$, by applying the forward and backward Bellman operators, A_{θ_k} and B_{θ_k} , L_{\max} times. Therefore, if the initial functions, $f(x, z)$ and $v(x, z)$, are close to the fixed points, $F(x, z; \theta_k)$ and $V(x, z; \theta_k)$, L_{\max} can be reduced. Then, the problem is what kind of the initial functions are close to the fixed points.

We suggest that $F(x, z; \theta_{k-1})$ and $V(x, z; \theta_{k-1})$ are set as the initial functions, $f(x, z)$ and $v(x, z)$. In most cases, θ_{k-1} is close to θ_k . When θ_{k-1} is close to θ_k , it is expected that $F(x, z; \theta_{k-1})$ and $V(x, z; \theta_{k-1})$ are close to $F(x, z; \theta_k)$ and $V(x, z; \theta_k)$. Therefore, by setting $F(x, z; \theta_{k-1})$ and $V(x, z; \theta_{k-1})$ as the initial functions $f(x, z)$ and $v(x, z)$, respectively, L_{\max} is expected to be reduced. Hence, MBEM can be more efficient than EM because MBEM can utilize the results of the previous iteration, $F(x, z; \theta_{k-1})$ and $V(x, z; \theta_{k-1})$, by this arbitrariness of the initial functions.

However, it is unclear how small L_{\max} can be compared to T_{\max} by setting $F(x, z; \theta_{k-1})$ and $V(x, z; \theta_{k-1})$ as the initial functions $f(x, z)$ and $v(x, z)$. Therefore, numerical evaluations are needed. Moreover, in the first iteration, we cannot use the results of the previous iteration, $F(x, z; \theta_{k-1})$ and $V(x, z; \theta_{k-1})$. Therefore, in the first iteration, we set $f(x, z) = p_0(x, z; \nu_k)$ and $v(x, z) = \bar{r}(x, z; \pi_k)$ because these initial functions guarantee $L_{\max} \leq T_{\max}$ from Proposition 6.

MBEM is summarized as Algorithm 3. The M step of Algorithm 3 is exactly the same as that of Algorithms 1 and 2, and only the E step is different. The time complexity of the E step in MBEM is $\mathcal{O}((|\mathcal{X}||\mathcal{Z}|)^2 L_{\max})$. MBEM does not use the inverse matrix, which resolves the drawback of BEM. Moreover, MBEM can reduce L_{\max} by the contractibility of the Bellman operators and the arbitrariness of the initial functions, which can resolve the drawback of EM.

Algorithm 3 Modified Bellman EM algorithm (MBEM)

```

 $k \leftarrow 0$ , Initialize  $\theta_k$ .
 $F(x, z; \theta_{k-1}) \leftarrow p_0(x, z; \nu_k)$ 
 $V(x, z; \theta_{k-1}) \leftarrow \tilde{r}(x, z; \pi_k)$ 
while  $\theta_k$  or  $J(\theta_k)$  do not converge do
  Calculate  $p(x', z' | x, z; \theta_k)$  by Equation (20).
  //—E step—//
   $F_0(x, z) \leftarrow F(x, z; \theta_{k-1})$ 
   $V_0(x, z) \leftarrow V(x, z; \theta_{k-1})$ 
   $L \leftarrow 0$ 
  repeat
     $F_{L+1}(x, z) \leftarrow A_{\theta_k} F_L(x, z)$ 
     $V_{L+1}(x, z) \leftarrow B_{\theta_k} V_L(x, z)$ 
     $L \leftarrow L + 1$ 
  until  $\max\{\|F_L - F_{L-1}\|_1, \|V_L - V_{L-1}\|_\infty\} < \frac{1-\gamma}{\gamma} \varepsilon$ 
  //—M step—//
  Update  $\theta_k$  to  $\theta_{k+1}$  by Equations (9)–(11).
   $k \leftarrow k + 1$ 
end while
return  $\theta_k$ 

```

6. Summary of EM, BEM, and MBEM

EM, BEM, and MBEM are summarized as in Table 1. The M step is exactly the same among these algorithms, and only the E step is different:

- EM obtains $F(x, z; \theta_k)$ and $V(x, z; \theta_k)$ by calculating the forward–backward algorithm up to T_{\max} . T_{\max} needs to be large enough to reduce the approximation errors of $F(x, z; \theta_k)$ and $V(x, z; \theta_k)$, which impairs the computational efficiency.
- BEM obtains $F(x, z; \theta_k)$ and $V(x, z; \theta_k)$ by solving the forward and backward Bellman equations. BEM can be more efficient than EM because BEM calculates the forward and backward Bellman equations instead of the forward–backward algorithm up to T_{\max} . However, BEM cannot always be more efficient than EM when the size of the state $|\mathcal{X}|$ or that of the memory $|\mathcal{Z}|$ is large because BEM calculates an inverse matrix to solve the forward and backward Bellman equations.
- MBEM obtains $F(x, z; \theta_k)$ and $V(x, z; \theta_k)$ by applying the forward and backward Bellman operators, A_{θ_k} and B_{θ_k} , to the initial functions, $f(x, z)$ and $v(x, z)$, L_{\max} times. Since MBEM does not need to calculate the inverse matrix, MBEM may be more efficient than EM even when the size of problems is large, which resolves the drawback of BEM. Although L_{\max} needs to be large enough to reduce the approximation errors of $F(x, z; \theta_k)$ and $V(x, z; \theta_k)$, which is the same problem as EM, MBEM can evaluate the approximation errors more tightly owing to the contractibility of A_{θ_k} and B_{θ_k} , and can utilize the results of the previous iteration, $F(x, z; \theta_{k-1})$ and $V(x, z; \theta_{k-1})$, as the initial functions, $f(x, z)$ and $v(x, z)$. These properties enable MBEM to be more efficient than EM.

Table 1. Summary of EM, BEM, and MBEM.

	EM	Bellman EM (BEM)	Modified Bellman EM (MBEM)
E step	forward–backward algorithm $F(x, z; \theta_k) \approx \sum_{t=0}^{T_{\max}} \gamma^t \alpha_t(x, z; \theta_k)$ $V(x, z; \theta_k) \approx \sum_{t=0}^{T_{\max}} \gamma^t \beta_t(x, z; \theta_k)$ $\mathcal{O}((\mathcal{X} \mathcal{Z})^2 T_{\max})$	forward and backward Bellman equations $F(\theta_k) = (I - \gamma P(\theta_k))^{-1} p(\nu_k)$ $V(\theta_k) = ((I - \gamma P(\theta_k))^{-1})^T r(\pi_k)$ $\mathcal{O}((\mathcal{X} \mathcal{Z})^3)$	forward and backward Bellman operators $F(x, z; \theta_k) \approx A_{\theta_k}^{L_{\max}} f(x, z)$ $V(x, z; \theta_k) \approx B_{\theta_k}^{L_{\max}} v(x, z)$ $\mathcal{O}((\mathcal{X} \mathcal{Z})^2 L_{\max})$
M step	Equations (9)–(11) $\mathcal{O}((\mathcal{X} \mathcal{Z})^2 \mathcal{Y} \mathcal{A})$	Equations (9)–(11) $\mathcal{O}((\mathcal{X} \mathcal{Z})^2 \mathcal{Y} \mathcal{A})$	Equations (9)–(11) $\mathcal{O}((\mathcal{X} \mathcal{Z})^2 \mathcal{Y} \mathcal{A})$

7. Numerical Experiment

In this section, we compare the performance of EM, BEM, and MBEM using numerical experiments of four benchmarks for DEC-POMDP: broadcast [40], recycling robot [15], wireless network [27], and box pushing [41]. Detailed settings such as the state transition probability, the observation probability, and the reward function are described at http://masplan.org/problem_domains, accessed on 22 June 2020. We implement EM, BEM, and MBEM in C++.

Figure 3 shows the experimental results. In all the experiments, we set the number of agent $N = 2$, the discount factor $\gamma = 0.99$, the upper bound of the approximation error $\varepsilon = 0.1$, and the size of the memory available to the i th agent $|Z^i| = 2$. The size of the state $|X|$, the action $|A^i|$, and the observation $|Y^i|$ are different for each problem, which are shown on each panel. We note that the size of the state $|X|$ is small in the broadcast (a,e,i) and the recycling robot (b,f,j), whereas it is large in the wireless network (c,g,k) and the box pushing (d,h,l).

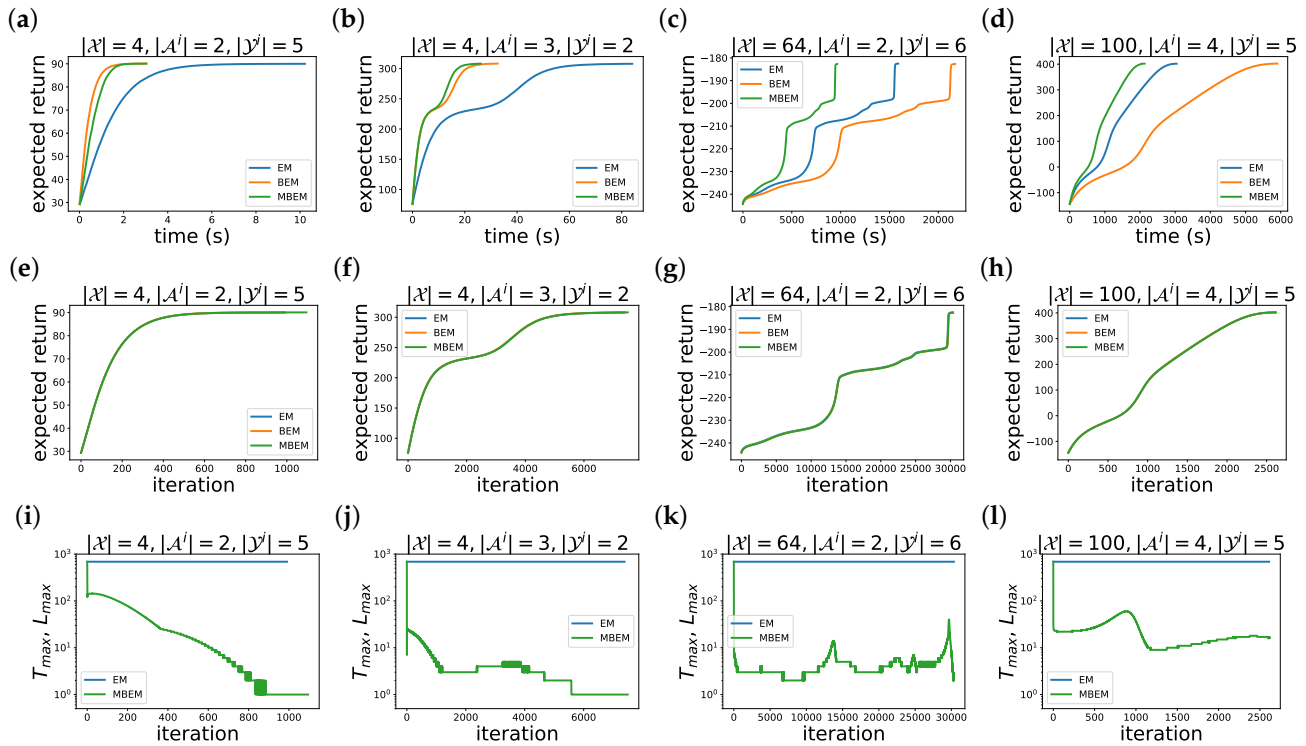


Figure 3. Experimental results of four benchmarks for DEC-POMDP: (a,e,i) broadcast; (b,f,j) recycling robot; (c,g,k) wireless network; (d,h,l) box pushing. (a–d) The expected return $J(\theta_k)$ as a function of the computational time. (e–h) The expected return $J(\theta_k)$ as a function of the iteration k . (i–l) T_{\max} and L_{\max} as functions of the iteration k . In all the experiments, we set the number of agent $N = 2$, the discount factor $\gamma = 0.99$, the upper bound of the approximation error $\varepsilon = 0.1$, and the size of the memory available to the i th agent $|Z^i| = 2$. The size of the state $|X|$, the action $|A^i|$, and the observation $|Y^i|$ are different for each problem, which are shown on each panel.

While the expected return $J(\theta_k)$ with respect to the computational time is different between the algorithms (a–d), that with respect to the iteration k is almost the same (e–h). This is because the M step of these algorithms is exactly the same. Therefore, the difference of the computational time is caused by the computational time of the E step.

The convergence of BEM is faster than that of EM in the small state size problems, i.e., Figure 3a,b. This is because EM calculates the forward–backward algorithm from $t = 0$ to $t = T_{\max}$, where T_{\max} is large. On the other hand, the convergence of BEM is slower than that of EM in the large state size problems, i.e., Figure 3c,d. This is because BEM calculates the inverse matrix.

The convergence of MBEM is faster than that of EM in all the experiments in Figure 3a–d. This is because L_{\max} is smaller than T_{\max} as shown in Figure 3i–l. While EM requires about 1000 calculations of the forward–backward algorithm to guarantee that the approximation error of $F(x, z; \theta_k)$ and $V(x, z; \theta_k)$ is smaller than ε , MBEM requires only about 10 calculations of the forward and backward Bellman operators. Thus, MBEM is more efficient than EM. The reason why L_{\max} is smaller than T_{\max} is that MBEM can utilize the results of the previous iteration, $F(x, z; \theta_{k-1})$ and $V(x, z; \theta_{k-1})$, as the initial functions, $f(x, z)$ and $v(x, z)$. It is shown from L_{\max} and T_{\max} in the first iteration. In the first iteration $k = 0$, L_{\max} is almost the same with T_{\max} because $F(x, z; \theta_{k-1})$ and $V(x, z; \theta_{k-1})$ cannot be used as the initial functions $f(x, z)$ and $v(x, z)$ in the first iteration. On the other hand, in the subsequent iterations $k \geq 1$, L_{\max} is much smaller than T_{\max} because MBEM can utilize the results of the previous iteration, $F(x, z; \theta_{k-1})$ and $V(x, z; \theta_{k-1})$, the initial functions $f(x, z)$ and $v(x, z)$.

8. Conclusions and Future Works

In this paper, we propose the Bellman EM algorithm (BEM) and the modified Bellman EM algorithm (MBEM) by introducing the forward and backward Bellman equations into the EM algorithm for DEC-POMDP. BEM can be more efficient than EM because BEM does not calculate the forward–backward algorithm up to the infinite horizon. However, BEM cannot always be more efficient than EM when the size of the state or that of the memory is large because BEM calculates the inverse matrix. MBEM can be more efficient than EM regardless of the size of problems because MBEM does not calculate the inverse matrix. Although MBEM needs to calculate the forward and backward Bellman operators infinite times, MBEM can evaluate the approximation errors more tightly owing to the contractibility of these operators, and can utilize the results of the previous iteration owing to the arbitrariness of initial functions, which enables MBEM to be more efficient than EM. We verified this theoretical evaluation by the numerical experiment, which demonstrates that the convergence of MBEM is much faster than that of EM regardless of the size of problems.

Our algorithms still leave room for further improvements that deal with the real-world problems, which often have a large discrete or continuous state space. Some of them may be addressed by the advanced techniques of the Bellman equations [1–4]. For example, MBEM may be accelerated by the Gauss–Seidel method [2]. The convergence rate of the E step of MBEM is given by the discount factor γ , which is the same as that of EM. However, the Gauss–Seidel method modifies the Bellman operators, which allows the convergence rate of MBEM to be smaller than the discount factor γ . Therefore, even if $F(x, z; \theta_{k-1})$ and $V(x, z; \theta_{k-1})$ are not close to $F(x, z; \theta_k)$ and $V(x, z; \theta_k)$, MBEM may be more efficient than EM by the Gauss–Seidel method. Moreover, in DEC-POMDP with a large discrete or continuous state space, $F(x, z; \theta_k)$ and $V(x, z; \theta_k)$ cannot be expressed exactly because it requires a large space complexity. This problem may be resolved by the value function approximation [34–36]. The value function approximation approximates $F(x, z; \theta_k)$ and $V(x, z; \theta_k)$ using parametric models such as neural networks. The problem is how to find the optimal approximate parameters. The value function approximation finds them by the Bellman equation. Therefore, the potential extensions of our algorithms may lead to the applications to the real-world DEC-POMDP problems.

Author Contributions: Conceptualization, T.T., T.J.K.; Formal analysis, T.T., T.J.K.; Funding acquisition, T.J.K.; Writing—original draft, T.T., T.J.K. All authors have read and agreed to the published version of the manuscript.

Funding: This research is supported by JSPS KAKENHI Grant Number 19H05799 and by JST CREST Grant Number JPMJCR2011.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: Not applicable.

Acknowledgments: We would like to thank the lab members for a fruitful discussion.

Conflicts of Interest: The authors declare no conflict of interest.

Appendix A. Proof in Section 3

Appendix A.1. Proof of Theorem 1

$J(\theta)$ can be calculated as follows:

$$\begin{aligned}
 J(\theta) &:= \mathbb{E}_{\theta} \left[\sum_{t=0}^{\infty} \gamma^t r(x_t, a_t) \right] \\
 &= \sum_{x_{0:\infty}, a_{0:\infty}} p(x_{0:\infty}, a_{0:\infty}; \theta) \sum_{t=0}^{\infty} \gamma^t r(x_t, a_t) = \sum_{t=0}^{\infty} \gamma^t \sum_{x_t, a_t} p(x_t, a_t; \theta) r(x_t, a_t) \\
 &= (1 - \gamma)^{-1} \sum_{T=0}^{\infty} p(T) \sum_{x_T, a_T} p(x_T, a_T; \theta) [(r_{\max} - r_{\min}) p(o = 1 | x_T, a_T) + r_{\min}] \\
 &= (1 - \gamma)^{-1} [(r_{\max} - r_{\min}) p(o = 1; \theta) + r_{\min}] \quad (\text{A1})
 \end{aligned}$$

where $x_{0:\infty} := \{x_0, \dots, x_{\infty}\}$ and $a_{0:\infty} := \{a_0, \dots, a_{\infty}\}$.

Appendix A.2. Proof of Proposition 1

In order to prove Proposition 1, we calculate $Q(\theta; \theta_k)$. It can be calculated as follows:

$$\begin{aligned}
 Q(\theta; \theta_k) &:= \mathbb{E}_{\theta_k} [\log p(o = 1, x_{0:T}, y_{0:T}, z_{0:T}, a_{0:T}, T; \theta) | o = 1] \\
 &= \mathbb{E}_{\theta_k} \left[\sum_{i=1}^N \left\{ \sum_{t=0}^T \log \pi^i(a_t^i | z_t^i) + \sum_{t=1}^T \log \lambda^i(z_t^i | z_{t-1}^i, y_t^i) + \log \nu^i(z_0^i) \right\} \middle| o = 1 \right] + C \\
 &= \sum_{i=1}^N \left(Q(\pi^i; \theta_k) + Q(\lambda^i; \theta_k) + Q(\nu^i; \theta_k) \right) + C, \quad (\text{A2})
 \end{aligned}$$

where

$$Q(\pi^i; \theta_k) := \mathbb{E}_{\theta_k} \left[\sum_{t=0}^T \log \pi^i(a_t^i | z_t^i) \middle| o = 1 \right], \quad (\text{A3})$$

$$Q(\lambda^i; \theta_k) := \mathbb{E}_{\theta_k} \left[\sum_{t=1}^T \log \lambda^i(z_t^i | z_{t-1}^i, y_t^i) \middle| o = 1 \right], \quad (\text{A4})$$

$$Q(\nu^i; \theta_k) := \mathbb{E}_{\theta_k} [\log \nu^i(z_0^i) | o = 1]. \quad (\text{A5})$$

C is a constant independent of θ .

Proposition A1 ([16]). $Q(\pi^i; \theta_k)$, $Q(\lambda^i; \theta_k)$, and $Q(\nu^i; \theta_k)$ are calculated as follows:

$$Q(\pi^i; \theta_k) \propto \sum_{a, z} \pi_k(a | z) \log \pi^i(a^i | z^i) \sum_{x, x', z'} p(x', z' | x, z, a; \lambda_k) F(x, z; \theta_k) (\bar{r}(x, a) + \gamma V(x', z'; \theta_k)), \quad (\text{A6})$$

$$Q(\lambda^i; \theta_k) \propto \sum_{z', z, y'} \lambda_k(z' | z, y') \log \lambda^i(z'^i | z^i, y'^i) \sum_{x', x} p(x', y' | x, z; \pi_k) F(x, z; \theta_k) V(x', z'; \theta_k), \quad (\text{A7})$$

$$Q(\nu^i; \theta_k) \propto \sum_z p_0(z; \nu_n) \log \nu^i(z^i) \sum_x p_0(x) V(x, z; \theta_k). \quad (\text{A8})$$

$F(x, z; \theta_k)$ and $V(x, z; \theta_k)$ are defined by Equations (12) and (13).

Proof. Firstly, we prove Equation (A6). $Q(\pi^i; \theta_k)$ can be calculated as follows:

$$\begin{aligned}
Q(\pi^i; \theta_k) &:= \mathbb{E}_{\theta_k} \left[\sum_{t=0}^T \log \pi^i(a_t^i | z_t^i) \middle| o = 1 \right] \\
&= \frac{1}{p(o = 1; \theta_k)} \sum_{T=0}^{\infty} \sum_{x_{0:T}, y_{0:T}, z_{0:T}, a_{0:T}} p(o = 1, x_{0:T}, y_{0:T}, z_{0:T}, a_{0:T}; \theta_k) \sum_{t=0}^T \log \pi^i(a_t^i | z_t^i) \\
&= \frac{1}{p(o = 1; \theta_k)} \sum_{T=0}^{\infty} \sum_{t=0}^T \sum_{x_t, z_t, a_t} p(o = 1, x_t, z_t, a_t; \theta_k) \log \pi^i(a_t^i | z_t^i) \\
&= \frac{1}{p(o = 1; \theta_k)} \sum_{T=0}^{\infty} \sum_{t=0}^T \sum_{x_t, z_t, a_t} p(T) p_t(x_t, z_t; \theta_k) \pi_k(a_t | z_t) p_t(o = 1 | x_t, z_t, a_t; \theta_k) \log \pi^i(a_t^i | z_t^i) \\
&= \frac{1}{p(o = 1; \theta_k)} \sum_{T=0}^{\infty} \sum_{t=0}^T \sum_{x, z, a} p(T) p_t(x, z; \theta_k) \pi_k(a | z) p_t(o = 1 | x, z, a; \theta_k) \log \pi^i(a^i | z^i) \\
&= \frac{1}{p(o = 1; \theta_k)} \sum_{x, z, a} \pi_k(a | z) \log \pi^i(a^i | z^i) \sum_{T=0}^{\infty} p(T) \sum_{t=0}^T p_t(x, z; \theta_k) p_t(o = 1 | x, z, a; \theta_k). \tag{A9}
\end{aligned}$$

Since $\sum_{T=0}^{\infty} \sum_{t=0}^T \dots = \sum_{t=0}^{\infty} \sum_{T=t}^{\infty} \dots$, we have

$$\begin{aligned}
Q(\pi^i; \theta_k) &= \frac{1}{p(o = 1; \theta_k)} \sum_{x, z, a} \pi_k(a | z) \log \pi^i(a^i | z^i) \sum_{t=0}^{\infty} p_t(x, z; \theta_k) \sum_{T=t}^{\infty} p(T) p_t(o = 1 | x, z, a; \theta_k) \\
&= \frac{1 - \gamma}{p(o = 1; \theta_k)} \sum_{x, z, a} \pi_k(a | z) \log \pi^i(a^i | z^i) \sum_{t=0}^{\infty} p_t(x, z; \theta_k) \sum_{T=t}^{\infty} \gamma^T p_t(o = 1 | x, z, a; \theta_k). \tag{A10}
\end{aligned}$$

Since $p_t(o = 1 | x, z, a; \theta_k) = p_0(o = 1 | x, z, a; T - t; \theta_k)$,

$$\begin{aligned}
Q(\pi^i; \theta_k) &= \frac{1 - \gamma}{p(o = 1; \theta_k)} \sum_{x, z, a} \pi_k(a | z) \log \pi^i(a^i | z^i) \sum_{t=0}^{\infty} p_t(x, z; \theta_k) \sum_{T=t}^{\infty} \gamma^T p_0(o = 1 | x, z, a; T - t; \theta_k) \\
&= \frac{1 - \gamma}{p(o = 1; \theta_k)} \sum_{x, z, a} \pi_k(a | z) \log \pi^i(a^i | z^i) \sum_{t=0}^{\infty} \gamma^t p_t(x, z; \theta_k) \sum_{T=0}^{\infty} \gamma^T p_0(o = 1 | x, z, a; T; \theta_k) \\
&= \frac{1 - \gamma}{p(o = 1; \theta_k)} \sum_{x, z, a} \pi_k(a | z) \log \pi^i(a^i | z^i) F(x, z; \theta_k) \sum_{T=0}^{\infty} \gamma^T p_0(o = 1 | x, z, a; T; \theta_k). \tag{A11}
\end{aligned}$$

$\sum_{T=0}^{\infty} \gamma^T p_0(o = 1 | x, z, a; T; \theta_k)$ can be calculated as follows:

$$\begin{aligned}
\sum_{T=0}^{\infty} \gamma^T p_0(o = 1 | x, z, a; T; \theta_k) &= \bar{r}(x, a) + \sum_{T=1}^{\infty} \gamma^T p_0(o = 1 | x, z, a; T; \theta_k) \\
&= \bar{r}(x, a) + \sum_{T=1}^{\infty} \gamma^T \sum_{x', z'} p_1(o = 1 | x', z'; T; \theta_k) p(x', z' | x, z, a; \lambda_k) \\
&= \bar{r}(x, a) + \gamma \sum_{x', z'} p(x', z' | x, z, a; \lambda_k) \sum_{T=0}^{\infty} \gamma^T p_0(o = 1 | x', z'; T; \theta_k) \\
&= \bar{r}(x, a) + \gamma \sum_{x', z'} p(x', z' | x, z, a; \lambda_k) V(x', z'; \theta_k). \tag{A12}
\end{aligned}$$

From Equations (A11) and (A12), we obtain

$$Q(\pi^i; \theta_k) = \frac{1 - \gamma}{p(o = 1; \theta_k)} \sum_{a, z} \pi_k(a | z) \log \pi^i(a^i | z^i) \sum_{x, x', z'} p(x', z' | x, z, a; \lambda_k) F(x, z; \theta_k) (\bar{r}(x, a) + \gamma V(x', z'; \theta_k)). \tag{A13}$$

Therefore, Equation (A6) is proved.

Secondly, we prove Equation (A7). $Q(\lambda^i; \theta_k)$ can be calculated as follows:

$$\begin{aligned}
Q(\lambda^i; \theta_k) &:= \mathbb{E}_{\theta_k} \left[\sum_{t=1}^T \log \lambda^i(z_t^i | z_{t-1}^i, y_t^i) \middle| o = 1 \right] \\
&= \frac{1}{p(o = 1; \theta_k)} \sum_{T=1}^{\infty} \sum_{x_{0:T}, y_{0:T}, z_{0:T}, a_{0:T}} p(o = 1, x_{0:T}, y_{0:T}, z_{0:T}, a_{0:T}, T; \theta_k) \sum_{t=1}^T \log \lambda^i(z_t^i | z_{t-1}^i, y_t^i) \\
&= \frac{1}{p(o = 1; \theta_k)} \sum_{T=1}^{\infty} \sum_{t=1}^T \sum_{x_{t-1:t}, y_t, z_{t-1:t}} p(o = 1, x_{t-1:t}, y_t, z_{t-1:t}, T; \theta_k) \log \lambda^i(z_t^i | z_{t-1}^i, y_t^i) \\
&= \frac{1}{p(o = 1; \theta_k)} \sum_{T=1}^{\infty} \sum_{t=1}^T \sum_{x_{t-1:t}, y_t, z_{t-1:t}} p(T) p_{t-1}(x_{t-1}, z_{t-1}; \theta_k) p(x_t, y_t | x_{t-1}, z_{t-1}; \pi_k) \\
&\quad \times \lambda_k(z_t | z_{t-1}, y_t) p_t(o = 1 | x_t, z_t, T; \theta_k) \log \lambda^i(z_t^i | z_{t-1}^i, y_t^i) \\
&= \frac{1}{p(o = 1; \theta_k)} \sum_{x, x', y', z, z'} \lambda_k(z' | z, y') \log \lambda^i(z'^i | z^i, y'^i) p(x', y' | x, z; \pi_k) \\
&\quad \times \sum_{T=1}^{\infty} p(T) \sum_{t=1}^T p_{t-1}(x, z; \theta_k) p_t(o = 1 | x', z', T; \theta_k). \tag{A14}
\end{aligned}$$

Since $\sum_{T=1}^{\infty} \sum_{t=1}^T \dots = \sum_{t=1}^{\infty} \sum_{T=t}^{\infty} \dots$, we have

$$\begin{aligned}
Q(\lambda^i; \theta_k) &= \frac{1}{p(o = 1; \theta_k)} \sum_{x, x', y', z, z'} \lambda_k(z' | z, y') \log \lambda^i(z'^i | z^i, y'^i) p(x', y' | x, z; \pi_k) \\
&\quad \times \sum_{t=1}^{\infty} p_{t-1}(x, z; \theta_k) \sum_{T=t}^{\infty} p(T) p_t(o = 1 | x', z', T; \theta_k) \\
&= \frac{1 - \gamma}{p(o = 1; \theta_k)} \sum_{x, x', y', z, z'} \lambda_k(z' | z, y') \log \lambda^i(z'^i | z^i, y'^i) p(x', y' | x, z; \pi_k) \\
&\quad \times \sum_{t=1}^{\infty} p_{t-1}(x, z; \theta_k) \sum_{T=t}^{\infty} \gamma^T p_t(o = 1 | x', z', T; \theta_k). \tag{A15}
\end{aligned}$$

Since $p_t(o = 1, x, z, T; \theta_k) = p_0(o = 1 | x, z, T - t; \theta_k)$,

$$\begin{aligned}
Q(\lambda^i; \theta_k) &= \frac{1 - \gamma}{p(o = 1; \theta_k)} \sum_{x, x', y', z, z'} \lambda_k(z' | z, y') \log \lambda^i(z'^i | z^i, y'^i) p(x', y' | x, z; \pi_k) \\
&\quad \times \sum_{t=1}^{\infty} p_{t-1}(x, z; \theta_k) \sum_{T=t}^{\infty} \gamma^T p_0(o = 1 | x', z', T - t; \theta_k) \\
&= \frac{(1 - \gamma)\gamma}{p(o = 1; \theta_k)} \sum_{x, x', y', z, z'} \lambda_k(z' | z, y') \log \lambda^i(z'^i | z^i, y'^i) p(x', y' | x, z; \pi_k) \\
&\quad \times \sum_{t=0}^{\infty} \gamma^t p_t(x, z; \theta_k) \sum_{T=0}^{\infty} \gamma^T p_0(o = 1 | x', z', T; \theta_k) \\
&= \frac{(1 - \gamma)\gamma}{p(o = 1; \theta_k)} \sum_{x, x', y', z, z'} \lambda_k(z' | z, y') \log \lambda^i(z'^i | z^i, y'^i) p(x', y' | x, z; \pi_k) F(x, z; \theta_k) V(x', z'; \theta_k). \tag{A16}
\end{aligned}$$

Therefore, Equation (A7) is proved.

Finally, we prove Equation (A8). $Q(v^i; \theta_k)$ can be calculated as follows:

$$\begin{aligned}
Q(v^i; \theta_k) &:= \mathbb{E}_{\theta_k} \left[\log v^i(z_0^i) \middle| o = 1 \right] \\
&= \frac{1}{p(o = 1; \theta_k)} \sum_{T=0}^{\infty} \sum_{x_{0:T}, y_{0:T}, z_{0:T}, a_{0:T}} p(o = 1, x_{0:T}, y_{0:T}, z_{0:T}, a_{0:T}; \theta_k) \log v^i(z_0^i) \\
&= \frac{1}{p(o = 1; \theta_k)} \sum_{T=0}^{\infty} \sum_{x_0, z_0} p(o = 1, x_0, z_0, T; \theta_k) \log v^i(z_0^i) \\
&= \frac{1}{p(o = 1; \theta_k)} \sum_{T=0}^{\infty} \sum_{x_0, z_0} p(T) p_0(x_0) v_k(z_0) p_0(o = 1 | x_0, z_0, T; \theta_k) \log v^i(z_0^i) \\
&= \frac{1}{p(o = 1; \theta_k)} \sum_{T=0}^{\infty} \sum_{x, z} p(T) p_0(x) v_k(z) p_0(o = 1 | x, z, T; \theta_k) \log v^i(z^i) \\
&= \frac{1}{p(o = 1; \theta_k)} \sum_z v_k(z) \log v^i(z^i) \sum_x p_0(x) \sum_{T=0}^{\infty} p(T) p_0(o = 1 | x, z, T; \theta_k) \\
&= \frac{1 - \gamma}{p(o = 1; \theta_k)} \sum_z v_k(z) \log v^i(z^i) \sum_x p_0(x) V(x, z; \theta_k). \tag{A17}
\end{aligned}$$

Therefore, Equation (A8) is proved. \square

Equations (9)–(11) can be calculated from Equations (A6)–(A8) using the Lagrange multiplier method [16]. Therefore, Proposition 1 is proved.

Appendix A.3. Proof of Proposition 2

The left-hand side of Equation (27) can be calculated as follows:

$$\begin{aligned}
&\left\| F(x, z; \theta_k) - \sum_{t=0}^{T_{\max}} \gamma^t \alpha_t(x, z; \theta_k) \right\|_{\infty} \\
&= \left\| \sum_{t=T_{\max}+1}^{\infty} \gamma^t \alpha_t(x, z; \theta_k) \right\|_{\infty} \leq \sum_{t=T_{\max}+1}^{\infty} \gamma^t \|\alpha_t(x, z; \theta_k)\|_{\infty} = \sum_{t=T_{\max}+1}^{\infty} \gamma^t = \frac{\gamma^{T_{\max}+1}}{1 - \gamma}. \tag{A18}
\end{aligned}$$

Therefore, if

$$\frac{\gamma^{T_{\max}+1}}{1 - \gamma} < \varepsilon \tag{A19}$$

is satisfied, Equation (27) is satisfied. From Equation (A19), we obtain Equation (26). Therefore, Equation (26) \Rightarrow Equation (27) is proved. Equation (26) \Rightarrow Equation (28) can be proved in the same way.

Appendix B. Proof in Section 4

Proof of Theorem 2

$F(x, z; \theta)$ can be calculated as follows:

$$\begin{aligned}
F(x, z; \theta) &:= \sum_{t=0}^{\infty} \gamma^t p_t(x, z; \theta) \\
&= p_0(x, z; \nu) + \sum_{t=1}^{\infty} \gamma^t p_t(x, z; \theta) \\
&= p_0(x, z; \nu) + \sum_{t=1}^{\infty} \gamma^t \sum_{x', z'} p(x, z | x', z'; \theta) p_{t-1}(x', z'; \theta) \\
&= p_0(x, z; \nu) + \gamma \sum_{x', z'} p(x, z | x', z'; \theta) \sum_{t=1}^{\infty} \gamma^{t-1} p_{t-1}(x', z'; \theta)
\end{aligned}$$

$$= p_0(x, z; \nu) + \gamma \sum_{x', z'} p(x, z | x', z'; \theta) F(x', z'; \theta). \quad (\text{A20})$$

Therefore, Equation (29) is proved. Equation (30) can be proved in the same way.

Appendix C. Proof in Section 5

Appendix C.1. Proof of Proposition 3

The left-hand side of Equation (37) can be calculated as follows:

$$\begin{aligned} \|A_{\theta}f - A_{\theta}g\|_1 &= \gamma \|P(\theta)f - P(\theta)g\|_1 \\ &\leq \gamma \|P(\theta)\|_1 \|f - g\|_1 \\ &= \gamma \|f - g\|_1. \end{aligned} \quad (\text{A21})$$

where $P_{ij}(\theta) := p((x', z') = i | (x, z) = j; \theta)$, $f_i := f((x, z) = i)$, and $g_i := g((x, z) = i)$. Equation (38) can be proved in the same way.

Appendix C.2. Proof of Proposition 4

We prove Equation (39) by showing $\lim_{L \rightarrow \infty} \|F(\theta) - A_{\theta}^L f\|_1 = 0$. Since $F(x, z; \theta)$ is the fixed point of A_{θ} ,

$$\|F(\theta) - A_{\theta}^L f\|_1 = \|A_{\theta}^L F - A_{\theta}^L f\|_1. \quad (\text{A22})$$

From the contractibility of A_{θ} ,

$$\|A_{\theta}^L F(\theta) - A_{\theta}^L f\|_1 \leq \gamma^L \|F(\theta) - f\|_1. \quad (\text{A23})$$

Since $\|F(\theta) - f\|_1$ is finite, the right-hand side is 0 when $L \rightarrow \infty$. Hence,

$$\lim_{L \rightarrow \infty} \|F(\theta) - A_{\theta}^L f\|_1 = 0 \quad (\text{A24})$$

is satisfied, and Equation (39) is proved. Equation (40) can be proved in the same way.

Appendix C.3. Proof of Proposition 5

From the definition of the norm,

$$\|F(\theta_k) - A_{\theta_k}^L f\|_{\infty} \leq \|F(\theta_k) - A_{\theta_k}^L f\|_1. \quad (\text{A25})$$

The right-hand side can be calculated as follows:

$$\begin{aligned} \|F(\theta_k) - A_{\theta_k}^L f\|_1 &\leq \|F(\theta_k) - A_{\theta_k}^{L+1} f\|_1 + \|A_{\theta_k}^{L+1} f - A_{\theta_k}^L f\|_1 \\ &= \|A_{\theta_k} F(\theta_k) - A_{\theta_k}^{L+1} f\|_1 + \|A_{\theta_k}^{L+1} f - A_{\theta_k}^L f\|_1 \\ &\leq \gamma \|F(\theta_k) - A_{\theta_k}^L f\|_1 + \gamma \|A_{\theta_k}^L f - A_{\theta_k}^{L-1} f\|_1. \end{aligned} \quad (\text{A26})$$

From this inequality, we have

$$\|F(\theta_k) - A_{\theta_k}^L f\|_1 \leq \frac{\gamma}{1 - \gamma} \|A_{\theta_k}^L f - A_{\theta_k}^{L-1} f\|_1. \quad (\text{A27})$$

Therefore, if $\|A_{\theta_k}^L f - A_{\theta_k}^{L-1} f\|_1 < (1 - \gamma)\gamma^{-1}\varepsilon$ is satisfied,

$$\|F(\theta_k) - A_{\theta_k}^L f\|_{\infty} < \varepsilon \quad (\text{A28})$$

holds. Equation (48) is proved in the same way.

Appendix C.4. Proof of Proposition 6

When $f(x, z) = p_0(x, z; \nu_k)$ and $v(x, z) = \bar{r}(x, z; \pi_k)$, the termination conditions of MBEM can be calculated as follows:

$$\|A_{\theta_k}^L f - A_{\theta_k}^{L-1} f\|_1 \leq \gamma^{L-1} \|A_{\theta_k} f - f\|_1 = \gamma^{L-1} \|\gamma P(\theta_k) p(\nu_k)\|_1 \leq \gamma^L \|P(\theta_k)\|_1 \|p(\nu_k)\|_1 = \gamma^L, \quad (\text{A29})$$

$$\|B_{\theta_k}^L v - B_{\theta_k}^{L-1} v\|_\infty \leq \gamma^L. \quad (\text{A30})$$

The calculation of $\|B_{\theta_k}^L v - B_{\theta_k}^{L-1} v\|_\infty$ is omitted because it is the same as that of $\|A_{\theta_k}^L f - A_{\theta_k}^{L-1} f\|_1$. If

$$\gamma^L < \frac{1-\gamma}{\gamma} \varepsilon \quad \left(\Rightarrow L > \frac{\log(1-\gamma)\varepsilon}{\log \gamma} - 1 \right) \quad (\text{A31})$$

is satisfied,

$$\|A_{\theta_k}^L f - A_{\theta_k}^{L-1} f\|_1 < (1-\gamma)\gamma^{-1}\varepsilon, \quad (\text{A32})$$

$$\|B_{\theta_k}^L v - B_{\theta_k}^{L-1} v\|_\infty < (1-\gamma)\gamma^{-1}\varepsilon \quad (\text{A33})$$

holds. Thus, L_{\max} satisfies

$$L_{\max} \leq \left\lceil \frac{\log(1-\gamma)\varepsilon}{\log \gamma} - 1 \right\rceil. \quad (\text{A34})$$

$\lceil \cdot \rceil : \mathbb{R} \rightarrow \mathbb{Z}$ is defined by $\lceil x \rceil := \min\{n \in \mathbb{Z} | n > x\}$. From Proposition 2, the minimum T_{\max} is given by the following equation:

$$T_{\max} = \left\lceil \frac{\log(1-\gamma)\varepsilon}{\log \gamma} - 1 \right\rceil. \quad (\text{A35})$$

Therefore, $L_{\max} \leq T_{\max}$ holds.

Appendix D. A Note on the Algorithm Proposed by Song et al.

In this section, we show that a parameter dependency is overlooked in the algorithm of [32]. We outline the derivation of the algorithm in [32] and discuss the parameter dependency. Since we use the notation in this paper, it is recommended to read the full paper before reading this section.

Firstly, we calculate the expected return $J(\theta)$ to derive the algorithm in [32]. Since [32] considers the case where $\nu(z) = \delta_{z, z_0}$, we also consider the same case in this section. The expected return $J(\theta)$ can be calculated as follows:

$$\begin{aligned} J(\theta) &:= \mathbb{E}_{\theta} \left[\sum_{t=0}^{\infty} \gamma^t r(x_t, a_t) \right] \\ &= \sum_{x_0, a_0} p(x_0) \pi(a_0 | z_0) \left[r(x_0, a_0) + \gamma \sum_{x_1, y_1, z_1} p(x_1 | x_0, a_0) p(y_1 | x_1, a_0) \lambda(z_1 | z_0, y_1) V(x_1, z_1; \theta) \right]. \end{aligned} \quad (\text{A36})$$

$V(x, z; \theta)$ is the value function, which is defined as follows:

$$V(x, z; \theta) := \mathbb{E}_{\theta} \left[\sum_{t=0}^{\infty} \gamma^t r(x_t, a_t) \middle| x_0 = x, z_0 = z \right]. \quad (\text{A37})$$

Equation (A36) can be rewritten as:

$$J(\theta) = \sum_{a_0^i} \pi^i(a_0^i|z_0^i) \tilde{r}(z_0^i, a_0^i; \pi^{-i}) + \sum_{a_0^i, y_1^i, z_1^i} \pi^i(a_0^i|z_0^i) \lambda^i(z_1^i|z_0^i, y_1^i) \tilde{V}(z_0^i, a_0^i, y_1^i, z_1^i; \theta) \quad (\text{A38})$$

where

$$\tilde{r}(z_0^i, a_0^i; \pi^{-i}) := \sum_{x_0, a_0^{-i}} p(x_0) \pi^{-i}(a_0^{-i}|z_0^{-i}) r(x_0, a_0), \quad (\text{A39})$$

$$\begin{aligned} \tilde{V}(z_0^i, a_0^i, y_1^i, z_1^i; \theta) \\ := \gamma \sum_{x_0, a_0^{-i}, x_1, y_1^{-i}, z_1^{-i}} p(x_0) \pi^{-i}(a_0^{-i}|z_0^{-i}) p(x_1|x_0, a_0) p(y_1|x_1, a_0) \lambda^{-i}(z_1^{-i}|z_0^{-i}, y_1^{-i}) V(x_1, z_1; \theta). \end{aligned} \quad (\text{A40})$$

Maximizing $J(\theta)$ is equivalent to maximizing $\log J(\theta)$, and $\log J(\theta)$ can be calculated as follows:

$$\begin{aligned} \log J(\theta) = \log \left\{ \sum_{a_0^i} \eta(z_0^i, a_0^i; \theta_k) \frac{\pi^i(a_0^i|z_0^i) \tilde{r}(z_0^i, a_0^i; \pi^{-i})}{\eta(z_0^i, a_0^i; \theta_k)} \right. \\ \left. + \sum_{a_0^i, y_1^i, z_1^i} \rho(z_0^i, a_0^i, y_1^i, z_1^i; \theta_k) \frac{\pi^i(a_0^i|z_0^i) \lambda^i(z_1^i|z_0^i, y_1^i) \tilde{V}(z_0^i, a_0^i, y_1^i, z_1^i; \theta)}{\rho(z_0^i, a_0^i, y_1^i, z_1^i; \theta_k)} \right\} \end{aligned} \quad (\text{A41})$$

where

$$\eta(z_0^i, a_0^i; \theta_k) = \frac{\pi_k^i(a_0^i|z_0^i) \tilde{r}(z_0^i, a_0^i; \pi_k^{-i})}{\xi(z_0^i; \theta_k)}, \quad (\text{A42})$$

$$\rho(z_0^i, a_0^i, y_1^i, z_1^i; \theta_k) = \frac{\pi_k^i(a_0^i|z_0^i) \lambda_k^i(z_1^i|z_0^i, y_1^i) \tilde{V}(z_0^i, a_0^i, y_1^i, z_1^i; \theta_k)}{\xi(z_0^i; \theta_k)}, \quad (\text{A43})$$

$$\xi(z_0^i; \theta_k) = \sum_{a_0^i} \pi_k^i(a_0^i|z_0^i) \tilde{r}(z_0^i, a_0^i; \pi_k^{-i}) + \sum_{a_0^i, y_1^i, z_1^i} \pi_k^i(a_0^i|z_0^i) \lambda_k^i(z_1^i|z_0^i, y_1^i) \tilde{V}(z_0^i, a_0^i, y_1^i, z_1^i; \theta_k). \quad (\text{A44})$$

By the Jensen's inequality, Equation (A41) can be calculated as follows:

$$\begin{aligned} \log J(\theta) \geq \sum_{a_0^i} \eta(z_0^i, a_0^i; \theta_k) \log \left\{ \frac{\pi^i(a_0^i|z_0^i) \tilde{r}(z_0^i, a_0^i; \pi^{-i})}{\eta(z_0^i, a_0^i; \theta_k)} \right\} \\ + \sum_{a_0^i, y_1^i, z_1^i} \rho(z_0^i, a_0^i, y_1^i, z_1^i; \theta_k) \log \left\{ \frac{\pi^i(a_0^i|z_0^i) \lambda^i(z_1^i|z_0^i, y_1^i) \tilde{V}(z_0^i, a_0^i, y_1^i, z_1^i; \theta)}{\rho(z_0^i, a_0^i, y_1^i, z_1^i; \theta_k)} \right\} =: Q(\theta; \theta_k). \end{aligned} \quad (\text{A45})$$

$\log J(\theta) = Q(\theta; \theta_k)$ is satisfied when $\theta = \theta_k$.

Then, θ_{k+1} is defined as follows:

$$\theta_{k+1} := \arg \max_{\theta} Q(\theta; \theta_k) \quad (\text{A46})$$

In this case, the following proposition holds:

Proposition A2 ([32]). $\log J(\theta_{k+1}) \geq \log J(\theta_k)$.

Proof. From Equation (A45), $\log J(\theta_{k+1}) \geq Q(\theta_{k+1}; \theta_k)$. From Equation (A46), $Q(\theta_{k+1}; \theta_k) \geq Q(\theta_k; \theta_k) = \log J(\theta_k)$. Therefore, $\log J(\theta_{k+1}) \geq \log J(\theta_k)$ is satisfied. \square

Therefore, since Equation (A46) monotonically increases $J(\theta)$, we can find θ^* , which locally maximizes $J(\theta)$. This is the algorithm in [32].

Then, the problem is how to calculate Equation (A46). It cannot be calculated analytically because θ dependency of $\tilde{V}(z_0^i, a_0^i, y_1^i, z_1^i; \theta)$ is too complex. However, ref. [32] overlooked the parameter dependency of $\tilde{r}(z_0^i, a_0^i; \pi^{-i})$ and $\tilde{V}(z_0^i, a_0^i, y_1^i, z_1^i; \theta)$, and therefore, it calculated Equation (A46) as follows:

$$\pi_{k+1}^i(a_0^i|z_0^i) = \frac{\eta(z_0^i, a_0^i; \theta_k) + \sum_{y_1^i, z_1^i} \rho(z_0^i, a_0^i, y_1^i, z_1^i; \theta_k)}{\sum_{a_0^i} \left[\eta(z_0^i, a_0^i; \theta_k) + \sum_{y_1^i, z_1^i} \rho(z_0^i, a_0^i, y_1^i, z_1^i; \theta_k) \right]}, \quad (\text{A47})$$

$$\lambda_{k+1}^i(z_1^i|z_0^i, y_1^i) = \frac{\sum_{a_0^i} \rho(z_0^i, a_0^i, y_1^i, z_1^i; \theta_k)}{\sum_{a_0^i, z_1^i} \rho(z_0^i, a_0^i, y_1^i, z_1^i; \theta_k)}. \quad (\text{A48})$$

However, Equations (A47) and (A48) do not correspond to Equation (A46), and therefore, the algorithm as a whole may not always provide the optimal policy.

References

1. Bertsekas, D.P. *Dynamic Programming and Optimal Control: Vol. 1*; Athena Scientific: Belmont, MA, USA, 2000.
2. Puterman, M.L. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*; John Wiley & Sons: Hoboken, NJ, USA, 2014.
3. Sutton, R.S.; Barto, A.G. *Introduction to Reinforcement Learning*; MIT Press: Cambridge, MA, USA, 1998; Volume 135.
4. Sutton, R.S.; Barto, A.G. *Reinforcement Learning: An Introduction*; MIT Press: Cambridge, MA, USA, 2018.
5. Kochenderfer, M.J. *Decision Making under Uncertainty: Theory and Application*; MIT Press: Cambridge, MA, USA, 2015.
6. Oliehoek, F. *Value-Based Planning for Teams of Agents in Stochastic Partially Observable Environments*; Amsterdam University Press: Amsterdam, The Netherlands, 2010.
7. Oliehoek, F.A.; Amato, C. *A Concise Introduction to Decentralized POMDPs*; Springer: Berlin/Heidelberg, Germany, 2016; Volume 1.
8. Becker, R.; Zilberstein, S.; Lesser, V.; Goldman, C.V. Solving transition independent decentralized Markov decision processes. *J. Artif. Intell. Res.* **2004**, *22*, 423–455. [\[CrossRef\]](#)
9. Nair, R.; Varakantham, P.; Tambe, M.; Yokoo, M. Networked distributed POMDPs: A synthesis of distributed constraint optimization and POMDPs. In Proceedings of the AAAI'05: Proceedings of the 20th National Conference on Artificial Intelligence, Pittsburgh, PA, USA, 9–13 July 2005; pp. 133–139.
10. Bernstein, D.S.; Givan, R.; Immerman, N.; Zilberstein, S. The complexity of decentralized control of Markov decision processes. *Math. Oper. Res.* **2002**, *27*, 819–840. [\[CrossRef\]](#)
11. Bernstein, D.S.; Hansen, E.A.; Zilberstein, S. Bounded policy iteration for decentralized POMDPs. In Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence (IJCAI), Edinburgh, UK, 30 July–5 August 2005; pp. 52–57.
12. Bernstein, D.S.; Amato, C.; Hansen, E.A.; Zilberstein, S. Policy iteration for decentralized control of Markov decision processes. *J. Artif. Intell. Res.* **2009**, *34*, 89–132. [\[CrossRef\]](#)
13. Amato, C.; Bernstein, D.S.; Zilberstein, S. Optimizing fixed-size stochastic controllers for POMDPs and decentralized POMDPs. *Auton. Agents Multi-Agent Syst.* **2010**, *21*, 293–320. [\[CrossRef\]](#)
14. Amato, C.; Bonet, B.; Zilberstein, S. Finite-state controllers based on mealy machines for centralized and decentralized pomdps. In Proceedings of the AAAI Conference on Artificial Intelligence, Atlanta, GA, USA, 11–15 July 2010.
15. Amato, C.; Bernstein, D.S.; Zilberstein, S. Optimizing memory-bounded controllers for decentralized POMDPs. *arXiv* **2012**, arXiv:1206.5258.
16. Kumar, A.; Zilberstein, S. Anytime planning for decentralized POMDPs using expectation maximization. In Proceedings of the 26th Conference on Uncertainty in Artificial Intelligence, Catalina Island, CA, USA, 8–11 July 2010; pp. 294–301.
17. Kumar, A.; Zilberstein, S.; Toussaint, M. Probabilistic inference techniques for scalable multiagent decision making. *J. Artif. Intell. Res.* **2015**, *53*, 223–270. [\[CrossRef\]](#)
18. Toussaint, M.; Storkey, A. Probabilistic inference for solving discrete and continuous state Markov Decision Processes. In Proceedings of the 23rd International Conference on Machine Learning, Pittsburgh, PA, USA, 25–29 June 2006; pp. 945–952.
19. Todorov, E. General duality between optimal control and estimation. In Proceedings of the 47th IEEE Conference on Decision and Control, Cancun, Mexico, 9–11 December 2008; pp. 4286–4292.
20. Kappen, H.J.; Gómez, V.; Opper, M. Optimal control as a graphical model inference problem. *Mach. Learn.* **2012**, *87*, 159–182. [\[CrossRef\]](#)
21. Levine, S. Reinforcement learning and control as probabilistic inference: Tutorial and review. *arXiv* **2018**, arXiv:1805.00909.
22. Sun, X.; Bischl, B. Tutorial and survey on probabilistic graphical model and variational inference in deep reinforcement learning. In Proceedings of the IEEE Symposium Series on Computational Intelligence (SSCI), Xiamen, China, 6–9 December 2019; pp. 110–119.
23. Bishop, C.M. *Pattern Recognition and Machine Learning*; Springer: Berlin/Heidelberg, Germany, 2006.
24. Toussaint, M.; Harmeling, S.; Storkey, A. *Probabilistic Inference for Solving (PO) MDPs*; Technical Report; Technical Report EDI-INF-RR-0934; School of Informatics, University of Edinburgh: Edinburgh, UK, 2006.

25. Toussaint, M.; Charlin, L.; Poupart, P. Hierarchical POMDP Controller Optimization by Likelihood Maximization. *UAI* **2008**, *24*, 562–570.
26. Kumar, A.; Zilberstein, S.; Toussaint, M. Scalable multiagent planning using probabilistic inference. In Proceedings of the 22nd International Joint Conference on Artificial Intelligence, Barcelona, Spain, 16–22 July 2011.
27. Pajarinen, J.; Peltonen, J. Efficient planning for factored infinite-horizon DEC-POMDPs. In Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence, Barcelona, Spain, 16–22 July 2011; Volume 22, p. 325.
28. Pajarinen, J.; Peltonen, J. Periodic finite state controllers for efficient POMDP and DEC-POMDP planning. *Adv. Neural Inf. Process. Syst.* **2011**, *24*, 2636–2644.
29. Pajarinen, J.; Peltonen, J. Expectation maximization for average reward decentralized POMDPs. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases, Proceedings of the European Conference, ECML PKDD 2013, Prague, Czech Republic, 23–27 September 2013*; Springer: Berlin/Heidelberg, Germany, 2013; pp. 129–144.
30. Wu, F.; Zilberstein, S.; Jennings, N.R. Monte-Carlo expectation maximization for decentralized POMDPs. In Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence, Beijing, China, 3–9 August 2013.
31. Liu, M.; Amato, C.; Anesta, E.; Griffith, J.; How, J. Learning for decentralized control of multiagent systems in large, partially-observable stochastic environments. In Proceedings of the AAAI Conference on Artificial Intelligence, Phoenix, AZ, USA, 12–17 February 2016; Volume 30.
32. Song, Z.; Liao, X.; Carin, L. Solving DEC-POMDPs by Expectation Maximization of Value Function. In Proceedings of the AAAI Spring Symposium, Palo Alto, CA, USA, 21–23 March 2016.
33. Kumar, A.; Mostafa, H.; Zilberstein, S. Dual formulations for optimizing Dec-POMDP controllers. In Proceedings of the AAAI, Phoenix, AZ, USA, 12–17 February 2016.
34. Bertsekas, D.P. Approximate policy iteration: A survey and some new methods. *J. Control Theory Appl.* **2011**, *9*, 310–335. [[CrossRef](#)]
35. Liu, D.R.; Li, H.L.; Wang, D. Feature selection and feature learning for high-dimensional batch reinforcement learning: A survey. *Int. J. Autom. Comput.* **2015**, *12*, 229–242. [[CrossRef](#)]
36. Mnih, V.; Kavukcuoglu, K.; Silver, D.; Rusu, A.A.; Veness, J.; Bellemare, M.G.; Graves, A.; Riedmiller, M.; Fidjeland, A.K.; Ostrovski, G.; et al. Human-level control through deep reinforcement learning. *Nature* **2015**, *518*, 529–533. [[CrossRef](#)]
37. Hallak, A.; Mannor, S. Consistent on-line off-policy evaluation. In Proceedings of the International Conference on Machine Learning, PMLR, Sydney, Australia, 6–11 August 2017; pp. 1372–1383.
38. Gelada, C.; Bellemare, M.G. Off-policy deep reinforcement learning by bootstrapping the covariate shift. In Proceedings of the AAAI Conference on Artificial Intelligence, Honolulu, HI, USA, 27 January–1 February 2019; Volume 33, pp. 3647–3655.
39. Levine, S.; Kumar, A.; Tucker, G.; Fu, J. Offline reinforcement learning: Tutorial, review, and perspectives on open problems. *arXiv* **2020**, arXiv:2005.01643.
40. Hansen, E.A.; Bernstein, D.S.; Zilberstein, S. Dynamic programming for partially observable stochastic games. In Proceedings of the AAAI, Palo Alto, CA, USA, 22–24 March 2004; Volume 4, pp. 709–715.
41. Seuken, S.; Zilberstein, S. Improved memory-bounded dynamic programming for decentralized POMDPs. *arXiv* **2012**, arXiv:1206.5295.