

Response to Reviewers' Comments

Manuscript Title: Review of Deep Reinforcement Learning for Offshore Wind Farm Maintenance Planning

Journal: Wind Energy Science

Dear Reviewers,

We sincerely thank you for the time and effort invested in reviewing our manuscript. We carefully considered each comment and revised the paper to improve structural clarity. Below, we provide an overview of the main changes and a one-to-one response to each comment.

Sincerely,

Marco Borsotti, Rudy R. Negenborn, and Xiaoli Jiang

Overview of the main changes made in the revised manuscript

- Improved Figure 2 to explicitly separate the three algorithm families presented.
- Section 2-4: while maintaining a topological separation of methodological categories, we revised the presentation to discuss the reviewed approaches chronologically within each category, so that the evolution is clearer.
- Clarified the manuscript structure to separate methods from applications. **Sections 2–4** introduce DRL algorithm families (single- and multi-agent) and problem formulations first.

Section 5 introduces an application-focused synthesis centered on integrated offshore wind-specific knowledge, highlighting how aerodynamics/wake effects, weather and sea-state accessibility, logistics, PHM information, and economic considerations are incorporated across studies. Moreover, Table 5 now includes an additional column group indicating which studies incorporate each category of offshore-specific knowledge.

Section 6 provides a discussion of benefits and drawbacks; **Section 7** outlines future research directions, with emphasis on multi-level maintenance actions.

One-to-one response to Reviewer #1

We thank the Reviewer for the positive assessment. We corrected typographical issues and consistency problems to improve readability.

One-to-one response to Reviewer #3

Comment 1: The flowchart in Figure 2 may lead to misunderstanding of the three methods, as they might be mistakenly regarded as parallel. It is recommended to redraw the chart to clearly show the similarities and differences among the three approaches.

Response: We agree that the prior visualization could be interpreted as three parallel pipelines. In Figure 2 of the revised manuscript, we clarified the distinction between the three algorithm families (value-based, policy-gradient, and actor–critic). We also revised the flowchart caption text to emphasize the common components and to highlight where the update rules differ across the three families.

Comment 2: It is suggested that the development process of the methods be presented in chronological order.

Response: We thank the reviewer for this comment. To address this, we reorganized sections 2-4, where we introduce the main methodological families (single-agent DRL, multi-agent DRL, and problem formulations). Now, within each methodological category, the reviewed approaches are discussed in chronological order where multiple studies are referenced. This preserves a clear topological separation of the methodological concepts while making the progression of DRL approaches over time easier to follow.

Comment 3: The overall structure of the paper is supposed to introduce methods first and then applications. However, the methods section is described in an application-oriented manner and the details of the methods are not sufficiently clear. Please consider combining Sections 5 and 7 to summarize the development of methods in O&M, followed by a comprehensive comparison of their advantages and disadvantages, as well as an outlook on future development.

Response: We agree and revised the structure accordingly.

Sections 2–4 provide a methods-oriented presentation, introducing DRL algorithm families (single-agent and multi-agent) and problem formulations (MDP/PO-MDP, graph-based, hierarchical) before moving to applications.

We consolidated the application-oriented material into **Section 5** (Applications of DRL Approaches for Offshore Wind O&M), which emphasizes how offshore-specific knowledge is incorporated across studies, the body of literature reviewed is then summarized by the updated Table 5 and summary figures.

Section 6 (Discussion) now summarizes both benefits and drawbacks (such as the predominance of simulation-based validation), and we close the discussion with a clear focus on the remaining modelling gap: the limited consideration of multi-level maintenance actions.

Section 7 (Future directions) expands on this gap and outlines research avenues to represent more realistic maintenance action spaces.

Once again, we thank the Reviewers for the constructive feedback. We believe the revised manuscript addresses the comments regarding paper structure and clearer separation between methods and applications, while also providing a consolidated discussion and outlook.

Sincerely,

Marco Borsotti, Rudy R. Negenborn, and Xiaoli Jiang