

Delft University of Technology

Assessing the Optimality of Decentralized Inspection and Maintenance Policies for Stochastically Degrading Engineering Systems

Bhustali, Prateek; Andriotis, Charalampos P.

DOI 10.1007/978-3-031-74650-5_13

Publication date 2025 **Document Version** Final published version

Published in Artificial Intelligence and Machine Learning

Citation (APA)

Bhustali, P., & Andriotis, C. P. (2025). Assessing the Optimality of Decentralized Inspection and Maintenance Policies for Stochastically Degrading Engineering Systems. In F. A. Oliehoek, M. Kok, & S. Verwer (Eds.), Artificial Intelligence and Machine Learning: 35th Benelux Conference, BNAIC/Benelearn 2023, Delft, The Netherlands, November 8-10, 2023, Revised Selected Papers (pp. 236-254). (Communications in Computer and Information Science; Vol. 2187 CCIS). Springer. https://doi.org/10.1007/978-3-031-74650-5_13

Important note

To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.

Green Open Access added to TU Delft Institutional Repository

'You share, we take care!' - Taverne project

https://www.openaccess.nl/en/you-share-we-take-care

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.



Assessing the Optimality of Decentralized Inspection and Maintenance Policies for Stochastically Degrading Engineering Systems

Prateek Bhustali^(\boxtimes) and Charalampos P. Andriotis

Faculty of Architecture and the Built Environment, TU Delft, Delft, The Netherlands {P.Bhustali,C.Andriotis}@tudelft.nl

Abstract. Long-term inspection and maintenance (I&M) planning, a multi-stage stochastic optimization problem, can be efficiently formulated as a partially observable Markov decision process (POMDP). However, within this context, single-agent approaches do not scale well for large multi-component systems since the joint state, action and observation spaces grow exponentially with the number of components. To alleviate this curse of dimensionality, cooperative decentralized approaches, known as decentralized POMDPs, are often adopted and solved using multi-agent deep reinforcement learning (MADRL) algorithms. This paper examines the centralization vs. decentralization performance of MADRL formulations in I&M planning of multi-component systems. Towards this, we set up a comprehensive computational experimental program focused on k-out-of-n system configurations, a common and broadly applicable archetype of deteriorating engineering systems, to highlight the manifestations of MADRL strengths and pathologies when optimizing global returns under varying decentralization relaxations.

Keywords: Inspection and maintenance planning \cdot Decentralized partially observable Markov decision processes \cdot Multi-agent deep reinforcement learning \cdot Actor-critic methods \cdot Stochastic deterioration

1 Introduction

Inspection and maintenance (I&M) planning of deteriorating engineering systems, such as bridges, roads, aircraft, etc., is an optimization problem of seeking a policy that minimizes the expected life-cycle cost over a given time horizon. The problem is particularly challenging for engineering systems with multiple heterogeneous components that exhibit distinct system-level and component-level behaviours because the space of possible policies is intractably large. To make the problem tractable, risk-based and condition-based approaches use heuristics to confine the policy space over which the objective is minimized [12, 44]. Although these methods provide fast and explainable policies, it is well understood that they can be far from optimal [4, 32]. Alternatively, since I&M is a problem of decision-making under uncertainty, several works leverage the partially observable Markov decision process (POMDP) framework to model the environment and state uncertainties inherent to real-world systems [7,8,24,26,29,32]. This enables us to leverage single-agent deep reinforcement learning (SADRL) algorithms, which unlike heuristics, access the entire policy space. They also circumvent many complexity limitations of traditional Bellman backup operator-based POMDP solvers [37,38] when applied to I&M planning problems [6,33].

Although, in theory, neural networks can tackle arbitrarily high-dimensional spaces, SADRL approaches do not scale well under practical computational constraints because the joint space of states, observations and action spaces grow *exponentially* with the number of agents. A natural way to address this is to decentralize the problem by assigning an agent to each component or subsystem, thus articulating the problem as a cooperative multi-agent task. This relaxation is an extension of the single-agent POMDP to multi-agent systems and is formally known as a decentralized POMDP (Dec-POMDP) [3,30,31]. Likewise, we can move from SADRL to multi-agent deep reinforcement learning (MADRL) to solve the Dec-POMDP.

On the one hand, decentralization in reinforcement learning is key to addressing scalability. On the other hand, it can introduce certain side effects, such as environment non-stationarity, equilibrium selection, multi-agent credit assignment, and other issues, which encumber the learning task and may hinder convergence to optimal policies [11,14,27]. These can become potentially predominant in I&M settings due to strong global reward signals, risk costs, and other dependencies common in the mathematical description of deteriorating engineering systems. It, therefore, remains to be understood how the various multi-agent learning pathologies manifest in various I&M planning settings.

In this work, we focus on understanding the optimality characteristics of the three major MADRL paradigms: centralized training with centralized execution (CTCE), centralized training with decentralized execution (CTDE) and decentralized training with decentralized execution (DTDE) in the context of I&M planning. Specifically, we study this through the lens of a 5-component k-out-of-n system, a common and broadly applicable archetype of deteriorating engineering systems, modelled as a (Dec-)POMDP. This allows us to conduct extensive numerical experiments with several MADRL architectures. Finally, we obtain empirical insights on the performance of decentralized algorithms, which can ultimately impede the convergence of MADRL algorithms to global optima.

2 Related Work

Some recent works in the maintenance planning context have leveraged the Dec-POMDP framework to tackle the component multiplicity challenge. Andriotis & Papakonstantinou exploit the independence of agent actions and factorize the policy network output agent-wise, effectively converting the Dec-POMDP into a multi-agent POMDP (MPOMDP) [4,5] (similar to [18]). Although action factorization curbs the exponential growth of the policy network output to linear in the number of agents, the policy input space can still grow exponentially, unless dependencies are properly broken down into likewise factorized representations [4,5,28].

CTDE and DTDE are two practical paradigms for solving Dec-POMDPs, which try to address this challenge [2,17,31,42]. Various CTDE approaches such as COMA [13], QMIX [36], FACMAC [35], VDN [39] have been proposed that can address the credit assignment problem. Leroy et al. extensively compare the performance of various CTDE approaches in the maintenance planning context for k-out-of-n systems with varying n and demonstrate empirically the limitations of current CTDE algorithms in a large system with more than n = 50 components [23].

An extensive empirical comparison of the CTDE and DTDE paradigms in cooperative settings has been carried out on standard MARDL environments by Papoudakis et al. to benchmark the efficacy of MADRL algorithms in practice [34]. Lyu et al. provide a theoretical and empirical study on the effectiveness of centralized and decentralized critics and show that both critics have the same policy gradient in expectation. However, the variance in centralized critics is at least as much as their decentralized counterparts. They empirically demonstrate the performance of the CTDE and DTDE paradigms in several environments to highlight shortcomings of training under centralized critics [25]. Guillaume et al. were among the first to compare all MADRL paradigms to solve Dec-POMDPs [9]. They empirically compare CTCE, CTDE and DTDE paradigms on benchmark problems and use RNN-based policy networks, among others, to encode the action-observation history.

In this work, we also compare all MADRL paradigms in the context of I&M planning, emphasizing on k-out-of-n deteriorating systems. However, unlike [9], we encode action-observation histories using beliefs over component states. We do this for three reasons: transition and observation models are often available in this context; belief-based policies are more explainable than RNN-based agents that learn to encode action-observation histories, which is particularly important in risk-sensitive settings; and learning stability is improved as training directly on the belief space is more robust to environment noise. The goal is to understand the strengths and shortcomings of MADRL approaches, specifically when solving real-world I&M planning problems.

3 Background

3.1 Dec-POMDPs

We decentralize the I&M planning problem for a multi-component system by assigning an agent to each component and requiring the agents to cooperatively minimize the inspection and maintenance cost over the system's lifetime. The objective can be extended to include risk thresholds, minimum component health states, budget constraints, etc. [5]. This decentralized cooperative multiagent setting is formally called a Dec-POMDP [30,31], defined by the tuple $\langle \mathcal{M}, S, \mathbf{A}, T, C, \mathbf{O}, \Omega, t_H, \gamma \rangle$, where $\mathcal{M} := \{1, \ldots, M\}$ is the set of agents, S is the state space, $\mathbf{A} = \times A_m$ is the joint action space, $T : S \times \mathbf{A} \times S \to [0, 1]$ is the transition model, $C : S \times \mathbf{A} \to \mathbb{R}$ is the system cost model, $\mathbf{O} = \times O_m$ is the joint observation space, $\Omega : \mathbf{O} \times S \times \mathbf{A} \to [0, 1]$ is the observation model, t_H is the time horizon, and $\gamma \in [0, 1)$ is the discount factor.

The solution to the Dec-POMDP is an *optimal joint policy*, π^* , that minimizes the expected sum of discounted costs:

$$\boldsymbol{\pi}^* = \arg\min_{\boldsymbol{\pi}} \mathbb{E}_{\mathbf{a} \sim \boldsymbol{\pi}} \left[\sum_{t=0}^{t_H - 1} \gamma^t c_t \big| \boldsymbol{\pi} \right]$$
(1)

where $c_t \in C(s_t, \mathbf{a}_t)$ is the cost following the joint action $\mathbf{a}_t \in \mathbf{A}$ prescribed by the joint policy $\boldsymbol{\pi} = \langle \pi_1, \ldots, \pi_M \rangle$ in state s_t . The policy $\boldsymbol{\pi}$ can be stochastic, mapping the agent's action-observation history to a probability distribution over actions, or deterministic, mapping the agent's observation history to agents' actions. It is known that every Dec-POMDP has at least one optimal deterministic joint policy [31].

In several I&M planning scenarios, the system components deteriorate independently, enabling a factorization of the state space, $S = \times S_m$, transition model, $T_m : S_m \times A_m \times S_m \rightarrow [0, 1]$ and observation model, $\Omega_m : S_m \times A_m \times O_m \rightarrow$ [0, 1] making it a Dec-POMDP with transition and observation independence [15, 21]. In scenarios where components are correlated or dependent, they can be transformed into independent representations through proper reconstruction of the underlying dynamic Bayesian network [28].

The deterioration models are often available or learnable offline, which enables us to maintain beliefs over the states of each agent [4,19,45]. When an inspection action is taken, we use the observation o_m to update the belief over a component's state s'_m using the Bayes' rule:

$$b_{m,t+1}(s'_m) = \frac{\Omega_m(o_m|s'_m, a_m) \cdot \sum_{s_m \in S_m} T_m(s'_m|s_m, a_m) b_{m,t}(s_m)}{\sum_{s'_m \in S_m} \Omega_m(o_m|s'_m, a_m) \sum_{s_m \in S_m} T_m(s'_m|s_m, a_m) b_{m,t}(s_m)}$$
(2)

An agent's policy, typically tied to a specific component, maps component beliefs to actions $\pi_m : B_m \to A_m$.

3.2 Baseline Performance

We establish baseline performance using the classical time-periodic inspections with condition-based maintenance (TPI-CBM) approach as in [5,16,24]. In this strategy, we inspect components at fixed time intervals (Δt_{insp}) and, at the inspection step, take maintenance actions based on observations (\boldsymbol{o}_{insp}), denoted by a sub-policy $\hat{\pi}(\boldsymbol{o}_{insp})$.

$$\boldsymbol{\pi}_{TC}^{*}(\Delta t_{\text{insp}}^{*}, \hat{\pi}^{*}(\boldsymbol{o}_{\text{insp}})) = \underset{\Delta t_{\text{insp}}, \hat{\pi}(\boldsymbol{o}_{\text{insp}})}{\arg\min} \mathbb{E}_{\mathbf{a} \sim \boldsymbol{\pi}_{TC}} \left[\sum_{t=0}^{t_{H}-1} \gamma^{t} c_{t} | \boldsymbol{\pi}_{TC} \right]$$
(3)

Intuitively, we formulate the inspection and maintenance planning problem as a combinatorial optimization problem by evaluating the objective defined in

Paradigm	Mathematical Framework	Algorithm	Observation	Action	Critic	Actor
CTCE	POMDP	$\mathbf{JAC}\ (\mathrm{SADRL})$		Joint		Shared
	MPOMDP	DCMAC DDMAC	Joint	Factored Factored	Centralized	Shared Separate
CTDE	Dec-POMDP	IACC IACC-PS	Independent	Independent Independent	Centralized	Separate Shared
DTDE		IAC IAC-PS	Independent	Independent Independent	Decentralized	Separate Shared

Table 1. Single- and multi-agent DRL architectures studied in this work.

Eq.(3) for each policy $\pi_{TC}(\Delta t_{\text{insp}}, \hat{\pi}(\boldsymbol{o}_{\text{insp}}))$ using Monte Carlo rollouts. Additionally, failed components are immediately repaired, enabling the strategy to retrogress to simple corrective maintenance when necessary.

3.3 Multi-agent Deep Reinforcement Learning (MADRL)

In this work, we study the 7 variants of MADRL as listed in Table 1. In all of them, we employ an off-policy actor-critic approach as in ACER [41] for sample efficiency. Following standard nomenclature [17], we categorize these algorithms based on information available during training and execution, as explained below and summarized in Table 1.

Centralized Training with Centralized Execution (CTCE). This paradigm essentially converts the Dec-POMDP either into a single-agent POMDP to directly learn the mapping between the joint belief space and joint action space or a multiagent POMDP (MPOMDP) to learn the mapping between the joint belief space and the agent-wise factored action space. Although the single-agent POMDP approaches are not scalable, they theoretically encompass global optima in their solution space. MPOMDP solutions relax the requirement for joint action spaces and can be seen as a semi-CTCE paradigm that can approximate well the POMDP solution space under mild conditions. Specifically, we consider three actor-critic algorithmic approaches:

- Joint actor-critic (JAC), where the actor learns the joint stochastic policy $\pi(\mathbf{a}|\mathbf{b};\theta)$, mapping the joint beliefs to joint actions, and the critic learns the value function $V^{\pi}(\mathbf{b};\phi)$ under that policy.
- **Deep centralized multi-agent actor-critic (DCMAC)**, where we exploit the factorizable nature of the joint action space and learn a policy for each agent that maps the joint belief to agents' actions, $\boldsymbol{\pi} = \langle \pi(a_m | \mathbf{b}; \theta) \rangle_{m=1}^M$, and a centralized critic $V^{\pi}(\mathbf{b}; \phi)$ guides each agent as in [4, 18].
- Deep decentralized multi-agent actor-critic (DDMAC): Like DCMAC, we exploit the factorizable action space but employ independent networks for each agent mapping the joint belief to component action, $\boldsymbol{\pi} = \langle \pi(a_m | \mathbf{b}; \theta_m) \rangle_{m=1}^M$, and a centralized critic $V^{\boldsymbol{\pi}}(\mathbf{b}; \phi)$ guides each agent as

introduced in [5]. This approach aims to model policy distributions independently, thus alleviating potential training complexities stemming from parameter sharing in the centralized paradigm.

Centralized Training with Decentralized Execution (CTDE). This paradigm assumes that the agents may have access to centralized information signals, such as joint observations, joint actions and critic gradients, only during training but act independently based on local observations during execution/inference. We consider two information accessibility scenarios:

- Independent actor centralized critic (IACC), where each agent only has access to its own component's beliefs and learns the local stochastic policy $\pi_m(a_m|b_m;\theta_m)$ guided by a centralized critic $V^{\pi}(\mathbf{b};\phi)$.
- IACC with parameter sharing (IACC-PS): A special case of IACC where agents share the same policy network $\pi(a_m|b_m, m; \theta)$ but the beliefs are indexed (using one-hot encoding) to enable the network to learn distinct policies.

Decentralized Training with Decentralized Execution (DTDE). Agents only have access to local information during training and execution and never have access to centralized information as in the previous cases. We consider two cases:

- Independent actor-critic (IAC), where each agent learns a stochastic policy mapping its belief to actions $\pi_m(a_m|b_m;\theta_m)$ and is guided by a decentralized critic $V^{\pi_m}(b_m;\phi_m)$.
- IAC with parameter sharing (IAC-PS), where agents share both policy networks $\pi(a_m|b_m, m; \theta)$ and critic networks $V^{\pi}(b_m, m; \phi)$ but the beliefs are indexed (using one-hot encoding) to enable distinct outputs.

Generally, the parameter-sharing approach can be extended to agents with heterogeneous action space cardinalities [40]. Several works have empirically demonstrated the benefits of parameter sharing in environments with homogeneous agents (in terms of observations and actions) [13, 18, 43]. However, Christianos et al. demonstrate empirically that these benefits are environment-specific, and such parameter sharing can become detrimental in certain environments [10]. Therefore, we study both variants in this work.

Decentralization presents several challenges due to the presence of other agents that impede convergence to optimal policies. From the perspective of a single agent, the presence of other agents affects [2, 14, 27]:

- Transitions: The environment is perceived as non-stationary by one agent due to the actions of other agents;
- Rewards: Reward received due to joint actions must be disentangled to identify an individual agent's contribution (known as multi-agent credit assignment problem), and exploratory actions by agents obfuscate the reward signal (also known as alter-exploration issue);

- Agent policy: Action selection must be coordinated when multiple optimal joint policies exist (referred to as equilibria selection). Agents can converge to a sub-optimal equilibrium because miscoordination due to unilateral deviation from an optimal equilibrium has lower gains/higher penalties than unilateral deviation from the sub-optimal equilibria. In such cases, a sub-optimal equilibrium is said to *shadow* the optimal equilibrium.

4 Experimental Setup

4.1 Environment: k-out-of-n System

This work studies the optimality characteristics of MADRL approaches, focusing on a k-out-of-n system with heterogeneous components. The k-out-of-n:G (G: good) is a common archetype of deteriorating systems in which the system is functional when at least k out of its n components are working and, therefore, $n \ge k$. The system has two notable special cases, namely k = n, a series system, and k = 1, representing a parallel component configuration. Examples of such systems include road networks, transmitters in communication networks, human kidneys, etc. [20, 22].

We consider a system with n = 5 components and assign an agent to each component, thus M = n. The state space of each component (S_m) describes its range of health states, $S_m := \{s^1 = \text{no-damage}, s^2 = \text{minor-damage}, s^3 =$ major-damage, $s^4 = \text{failure}\}$ and similarly the component-wise actions are given by $A_m := \{a^1 = \text{do-nothing}, a^2 = \text{repair}, a^3 = \text{inspect}\}$. Each component m has a unique and stationary deterioration model T_m^d , the natural deterioration of the environment (e.g. ageing due to corrosion, fatigue or other stressors):

$$\begin{bmatrix} T_1^d & T_2^d & T_3^d \\ 0.82 \ 0.13 \ 0.05 \ 0 \\ 0 & 0.87 \ 0.09 \ 0.04 \\ 0 & 0 & 0.91 \ 0.09 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 0.72 \ 0.19 \ 0.09 \ 0 \\ 0 & 0.78 \ 0.18 \ 0.04 \\ 0 & 0 & 0.85 \ 0.15 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 0.79 \ 0.17 \ 0.04 \ 0 \\ 0 & 0.85 \ 0.09 \ 0.06 \\ 0 & 0 & 0.91 \ 0.09 \\ 0 & 0 & 0 & 1 \end{bmatrix} \\ \begin{bmatrix} T_4^d & T_5^d \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 0.80 \ 0.12 \ 0.08 \ 0 \\ 0 & 0.83 \ 0.12 \ 0.05 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 0.88 \ 0.12 \ 0 & 0 \\ 0 & 0.93 \ 0.07 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

The above synthetic transition models capture the main characteristics of a deterioration model, namely, they are upper-triangular (component state cannot be improved without repair) and unit-mass at the terminal state (failure is an absorbing state).

The *repair* action restores the component back to state s^1 , but this action only succeeds with a probability r_m and the transition model corresponding to this action is given by:

$$T_m^r(s_m'|s_m, repair) := \begin{bmatrix} 1 & 0 & 0 & 0 \\ r_m & 1 - r_m & 0 & 0 \\ r_m & 0 & 1 - r_m & 0 \\ r_m & 0 & 0 & 1 - r_m \end{bmatrix}$$
(4)

where r_m is the repair accuracy of agent m, reflecting the uncertainty in the duration the action requires to be completed, or other. We use the following value of r_m for each of the m components 1, 0.9, 0.95, 0.85, 0.8, respectively. The chance of unsuccessful repair actions aims to capture the inevitable human error during maintenance activities in practice.

We summarize the transition model for a component as follows:

$$T_m := \begin{cases} T_m^d, & \text{if } a = \text{ do-nothing or } a = \text{ inspect} \\ T_m^r \times T_m^d, & \text{if } a = \text{ repair} \end{cases}$$
(5)

The cost model is divided into component-level costs (repair and inspection costs) and system-level costs. For each component m, the repair costs c_m^{repair} are 30, 90, 80, 250, 350, respectively, and inspections costs $c_m^{inspect}$ are 20, 40, 25, 50, 100, respectively. System failure leads to a penalty $c^{failure}$ equal to 3 times the sum of component repair costs. The cost model can be summarized as follows:

$$C(s,a) = \mathbb{1}_{failure}(s) \cdot c^{failure} + \sum_{m=1}^{M} \mathbb{1}_{a^2}(a_m) \cdot c_m^{repair} + \mathbb{1}_{a^3}(a_m) \cdot c_m^{inspect}$$
(6)

where $\mathbb{1}_{y}(x)$ denotes the indicator function, taking a value of 1 if x=y, and 0 otherwise. This penalty necessitates the agents to cooperatively coordinate maintenance actions to ensure global system functionality.

Each component has a unique observation model Ω_m through which the respective agent obtains imperfect observations of its true state when the inspection action is chosen and is given as follows:

$$\Omega_m := \begin{bmatrix} p_m & 1 - p_m & 0 & 0\\ \frac{(1-p_m)}{2} & p_m & \frac{(1-p_m)}{2} & 0\\ 0 & 1 - p_m & p_m & 0\\ 0 & 0 & 0 & 1 \end{bmatrix}$$
(7)

where p_m describes the observation accuracy of component m for the inspection action and takes values of 0.8, 0.85, 0.9, 0.95, and 0.8 for the five agents. Component failure, however, is assumed to be self-announcing, i.e. information about failed components is noise and cost-free. The objective is to minimize the operation cycle cost of the system over $t_H = 50$ with $\gamma = 0.99$.

4.2 MADRL Algorithms and Baselines

For brevity, we only describe the algorithm for JAC in detail in Appendix 1. The same off-policy features are used for all seven multi-agent approaches. All MADRL algorithms hyperparameters are tuned on the 4-out-of-5 setting, and the Table 2. Comparison of the best performance observed for each algorithm (with \pm indicating the 95% confidence interval when evaluating the optimization objective (1) with 10,000 MC rollouts. Bold indicates best in the respective k-out-of-n setting.

		1-out-of-5	2-out-of-5	3-out-of-5	4-out-of-5	5-out-of-5
Baseline	TPI-CBM	$1485.74 (\pm 8.43)$	$1498.8 (\pm 8.62)$	$1654.19 (\pm 14.11)$	$3998.41 (\pm 43.51)$	$19728.42 (\pm 95.45)$
CTCE	JAC	$365.51 \ (\pm 10.90)$	735.96 (±13.59)	$1686.41 (\pm 18.56)$	3413.96 (±33.88)	12595.17 (± 67.70)
	DCMAC	$363.64~(\pm 10.96)$	735.56 (± 14.08)	$1789.00 \ (\pm 21.25)$	$3396.92 \ (\pm 32.64)$	12509.88 (± 66.73)
	DDMAC	$364.10~(\pm 10.67)$	732.62 (± 13.49)	$1858.60 \ (\pm 24.16)$	$3384.11\ (\pm 32.95)$	$12764.02 \ (\pm 67.13)$
CTDE	IACC	298.01 (± 7.96)	$820.21 \ (\pm 13.80)$	1600.01 (± 15.00)	$3460.20 \ (\pm 35.57)$	$12685.70 \ (\pm 66.53)$
	IACC-PS	$\textbf{285.64}~(\pm 7.99)$	833.55 (±10.28)	1578.24 (± 14.88)	$3421.71 \ (\pm 35.39)$	12543.43 (± 69.51)
DTDE	IAC	$563.99 (\pm 21.91)$	$1422.94 (\pm 18.13)$	$1603.05 \ (\pm 14.76)$	$3421.37 \ (\pm 32.43)$	$13291.65 (\pm 73.41)$
	IAC-PS	$296.57 (\pm 8.42)$	850.90 (±16.11)	$1646.18 (\pm 14.46)$	3395.78 (±34.13)	$13370.26 (\pm 76.00)$

best-performing ones are reported in Table 4 (see appendix). We use these tuned hyperparameters and train ten instances of each agent with random seeds on all k-out-of-n settings. To pick the best policy, we evaluate the agent periodically every 4,000 training episodes using 10,000 Monte Carlo rollouts. A large number of rollouts is needed due to the high variance of expected discounted cost Eq. (1), as also highlighted in [23].

For the TPI-CBM heuristic, the policy space has $t_H \times (|S_m|-2) = 50 \times 3 = 150$ policies, assuming policy uniformity over components. We subtract 2 since the initial and final states' actions are fixed, i.e., do-nothing and repair, respectively. However, if we were to allow different component-wise policies, the search would grow to $150^5 = 7.5 \times 10^{10}$, rendering the combinatorial optimization of Eq. (3) intractable since evaluating the objective requires a large number of Monte Carlo evaluations. Therefore, we optimize the heuristic, assuming the same policy for all components and find an optimum over the smaller search space of 150 policies.

All code and experimental data are publicly available: https://github.com/ omniscientoctopus/optimality-of-decentralization

5 Results and Discussion

We summarize the performance of the MADRL algorithms in all k-out-of-n settings using box plots and the test performance of agents during training in Fig. 1. Additionally, we report the best performance observed for each algorithm in Table 2. For a more detailed comparison, we report mean performance with 95% confidence intervals over ten random seeds in Table 3 in the appendix. To enable a concise performance comparison of the paradigms, we aggregate the results across settings using mean, interquartile mean (IQM) and median in Fig. 2 as demonstrated by Agarwal et al. in [1]. Before aggregating the performance, we normalize the results with the respective baselines in each setting.

CTCE: Figure 1 shows that JAC, DCMAC, and DDMAC are all able to consistently outperform the heuristic in most settings, and the equivalence of their individual aggregate performance is shown in Fig. 2, where the confidence intervals have a significant overlap. However, in the 3-out-of-5 setting, while their best instance remains competitive, we observe discrepancies in their performance, as reported in Table 2. Sub-optimalities in the CTCE paradigm have been previously reported by Bono et al. and have been attributed to the difficult exploration



Fig. 1. Performance of CTCE, CTDE and DTDE algorithms in different k-out-of-n settings, normalized w.r.t. the heursitic (h) (lower is better). Left: Performance of the algorithms evaluated every 4,000 episodes during training, using 10,000 rollouts. The curve for each algorithm is aggregated over ten training instances. The bold line and the shaded region indicate the median and interquartile range respectively. **Right**: Each box plot summarizes the performance of the best policies across ten training instances, where the box represents the interquartile range and circles denote the outliers.



Comparison of MADRL algorithms (lower is better)

Fig. 2. Aggregate performance over the different k-out-of-n settings: mean, median and interquartile mean (IQM) (lower is better). Results for individual settings are normalized with respective heuristic baselines. Vertical lines indicate point estimates, and bands show 95% confidence intervals.

problem imposed by the joint state and action spaces [9]. We concur with this observation, especially because, unlike [9], we encode action-observation histories using beliefs and thus eliminating any additional complexities arising due to alternate approximations. The observed discrepancy cannot be merely an artifact of low neural network representational capacity as long as the same network architecture performs well in other settings. Another observation underpinning this claim is the best-performing training instance (random seed) of the IACC, IACC-PS and IAC algorithms, reported in Table 2, that are able to surpass the performance of CTCE algorithms and prove the existence of better policies.

CTDE: Both IACC and IACC-PS exhibit performance equivalent to CTCE algorithms in the 5-out-of-5 setting, but begin to exhibit variability as we move towards the 1-out-of-5 setting as shown in Fig. 1. However, surprisingly, the best training instances of IACC and IACC-PS algorithms outperform JAC, the best CTCE training instance, in the 3-out-of-5 setting by about 7%. We examine this closely by comparing the histogram of returns from 10,00 rollouts for the IACC-PS and JAC policies in Fig. 5 and contrast it with the corrective repair heuristic policy. We clearly observe how the environment under JAC policy is slightly more likely to witness at least one system failure. Additionally, we also plot individual rollouts under these polices in Figs. 3 and 4, respectively. In most time steps, the IACC-PS policy chooses corrective repair actions, however, it is able to marginally outperform the corrective repair heuristic by avoiding costly repair actions closer to the end of the life-cycle, such as for component 4 after t = 40. In contrast, the JAC agent chooses to monitor component 3 via inspections, and chooses a combination of predictive and corrective repair actions to keep



Fig. 3. Sample rollout of the best JAC policy in the 3-out-of-5 setting. The JAC agent chooses to monitor component 3 via inspections, and uses a combination of predictive and corrective repair actions to keep the system functional.



Fig. 4. Sample rollout of the best IACC-PS policy in the 3-out-of-5 setting. In most time steps, the agents choose simple corrective repair actions, however, they are able to marginally outperform the corrective replace heuristic policy by avoiding costly repair actions closer to the end of the life-cycle. We observe this for component 4 after t = 40, where the agent chooses do-nothing despite component failure.

the system functional. A possible explanation is that this component combines aggressive deterioration and low inspection to repair cost ratio.

DTDE: In terms of aggregate performance, fully decentralized algorithms are outperformed by both CTCE and CTDE algorithms, as shown in Fig. 2. Similar to CTDE algorithms, DTDE algorithms exhibit increasing variability in performance as we move from the 5-out-of-5 setting to the 1-out-of-5 setting, i.e. towards environments requiring more coordination among agents, as shown in Fig. 1. However, the performance of DTDE algorithms is more adversely



Fig. 5. Histograms of the discounted episodic costs of the best performing training instances (as highlighted in Table 2) in the 3-out-of-5 setting compared against a corrective repair heuristic policy over 10,000 environment rollouts. The modes in the histograms correspond to episodes in which system failures occur. Note that the IACC-PS policy closely resembles the corrective repair policy, but is marginally better as explained in Fig. 4.



Fig. 6. Sample rollout of an outlier IAC policy in the 1-out-of-5 setting (indicated with † in Fig. 1), where the agents succumb to multi-agent learning pathologies, resulting in a poor policy. Agents 1,3,4,5 choose do-nothing actions even during system failure, and only agent 2 takes necessary repair actions. (Color figure online)

impacted as the need for cooperation in the environment increases. We illustrate this using a sample rollout of an outlier policy in the 1-out-of-5 setting (indicated with † in Fig. 1) in Fig. 6. Although the best performing instance is able to outperform the heuristic, this training instance succumbs to multi-agent learning pathologies and results in a poor policy. The performance of the parameter-sharing variants show similar variability, but can sometimes stumble into good policies, as reported in Table 2 for the 1-out-of-5, 4-out-of-5 and 5-out-of-5 settings.

Overall, as we move from the 5-out-of-5 setting to the 1-out-of-5 setting, environments demand more coordination among agents. In Fig. 2, we observe a clear trend: CTCE algorithms are unaffected by coordination requirements and thus show little variability in aggregate performance. In contrast, DTDE algorithms show large variability highlighting their susceptibility to multi-agent learning pathologies. CTDE algorithms are able to mitigate these and foster coordination among agents, thus performing better than their DTDE counterparts. Lastly, in Fig. 2, we also observe that the performance of the parameter sharing variants is comparable to their independent counterparts, in addition to being significantly more computationally efficient to train.

6 Conclusions

Multi-agent deep reinforcement learning provides a scalable approach to optimal inspection and maintenance (I&M) planning for multi-component systems. In this work, we study the three major multi-agent learning paradigms, i.e., centralized training with centralized execution (CTCE), centralized training with decentralized execution (CTDE) and decentralized training with decentralized execution (DTDE), through the lens of a k-out-of-5 system to investigate their efficacy and limitations in I&M planning. Our key conclusions are listed below:

- While CTCE methods are shown to be superior overall, they also remain susceptible to sub-optimal policies due to exploration challenges induced by the joint state/action spaces. More importantly, the large joint spaces hinder scalability of these algorithms in real-world I&M problems.
- CTDE and DTDE approaches, although more scalable, face challenges stemming from decentralization, such as the emergence of non-stationary environments, multi-agent credit assignment issues, shadowed equilibria, etc. Vanilla CTDE approaches are observed to mitigate these pathologies and consistently outperform their decentralized counterparts. However, due to the large variability in performance, DTDE algorithms are sometimes still able to converge to competitive policies.
- The overall performance of parameter-sharing variants in both CTDE and DTDE paradigms is comparable to their independent counterparts, despite heterogeneity in system components in terms of costs, deterioration characteristics, and action effects. Additionally, parameter-sharing variants are significantly more computationally efficient to train.

 Finally, with increasing number of redundant components in the system, more cooperation among agents is needed to maximally exploit these redundancies. As such, environments with redundant system components can serve as a good benchmark to test agent cooperation in the context of I&M planning.

We aim to continue this line of work and study more rigorously the influence of component connectivity configurations in deteriorating engineering systems, such as series-parallel systems, and homogeneity/heterogeneity of system components on policy learning.

Acknowledgements. This material is based upon work supported by the TU Delft AI Labs program. The authors gratefully acknowledge this support.

A Appendix

Table 3. Mean performance of the algorithms aggregated over ten training instances (random seeds) with \pm indicating the 95% confidence interval. Bold indicates best in the respective k-out-of-n setting.

	1-out-of-5		2-out-of-5	3-out-of-5	4-out-of-5	5-out-of-5	
Baseline	TPI-CBM	1485.74	1498.8	1654.19	3998.41	14780.43	
CTCE	JAC	372.46 (±2.51)	$1070.24 \ (\pm 178.81)$	2487.57 (±285.79)	$3488.80 (\pm 11.30)$	$12698.22 (\pm 50.20)$	
	DCMAC	377.01 (±5.83)	$1146.16\ (\pm 162.51)$	$2774.12 (\pm 256.45)$	$3451.70 (\pm 9.68)$	$12642.19 (\pm 56.38)$	
	DDMAC	$379.71 (\pm 21.31)$	$1194.30~(\pm 188.91)$	$2965.18 (\pm 152.77)$	$3517.82 \ (\pm 12.38)$	13097.86 (±160.00)	
CTDE	IACC	$13467.32 (\pm 940.42)$	$13844.57 (\pm 988.96)$	$14329.28 (\pm 850.25)$	9952.01 (±723.77)	$20830.69 (\pm 277.56)$	
	IACC-PS	13311.32 (± 1326.17)	13320.42 (± 1316.56)	11802.21 (± 2468.70)	$7449.54 (\pm 1025.22)$	17457.71 (± 2312.53)	
DTDE	IAC	4797.06 (±3042.81)	4892.51 (±2266.34)	$3917.98 (\pm 1273.71)$	4319.42 (±49.72)	$13528.65 (\pm 278.93)$	
	IAC-PS	$5442.69 (\pm 3246.56)$	7326.18 (±2951.19)	$4051.89\ (\pm 2519.96)$	$3632.35 (\pm 86.58)$	$13505.44 \ (\pm 131.70)$	

 Table 4. Tuned hyperparameters of the algorithms used to train agents on various k-out-of-n settings

Paradigm			CTCE	CTCE CTE			DTDE			
Algorithm		JAC	DCMAC	DDMAC	IACC	IACC-PS	IAC	IAC-PS		
Episodes (E)					100,000					
Timesteps			$E \times t_H = 100,000 \times 50 = 5M$							
Architecture	Actor	[21, 64, 64, 243]	[21, 32, 32, 15]	[21, 16, 16, 3] x 5	[5, 16, 16, 3] x 5	[10, 32, 32, 3]	[5, 16, 16, 3] x 5	[10, 32, 32, 3]		
	# parameters	21,363	2,255	3,375	2,095	1,507	2,095	1,507		
	Critic	[21, 64, 64, 1]	[21, 64, 64, 1]	[21, 64, 64, 1]	[21, 64, 64, 1]	[21, 64, 64, 1]	$[5, 16, 16, 1] \ge 5$	[10, 64, 64, 1]		
	# parameters	5,633	5,633	5,633	5,633	5,633	1,925	4,929		
	Total # parameters	26,996	7,888	9,008	7,728	7,140	4,020	6,436		
learning batch size				64 (timesteps)						
	optimizer				Adam					
	initial lr	0.0001	0.0001	0.0001	0.0005	0.0005	0.0001	0.0005		
Actor	decay factor	0.1								
Actor	decay episodes	20,000	20,000	20,000	10,000	10,000	25,000	25,000		
	decay type	linear								
Critic	initial_lr	0.005	0.005	0.001	0.005	0.001	0.001	0.001		
	decay factor	0.1								
	decay episodes	20,000	20,000	20,000	10,000	25,000	25,000	25,000		
	decay type	linear								
ϵ -greedy strategy	ϵ -start				1					
	ε-end	0.005	0.005	0.005	0.001	0.005	0.001	0.001		
	decay episodes	20,000	20,000	20,000	10,000	10,000	25,000	25,000		
	decay type		linear							
Replay Buffer	timesteps	500.000	500,000	500.000	10.000	10.000	10.000	10.000		

As described in the main text, we use the 4-out-of-5 setting for hyperparameter optimization and report the tuned hyperparameters in Table 4. We note that decentralized agents often exhibit instabilities when using large replay buffers, thus, their replay buffers are much smaller than their centralized counterparts. This is because random samples from a large replay buffer can correspond to policies significantly different from the current policy, forcing agents to modify their recently learned policy drastically.

Algorithm 1: Joint Actor Critic (JAC) **Initialize:** Actor $\pi(\mathbf{a}|\mathbf{b};\theta)$, Critic $V^{\pi}(\mathbf{b};\phi)$, Experience Replay buffer \mathcal{D} , exploration scheduler ε , learning rates α_a^0, α_c^0 , learning rate schedulers ζ_a, ζ_c , optimizers η_a, η_c , batch size B, Episodes E, global time step t' = 0, exploration probability ϵ_0 . 1 for e = 1 to E do $\mathbf{b}_0 = \mathrm{env}()$ // get initial belief 2 for t = 0 to $t_H - 1$ do 3 if explorationstrategy(ϵ'_t) = random then 4 $\mathbf{a}_t \sim \mathcal{U}(\mathbf{A}), \mathbf{A}$ is the joint action space $\mathbf{5}$ 6 else Sample action: $(\mathbf{a}_t, \mathbf{p}_t) = \boldsymbol{\pi}(\mathbf{b}_t; \theta)$ 7 Environment step: $(\mathbf{b}_{t+1}, c_t, \mathbf{d}_t) = \operatorname{env}(\mathbf{a}_t)$ // $d_t := \mathbb{1}_{t_H-1}(t)$ 8 Store experience: $(\mathbf{b}_t, \mathbf{a}_t, p_t, \mathbf{b}_{t+1}, r_t, \mathbf{d}_t) \cup \mathcal{D}$ 9 if $t' > 10 \times B$ then 10 Sample batch from buffer: $\mathcal{B} = \{(\mathbf{b}_t, \mathbf{a}_t, p_t, \mathbf{b}_{t+1}, c_t, \mathbf{d}_t)\}_{i=1}^{\mathsf{B}} \sim \mathcal{D}$ 11 $\pi(\mathbf{a}|\mathbf{b}; \theta^{t+1}), V^{\pi}(\mathbf{b}; \phi^{t+1}) = \texttt{Train()}$ 12 $\mathbf{b}_t \longleftarrow \mathbf{b}_{t+1}$ 13 $\epsilon_{t'+1} \longleftarrow \varepsilon(\epsilon'_t)$ // update exploration parameter 14 $\alpha_a^{t'+1} \longleftarrow \zeta_a(\alpha_a^{t'})$ // update actor learning rate 15 $\begin{array}{c} \alpha_c^{t'+1} \longleftarrow \zeta_c(\alpha_c^{t'}) \\ t' \longleftarrow t'+1 \end{array}$ // update critic learning rate 16 17 18 Train(): For each sample i in batch, 19 $\texttt{future} = \texttt{sg}(V^{\pi}(\mathbf{b}'_i; \phi)) \cdot \mathtt{d}_i$ 20 // stop gradients: $sg(\cdot)$ Advantage: $\mathbf{A}_{i} = c_{i} + \gamma \cdot \mathtt{future} - V^{\pi}(\mathbf{b}_{i}; \phi)$ // TD error $\mathbf{21}$ $\rho_i = \log \pi(\mathbf{a}_i | \mathbf{b}_i; \theta)$ 22 Importance sampling weights: $w_i = \frac{\operatorname{sg}(\pi(\mathbf{a}_i | \mathbf{b}_i; \theta))}{2}$ $\mathbf{23}$ Clipping weights: $w_i \leftarrow \min(w_i, \bar{w})$ 24 Actor gradients: $g_{\theta} = \frac{1}{B} \sum_{i=1}^{B} w_i \cdot (\nabla_{\theta} \rho_i \cdot \mathbf{A}_i)$ Critic gradients: $g_{\phi} = \frac{1}{B} \sum_{i=1}^{B} w_i \cdot \nabla_{\phi} \mathbf{A}_i$ $\mathbf{25}$ $\mathbf{26}$ $\theta' \longleftarrow \eta_a(\mathbf{g}_{\theta}, \theta, \alpha_a)$ 27 $\mathbf{28}$ $\phi' \longleftarrow \eta_c(\mathbf{g}_{\phi}, \phi, \alpha_c)$ return $\boldsymbol{\pi}(\mathbf{a}|\mathbf{b}; \boldsymbol{\theta}'), \, V^{\boldsymbol{\pi}}(\mathbf{b}; \boldsymbol{\phi}')$ 29

For training, we employ a variant of the actor-critic with experience replay (ACER) algorithm for learning [41] as introduced for I&M in [4,5] and outlined above. To minimize the variance caused by the importance sampling weights, we clip the values w_i by setting $\bar{w} = 2$.

References

- Agarwal, R., Schwarzer, M., Castro, P.S., Courville, A.C., Bellemare, M.: Deep reinforcement learning at the edge of the statistical precipice. In: Advances in Neural Information Processing Systems, vol. 34, pp. 29304–29320. Curran Associates, Inc. (2021). https://proceedings.neurips.cc/paper/2021/hash/ f514cec81cb148559cf475e7426eed5e-Abstract.html
- Albrecht, S.V., Christianos, F., Schäfer, L.: Multi-Agent Reinforcement Learning: Foundations and Modern Approaches. MIT Press, Cambridge (2023)
- Amato, C., Chowdhary, G., Geramifard, A., Ure, N.K., Kochenderfer, M.J.: Decentralized control of partially observable Markov decision processes. In: 52nd IEEE Conference on Decision and Control. pp. 2398–2405. IEEE, Firenze (2013). http://ieeexplore.ieee.org/document/6760239/
- Andriotis, C.P., Papakonstantinou, K.G.: Managing engineering systems with large state and action spaces through deep reinforcement learning. Reliab. Eng. Syst. Saf. 191(March), 106483 (2019). https://doi.org/10.1016/j.ress.2019.04.036
- Andriotis, C.P., Papakonstantinou, K.G.: Deep reinforcement learning driven inspection and maintenance planning under incomplete information and constraints. Reliab. Eng. Syst. Saf. 212(March), 107551 (2021). https://doi.org/10. 1016/j.ress.2021.107551
- Andriotis, C.P., Papakonstantinou, K.G., Chatzi, E.N.: Value of structural health information in partially observable stochastic environments. Struct. Saf. 93, 102072 (2021). https://doi.org/10.1016/J.STRUSAFE.2020.102072
- Arcieri, G., Hoelzl, C., Schwery, O., Straub, D., Papakonstantinou, K.G., Chatzi, E.: Bridging POMDPs and Bayesian decision making for robust maintenance planning under model uncertainty: an application to railway systems. Reliab. Eng. Syst. Saf. 239, 109496 (2023). https://www.sciencedirect.com/science/article/pii/ S0951832023004106
- Bismut, E., Straub, D.: Optimal adaptive inspection and maintenance planning for deteriorating structural systems. Reliab. Eng. Syst. Saf. 215, 107891 (2021). https://www.sciencedirect.com/science/article/pii/S0951832021004063
- Bono, G., Dibangoye, J.S., Matignon, L., Pereyron, F., Simonin, O.: Cooperative multi-agent policy gradient. In: Berlingerio, M., Bonchi, F., Gärtner, T., Hurley, N., Ifrim, G. (eds.) ECML PKDD 2018. LNCS (LNAI), vol. 11051, pp. 459–476. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-10925-7_28
- Christianos, F., Papoudakis, G., Rahman, A., Albrecht, S.V.: Scaling multi-agent reinforcement learning with selective parameter sharing (2021). http://arxiv.org/ abs/2102.07475, arXiv:2102.07475 [cs]
- Claus, C., Boutilier, C.: The dynamics of reinforcement learning in cooperative multiagent systems. In: Proceedings of the Fifteenth National/Tenth Conference on Artificial Intelligence/innovative Applications of Artificial Intelligence, pp. 746– 752. AAAI 1998/IAAI 1998, American Association for Artificial Intelligence, USA (1998)

- Deodatis, G., Fujimoto, Y., Ito, S., Spencer, J., Itagaki, H.: Non-periodic inspection by Bayesian method I. Probab. Eng. Mech. 7(4), 191–204 (1992). https://www. sciencedirect.com/science/article/pii/026689209290023B
- Foerster, J., Farquhar, G., Afouras, T., Nardelli, N., Whiteson, S.: Counterfactual multi-agent policy gradients (2017). http://arxiv.org/abs/1705.08926, arXiv:1705.08926 [cs]
- Fulda, N., Ventura, D.: Predicting and preventing coordination problems in cooperative Q-learning systems. In: Proceedings of the 20th International Joint Conference on Artifical Intelligence, pp. 780–785. IJCAI 2007, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA (2007)
- Goldman, C.V., Zilberstein, S.: Decentralized control of cooperative systems: categorization and complexity analysis. J. Artif. Intell. Res. 22, 143–174 (2004). https://doi.org/10.1613/jair.1427
- Grall, A., Bérenguer, C., Dieulle, L.: A condition-based maintenance policy for stochastically deteriorating systems. Reliab. Eng. Syst. Saf. 76(2), 167–180 (2002). https://doi.org/10.1016/S0951-8320(01)00148-X
- Gronauer, S., Diepold, K.: Multi-agent deep reinforcement learning: a survey. Artif. Intell. Rev. 55(2), 895–943 (2022). https://doi.org/10.1007/s10462-021-09996-w
- Gupta, J.K., Egorov, M., Kochenderfer, M.: Cooperative multi-agent control using deep reinforcement learning. In: Sukthankar, G., Rodriguez-Aguilar, J.A. (eds.) AAMAS 2017. LNCS (LNAI), vol. 10642, pp. 66–83. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-71682-4_5
- Kaelbling, L.P., Littman, M.L., Cassandra, A.R.: Planning and acting in partially observable stochastic domains. Artif. Intell. 101(1–2), 99–134 (1998). https://doi. org/10.1016/S0004-3702(98)00023-X
- 20. Kapur, K.C., Pecht, M.: Reliability Engineering, 1st edn. Wiley, Hoboken (2014)
- Kochenderfer, M.J., Wheeler, T.A., Wray, K.H.: Algorithms for Decision Making. MIT Press, Cambridge (2022)
- Kuo, W., Zuo, M.: Optimal Reliability Modeling: Principles and Applications. Wiley, Hoboken (2003). https://catalogimages.wiley.com/images/db/pdf/ 047139761X.07.pdf
- Leroy, P., Morato, P.G., Pisane, J., Kolios, A., Ernst, D.: IMP-MARL: a suite of environments for large-scale infrastructure management planning via MARL (2023), http://arxiv.org/abs/2306.11551, arXiv:2306.11551 [cs, eess] version: 1
- Luque, J., Straub, D.: Risk-based optimal inspection strategies for structural systems using dynamic Bayesian networks. Struct. Saf. 76, 68–80(June 2017) (2019). https://doi.org/10.1016/j.strusafe.2018.08.002
- Lyu, X., Baisero, A., Xiao, Y., Daley, B., Amato, C.: On centralized critics in multiagent reinforcement learning. J. Artif. Intell. Res. 77, 295–354 (2023). https:// www.jair.org/index.php/jair/article/view/14386
- Madanat, S., Ben-Akiva, M.: Optimal inspection and repair policies for infrastructure facilities. Transp. Sci. 28(1), 55–62 (1994). https://doi.org/10.1287/trsc.28.1. 55,https://pubsonline.informs.org/doi/10.1287/trsc.28.1.55
- Matignon, L., Laurent, G.J., Le Fort-Piat, N.: Independent reinforcement learners in cooperative Markov games: a survey regarding coordination problems. Knowl. Eng. Rev. 27(1), 1–31 (2012). https://www.cambridge.org/core/product/identifier/S0269888912000057/type/journal_article
- 28. Morato, P.G., Andriotis, C.P., Papakonstantinou, K.G., Rigo, P.: Inference and dynamic decision-making for deteriorating systems with probabilistic dependencies through Bayesian networks and deep reinforcement learning. Reliab. Eng.

Syst. Saf. 235, 109144 (2023). https://www.sciencedirect.com/science/article/pii/S0951832023000595

- Morato, P.G., Papakonstantinou, K.G., Andriotis, C.P., Nielsen, J.S., Rigo, P.: Optimal inspection and maintenance planning for deteriorating structural components through dynamic Bayesian networks and Markov decision processes. Struct. Saf. 94(August 2021), 102140 (2022). https://doi.org/10.1016/j.strusafe. 2021.102140
- Oliehoek, F.A., Amato, C.: A Concise Introduction to Decentralized POMDPs. SpringerBriefs in Intelligent Systems, Springer, Cham (2016)
- Oliehoek, F.A., Spaan, M.T.J., Vlassis, N.: Optimal and approximate Q-value functions for decentralized POMDPs. J. Artif. Intell. Res. 32, 289–353 (2008). https:// doi.org/10.1613/jair.2447
- Papakonstantinou, K.G., Shinozuka, M.: Planning structural inspection and maintenance policies via dynamic programming and Markov processes. Part II: POMDP implementation. Reliab. Engi. Syst. Saf. 130, 214–224 (2014). https://doi.org/10. 1016/j.ress.2014.04.006
- Papakonstantinou, K.G., Andriotis, C.P., Shinozuka, M.: POMDP and MOMDP solutions for structural life-cycle cost minimization under partial and mixed observability. Struct. Infrastruct. Eng. 14(7), 869–882 (2018). https://doi.org/10.1080/ 15732479.2018.1439973
- Papoudakis, G., Christianos, F., Schäfer, L., Albrecht, S.V.: Benchmarking multiagent deep reinforcement learning algorithms in cooperative tasks (2021). http:// arxiv.org/abs/2006.07869, arXiv:2006.07869 [cs, stat]
- 35. Peng, B., et al.: FACMAC: factored multi-agent centralised policy gradients (2021). http://arxiv.org/abs/2003.06709, arXiv:2003.06709 [cs, stat]
- Rashid, T., Samvelyan, M., de Witt, C.S., Farquhar, G., Foerster, J., Whiteson, S.: QMIX: monotonic value function factorisation for deep multi-agent reinforcement learning (2018). http://arxiv.org/abs/1803.11485, arXiv:1803.11485 [cs, stat]
- Shani, G., Pineau, J., Kaplow, R.: A survey of point-based POMDP solvers. Auton. Agent. Multi-Agent Syst. 27, 1–51 (2013). https://doi.org/10.1007/s10458-012-9200-2
- Spaan, M.T.J., Vlassis, N.: Perseus: randomized point-based value iteration for POMDPs. J. Artif. Intell. Res. 24, 195–220 (2005)
- Sunehag, P., et al.: Value-decomposition networks for cooperative multi-agent learning (2017). http://arxiv.org/abs/1706.05296
- Terry, J.K., Grammel, N., Son, S., Black, B.: Parameter sharing for heterogeneous agents in multi-agent reinforcement learning (2022). http://arxiv.org/abs/2005. 13625
- 41. Wang, Z., et al.: Sample efficient actor-critic with experience replay (2017). http://arxiv.org/abs/1611.01224
- Wong, A., Bäck, T., Kononova, A.V., Plaat, A.: Deep multiagent reinforcement learning: challenges and directions (2022). http://arxiv.org/abs/2106.15691, arXiv:2106.15691 [cs]
- Yu, C., et al.: The surprising effectiveness of PPO in cooperative, multi-agent games (2022). http://arxiv.org/abs/2103.01955, arXiv:2103.01955 [cs]
- Zhu, B., Frangopol, D.M.: Risk-based approach for optimum maintenance of bridges under traffic and earthquake loads. J. Struct. Eng. 139(3), 422–434 (2013). https://doi.org/10.1061/(ASCE)ST.1943-541X.0000671
- Åström, K.J.: Optimal control of Markov processes with incomplete state information. J. Math. Anal. Appl. 10(1), 174–205 (1965). https://doi.org/10.1016/0022-247X(65)90154-X