Reply to the reviewers

We would like to thank all reviewers for their thorough analysis of our manuscript and their constructive comments. We have carefully considered each comment and provide below our answers. We will revise our manuscript accordingly and believe that this revision will address all reviewers' concerns.

Anonymous Referee #2

Report #2

Accepted as is.

Anonymous Referee #3

Report #1

Thank you for all the corrections, I find the paper clearer, but I still have some comments.

***Referee comment**: Line 11*
I am still uncomfortable with the "faster than real-time power forecasting" formulation.
I propose: "a key advantage of the data-driven approach is the high prediction speed compared to physics-based methods, such that it can be employed for applications where quasi-immediate power forecasting on different time scales is required".

Response of authors: Thank you for your suggestion. In the revised manuscript we will reformulate "faster than real-time power forecasting" in the abstract as follows: "enabling its use in applications that require forecasting multiple scenarios in real-time."

Further the document we specify the prediction speed more precisely ("a few milliseconds on PC") and we refer to the "reinforcement control setting where fast evaluations of many possible controller actions are required".

***Referee comment**: Line 235*
You introduce 2.2.3 before 2.2.2, please correct the order.

Response of authors: We will correct the order in the revised document.

***Referee comment**: Line 350*
I understand why you put the farm internal wake loss into sub-Section 2.2.
But I am still uncomfortable by defining the farm internal wake loss as a "machine learning model".
It uses the outputs of 2 different wake models, but it is not itself a ML model.
You do not have another "ML model" specific to the wake losses, you do not have a loss function including the wake losses, etc.
So I would put the farm internal wake loss outside the 2.2 ML models sub=Section.

Response of authors: We agree and will put the farm internal wake loss outside of the sub-section with all ML models.

***Referee comment**: Line 371*
I would put 2.2.6 at the beginning of the sub-Section 2.2.
So you would introduce 2.2 directly by presenting the full model.
You can also discuss the way you compute the farm internal wake loss.
And only then, you detail each model individually, 2.2.1 wind farm power model, etc.

Response of authors: Thank you for your suggestion. In the revised manuscript, we will present the overall pipeline indeed before detailing every ML model separately. In line with your previous comment, we will put the farm internal wake loss outside of the ML-models sub-section.

***Referee comment**: Line 390*
General comment: I feel that the results section is a little too big compared to the size of the paper. I would, if possible, reduce this section by focusing on the most important results and put less important results into an annex.

Response of authors: In the revised manuscript, we will shift the tables with the hyper-parameters of the ML models to an annex. While these details are valuable for reproducing our work, they are less important to the overall understanding and flow of the manuscript.